

Persistent homology of time-dependent functional networks constructed from coupled time series

Bernadette J. Stolz, Heather A. Harrington, and Mason A. Porter

Citation: *Chaos* **27**, 047410 (2017); doi: 10.1063/1.4978997

View online: <http://dx.doi.org/10.1063/1.4978997>

View Table of Contents: <http://aip.scitation.org/toc/cha/27/4>

Published by the [American Institute of Physics](#)

Welcome to a

Smarter Search 

PHYSICS
TODAY

with the redesigned
Physics Today Buyer's Guide

Find the tools you're looking for today!

Persistent homology of time-dependent functional networks constructed from coupled time series

Bernadette J. Stolz,¹ Heather A. Harrington,¹ and Mason A. Porter^{1,2,3}

¹Mathematical Institute, University of Oxford, Oxford OX2 6GG, United Kingdom

²Department of Mathematics, University of California Los Angeles, Los Angeles, California 90095, USA

³CABDyN Complexity Centre, University of Oxford, Oxford OX1 1HP, United Kingdom

(Received 22 July 2016; accepted 4 January 2017; published online 28 April 2017)

We use topological data analysis to study “functional networks” that we construct from time-series data from both experimental and synthetic sources. We use persistent homology with a weight rank clique filtration to gain insights into these functional networks, and we use persistence landscapes to interpret our results. Our first example uses time-series output from networks of coupled Kuramoto oscillators. Our second example consists of biological data in the form of functional magnetic resonance imaging data that were acquired from human subjects during a simple motor-learning task in which subjects were monitored for three days during a five-day period. With these examples, we demonstrate that (1) using persistent homology to study functional networks provides fascinating insights into their properties and (2) the position of the features in a filtration can sometimes play a more vital role than persistence in the interpretation of topological features, even though conventionally the latter is used to distinguish between signal and noise. We find that persistent homology can detect differences in synchronization patterns in our data sets over time, giving insight both on changes in community structure in the networks and on increased synchronization between brain regions that form loops in a functional network during motor learning. For the motor-learning data, persistence landscapes also reveal that on average the majority of changes in the network loops take place on the second of the three days of the learning process. *Published by AIP Publishing.* [<http://dx.doi.org/10.1063/1.4978997>]

Computational topology is a family of methods that are based on topological ideas (e.g., they often arise from algebraic topology) and give insights into topological invariants, such as connectedness or holes in high-dimensional data sets.^{1–3} Such efforts have come to be called topological data analysis, and a method known as persistent homology (PH) has been particularly helpful for understanding shapes and their persistence over multiple scales.^{4,5} Traditionally, PH has been applied to point-cloud data, though it has also been applied to networks in many applications, ranging from granular materials (see, e.g., Ref. 6) to functional brain networks.^{7,8} We employ these topological tools, which are designed to yield global, “higher-order” insights that go beyond revelations from pairwise connections (which are the norm in network science), in a study of functional networks constructed from both empirical and synthetic time-series data. We use persistence landscapes to show that the topological tools can (1) capture dynamics of networks constructed from the data and (2) identify mesoscale features that we relate to community structure^{9,10} in the associated functional networks. To help readers optimally understand these insights, we also present an intuitive introduction to PH and how to apply it to networks.

Analysis of neuronal networks is crucial for understanding the human brain.^{11–16} Every neuron consists of a cell body and one long axon, which is responsible for propagating signals to other cells.¹⁷ Neurons or (on a larger scale) different brain regions can be construed as nodes of a network, whose edges represent either structural or functional connections between those nodes. Examining neuronal data using a network-based approach allows one to use mathematical tools from subjects such as graph theory to better understand structural and functional aspects of neuronal interactions, identify key regions in the brain that are involved in physiological and pathological processes, and compare the structure of neuronal interactions with those of other complex systems. For example, data analysis using network theory has led to the insight that the brain has underlying modular structures, with small subunits that are able to carry out specific functions while minimally influencing other parts of the brain.^{11,14,18}

The standard methods from network theory are based on pairwise connections, which one can use to study the micro-scale, mesoscale, and macroscale structures.¹⁹ An alternative approach for studying networks²⁰ is to use methods from *computational topology*, which explicitly incorporates “higher-order” structures beyond pairwise connections and includes algorithmic methods for understanding topological invariants, such as connectedness, loops, or holes in high-dimensional data structures^{1–3} (see Section II B). Although one can also represent higher-order structures using formalisms such as hypergraphs²¹ (see, e.g., a recent paper²² by Bassett *et al.*), those other approaches may not be the most convenient means

I. INTRODUCTION

The human brain consists of many billions of neurons, whose major task is to receive, conduct, and transmit signals.

for optimally conveying information about the shape or scale of mesoscale structures in a network. Other recent work concerns clustering in networks using higher-order structures.²³

Methods from computational topology enable one to understand global low-dimensional structures in networks, and they have led to insights in an increasingly large number of applications⁵ in diverse topics, ranging from granular materials⁶ and contagions on networks²⁴ to path planning²⁵ and collective behavior in animals.²⁶ In particular, *persistent homology* (PH), a mathematical formalism to explore the persistence of topological structures in data sets, has become increasingly prominent in neuroscience in the last few years.^{7,8} Among other applications, it has been used to determine differences in brain networks of children with hyperactivity disorders and autism spectrum compared to normal situations,²⁷ study the effect of the psychoactive component of “magic mushrooms” (psilocybin mushrooms) on functional brain networks of humans,²⁸ analyze covariates that influence neural spike-train data,²⁹ and study structural and functional organization of neural microcircuits.³⁰ Other neuronal applications have included consideration of place cells in the hippocampus of rats during spatial navigation,^{31–33} analysis of mathematical models of transient hippocampal networks,³⁴ and a demonstration that topological features of networks of brain arteries in humans are correlated with their age.³⁵ We also note that PH is not the only topological method that has been used to study the human brain or time series. More than fifty years ago, for example, Zeeman³⁶ used tolerance spaces and Vietoris homology theory to study aspects of visual perception. In the 1990s, Muldoon *et al.*³⁷ developed a method to study the topology of manifolds that underlie time-series data.

In the present investigation, we use time-series data to construct so-called *functional networks*^{11,12,38,39} (but note that one can study coupled time series using a variety of different approaches^{40–42}). Functional brain networks consist of a set of nodes (e.g., brain regions) and a set of weighted edges between nodes, where the edge weights quantify the similarity of the associated time series according to a chosen measure. A functional network contrasts with a “structural network,” which refers to the underlying physical connections (e.g., anatomical connections) between nodes. For example, neurons are connected to each other in structural networks, but one can analyze the similarity in their firing patterns through functional networks. We use the term “functional network” in a more general way: by constructing a matrix of similarities between coupled time series using some measure (and enforcing the diagonal entries to be 0), one obtains a functional network whose weighted adjacency matrix (sometimes also called an “association matrix”) $\tilde{A} = (\tilde{a}_{ij})_{i,j=1}^N$ has elements that indicate the similarity between the time series of entities i and j . Studying functional networks is common in neuroscience, and they are also used in a wealth of other applications (e.g., finance,⁴³ voting among legislators,⁴⁴ and climate⁴⁵). Importantly, the time series can come either from empirical data or from the output of a dynamical system (or stochastic process), and the latter can be helpful for validating methods for network analysis.⁴⁶ In our paper, we will consider time series either from coupled oscillators (i.e., as the output of a

dynamical system) or from a set of spatially distinct brain regions defined by a fixed anatomical atlas. In the context of functional brain networks, the adjacency-matrix element a_{ij} arises as a measure of “functional connectivity” (i.e., behavioral similarity) between the time series for nodes (i.e., brain regions) i and j . There are many different ways to measure similarity of time series,^{12,47,48} and that can be a major issue when it comes to interpreting results. Comparing the networks that arise from different similarity measures is beyond the scope of our work, so we will simply use two common measures (pairwise synchrony and wavelet coherence) of time-series similarity. However, the methods that we employ can be applied to functional networks that are constructed using any measure of similarity between time series.

In many studies based on experimental data, functional networks are used to construct binary graphs (i.e., unweighted graphs).¹² To do this, one typically applies a global threshold $\xi \in \mathbb{R}^+$ to a weighted adjacency matrix to obtain a binary adjacency matrix $A = (a_{ij})_{i,j=1}^N$ associated with an unweighted graph. The adjacency-matrix elements are then

$$a_{ij} = \begin{cases} 1, & \text{if } \tilde{a}_{ij} \geq \xi, \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

The choice of threshold has a strong influence on the resulting matrix, and it thereby exerts a major influence on the structure of the associated graph.¹² Some approaches to address this issue include determining a single “optimal” threshold, thresholding the weighted adjacency matrix at different values,^{49,50} examining the network properties as a function of threshold, or not thresholding at all and considering the weighted adjacency matrix itself.^{12,14} (One can also threshold a weighted adjacency matrix by setting sufficiently small entries to 0 but keeping the values of the other entries.) If one is thresholding and binarizing data, there is no guarantee that there exists an interval of thresholds that yield networks with qualitatively similar properties, and arbitrarily throwing away data can be problematic even when such intervals do exist. For example, parameters such as graph size (i.e., number of nodes) need to be taken into account when interpreting results on thresholded, binarized networks.⁵¹ An advantage of using persistent homology is that one can examine a graph “filtration” (see Section II C) generated by multiple—ideally all—possible global thresholds and systematically analyze the persistence of topological features across these thresholds. Such a filtration can also be created using decreasing local thresholds.

In our topological analysis, we focus on “loops” in a network. For our purposes, a loop in a graph is a set of at least four edges that are connected in a way that forms a topological circle.⁵² Loops are thus 1-dimensional topological features. We choose to focus on loops rather than features with dimension 0, which correspond to connected components of a graph, are topologically simpler, and can be studied using many other approaches (e.g., through the number of 0 elements in the spectrum of the combinatorial graph Laplacian²¹). It has been demonstrated in other applications




FIG. 1. Examples of (from left to right) a 0-simplex, a 1-simplex, a 2-simplex, and a 3-simplex. [Figure adapted from Ref. 3.]

(e.g., contagions on networks²⁴) that loops are important topological features of graphs, and a recent study⁵³ demonstrated the importance of loops (and related higher-dimensional objects) in structural neuronal networks. Structural and functional neuronal networks are related and share some network features,¹¹ so we expect loops to provide interesting insights.

The remainder of our paper is organized as follows. In Section II, we give a brief and intuitive introduction to persistent homology, the weight rank clique filtration (WRCF), and persistence landscapes. In Section III, we introduce our first example, the Kuramoto model of nonlinearly coupled oscillators; and we present results from our application of persistent homology to time-series data produced by coupled Kuramoto oscillators. In Section IV, we introduce and analyze our second example, which consists of time-series functional magnetic resonance imaging (fMRI) data from experiments of humans performing a simple motor-learning task. We present our conclusions in Section V, and we provide a mathematical introduction to persistent homology in the [Supplementary Material](#).

II. PERSISTENT HOMOLOGY

Persistent homology (PH)^{1–3} is a method from computational topology that quantifies global topological structures (e.g., connectedness and holes) in high-dimensional data. One can think of PH as looking at the “shape” of data in a given dimension using a set of different lenses. Each lens conveys topological features inside data at a different resolution. One then construes structures that persist over a range of different lenses to represent one or more significant features of the data. Structures that are observed only through a

small number of lenses are commonly construed as noise,^{54,55} especially in settings in which the data are sampled from a manifold. For empirical data, the relationship between low persistence of a feature and whether it constitutes noise in a data set rather than signal has not yet been verified statistically, but we will illustrate a situation in which some short-lived structures represent important features and possibly genuine geometrical (not just topological) features of data in Sections III and IV.

In this section, we provide an intuitive introduction to the mathematical concepts behind PH. In the [Supplementary Material](#), we give a mathematically rigorous introduction (including precise definitions).

A. Simplicial complexes

One can study the properties of a topological space^{56,57} by partitioning it into smaller and topologically simpler pieces, which when reassembled include the same aggregate topological information as the original space. The most trivial topological space $X = \{\emptyset, x\}$ consists of the empty set \emptyset and a single point x . If we want to simplify the description of the topological properties of X , we would simply choose a single node to represent it. However, a node or even a collection of nodes does not allow one to capture the topological properties of more complicated spaces, such as a 2-sphere or the surface of the earth. In such cases, one needs a simple object that carries the information that the space is connected but also encloses a hole. For example, one could use a tetrahedron, which is an example of a mathematical object called a *simplex*.

The building blocks that one uses to approximate topological spaces are called k -simplices, where the parameter k indicates the dimension of the simplex. Every k -simplex includes $k + 1$ independent nodes: a point is a 0-simplex, an edge is a 1-simplex, a triangle is a 2-simplex, and a tetrahedron is a 3-simplex (see Fig. 1). Observe that the lower-dimensional simplices are contained in the higher-dimensional simplices. This allows one to build higher-dimensional simplices using lower-dimensional ones. The lower-dimensional simplices form the so-called *faces* of the associated higher-dimensional objects.

One combines different simplices into a *simplicial complex* to capture different aspects of a topological space. For every simplex that is part of a simplicial complex, we




FIG. 2. Panels (a)–(c) give examples of simplicial complexes, and panel (d) gives an example of an object that is not a simplicial complex. The blue color indicates a 2-simplex. Example (a) illustrates that simplicial complexes are not necessarily also simplices. The three edges do not form a 2-simplex; instead, they form a simplicial complex that consists of 1-simplices. In examples (b) and (c), all 1-simplices and 2-simplices are connected to each other via 0-simplices (i.e., the intersections of all present simplices are 0-simplices). Example (d) is a collection of simplices that violates the definition of a simplicial complex, because the intersection between the two triangles does not consist of a complete edge that is shared by both simplices (as it includes only approximately 75% of the base edge of the upper triangle). Note that any combination of the three simplicial complexes (a), (b), and (c) is also a simplicial complex.

demand that all of its faces are also contained in the simplicial complex. Additionally, two simplices that are part of a simplicial complex are allowed to intersect only in common faces. In Fig. 2, we show several examples of simplicial complexes and one example that is not a simplicial complex.

We take the *dimension* of a simplicial complex to be the dimension of its highest-dimensional simplex. One can use simplicial complexes to represent topological spaces if and only if there exists a continuous deformation that can stretch and bend the simplicial complex into the topological space, and only then are topological properties of the topological space preserved by the simplicial complex.

B. Homology and Betti numbers

If one is interested in the nature of a simplicial complex of dimension k , one can either consider the full complex, which can be very large, or one can examine different subsets of simplices that are contained in the complex. For example, the set of all 1-simplices consists of a collection of edges, some of which may be connected or even form a loop. However, one can consider a range of different topological features—e.g., in some cases, a collection of edges surrounding a hole or void could be more interesting than individual edges—and one typically seeks features that are invariant if one squeezes or stretches the edges. *Homology* is a formal way to quantitatively detect topological invariants in a given dimension to give insight into the nature of a topological space. By using homology, one can, for example, distinguish a 2-sphere from a torus. For a simplicial complex of dimension k , one can define a vector space known as the p th *homology group* for every dimension $p \in \{0, \dots, k\}$. In dimension 1, for example, the elements of the homology group are called “loops.” The elements of the homology group can be divided into different *homology classes*, which each represent a hole in the topological space. For instance, in dimension 1, loops in the same homology class all surround the same 1-dimensional hole. We give an example of two loops that surround the same hole in Fig. 3. The homology classes yield a family of vector spaces, whose dimensions are called *Betti numbers*, associated to a simplicial complex. One can interpret the first three Betti numbers, β_0 , β_1 , and β_2 , to represent, respectively, the number of connected components, the number of 1-dimensional holes, and the number of 2-dimensional holes in a simplicial complex. As we pointed out in Section I, we focus on the number of loops (i.e., on β_1) in our network analysis rather than on β_0 , which corresponds to the number of connected components in a graph. One can study connected components in graphs using many other approaches, such as by calculating the number of 0 eigenvalues in the spectrum of the combinatorial graph Laplacian.²¹

C. Filtrations

Although homology gives information about a single simplicial complex, it typically is more relevant to study the topological features across sequences (called *filtrations*) of simplicial complexes. A filtration^{2,54,55} of a simplicial complex Σ is a sequence of embedded simplicial complexes,




FIG. 3. Example of a loop in a simplicial complex. The green and the blue loops both surround the same 1-dimensional hole and are therefore considered to be representatives of the same homology class.

$$\emptyset = \Sigma_0 \subseteq \Sigma_1 \subseteq \Sigma_2 \subseteq \dots \subseteq \Sigma_k = \Sigma, \quad (2)$$

starting with the empty complex and ending with the entire simplicial complex. One can use homology to study topological features (e.g., loops) in every step of the filtration and determine how persistent they are with respect to a given filtration. A topological feature h is *born* at Σ_m if the homology group of Σ_m is the first homology group to include the feature. Similarly, a topological feature *dies* in Σ_n if it is present in the homology group of Σ_{n-1} but not in that of Σ_n . One then defines the *persistence* p of the topological feature as

$$p = n - m.$$

Persistence was first used as a measure to rank topological features by their lifetime¹ in a filtration in \mathbb{R}^3 .

There are many ways to define simplicial complexes and filtrations on weighted graphs, and the choice of filtration tends to be motivated either by the type of questions to be answered or by consideration of computation time.

1. Weight rank clique filtration

Although we focus on network data, we note that PH has been applied much more often to data in the form of point clouds.^{54,55} The simplest way to create a sequence of embedded graphs (e.g., a filtration) from a weighted network is to filter by weights.⁵⁸ To do this, one creates a sequence of embedded (binary) graphs by ranking all edge weights ν_i in descending order. In filtration step t , one retains an edge if and only if its weight is at least ν_t . To construct the filtration, one repeats this procedure until the graph is complete in the

last step. Using this method, one is able to study 0-simplices (i.e., nodes) and 1-simplices (i.e., edges). The weight rank clique filtration (WRCF),⁵⁹ which we will use in our analysis and which has been applied previously for examining weighted neuronal networks,^{28,33,59} extends this definition to include higher-dimensional simplices. One constructs a WRCF as follows:

1. Define filtration step 0 as the set of all nodes.
2. Rank all edge weights $\{\nu_1, \dots, \nu_{\text{end}}\}$, with $\nu_1 = \nu_{\text{max}}$ and $\nu_{\text{end}} = \nu_{\text{min}}$. (We will use τ to denote the number of distinct weights in a graph.)
3. In filtration step t , threshold the graph at weight ν_t to create a binary graph.
4. Find all maximal c -cliques for $c \in \mathbb{N}$, and define them to be c -simplices.

This is a valid simplicial complex: every $(c + 1)$ -clique in the graph guarantees the existence of a c -face on that clique, because cliques are closed under both intersection and taking subsets. Consequently, they satisfy the requirements for a simplicial complex. This type of simplicial complex on a graph is called a *clique complex*.

One can visualize the persistence of homology classes of a filtration of a simplicial complex using *barcodes*.⁵⁴ A barcode for a given dimension is a collection $\{b_l, d_l\}_{l=1}^m$ of interval endpoints, where every interval (b_l, d_l) represents a topological feature l of the given dimension (examples of such features include connected components and loops), b_l denotes the birth time of feature l with respect to the filtration step, and d_l denotes its death time. The length $d_l - b_l$ of the bar measures the persistence of the feature. In Fig. 4, we show an example of a WRCF and its corresponding barcode.

D. Persistence landscapes

As an alternative topological summary to barcodes, one can use *persistence landscapes*,^{60,61} which consist of

piecewise-linear functions in a separable Banach space. For a given barcode interval (b, d) , one defines the function

$$f_{(b,d)} = \begin{cases} 0, & \text{if } x \notin (b, d), \\ x - b, & \text{if } x \in \left(b, \frac{b+d}{2}\right], \\ -x + d, & \text{if } x \in \left(\frac{b+d}{2}, d\right). \end{cases} \quad (3)$$

For a barcode $\{b_l, d_l\}_{l=1}^m$ and $q \geq 0$, the q -th *persistence landscape* is given by the set of functions

$$\lambda_q : \mathbb{R} \rightarrow \mathbb{R}, \\ \lambda_q(x) = q\text{th-largest value of } \{f_{(b_l, d_l)}(x)\}_{l=1}^m. \quad (4)$$

If the q th-largest value does not exist, then $\lambda_q(x) = 0$. One can think of the 0th persistence landscape as being the outline of the collection of peaks created by the images of the collection of functions f associated to a barcode. To obtain the 1st persistence landscape, one peels away this topmost “layer” of peaks and then considers the outline of the remaining collection of peaks. This gives the 1st persistence landscape, and one continues in this manner to obtain subsequent persistence landscapes. The *persistence landscape* λ of the barcode $\{b_l, d_l\}_{l=1}^m$ is then defined as the sequence $\{\lambda_q\}$ of functions λ_q (where λ_q is the q th layer of the landscape).

Even though persistence landscapes visualize the same information as barcodes for an individual filtration and one can construct a bijective correspondence between the two objects, the former have distinct advantages over the latter. For example, one can calculate a unique “average landscape” for a set of persistence landscapes from several filtrations by taking the mean over the function values for every landscape layer (see Fig. 5). This is not possible for barcodes, as they are not elements of a Banach space. For an average landscape, it is thus not possible to find a corresponding average

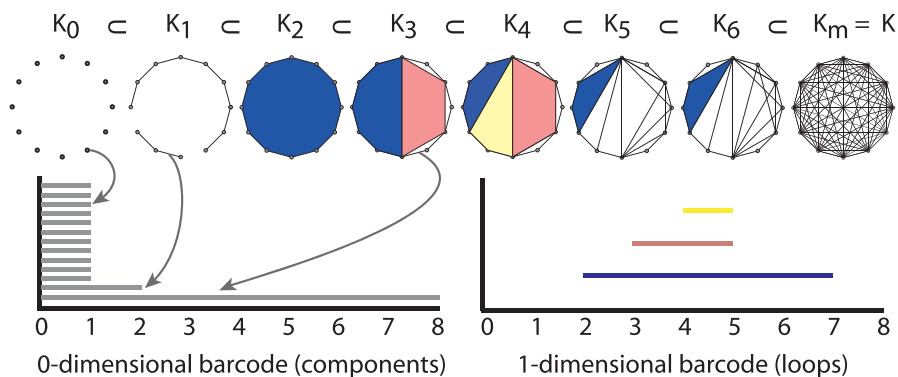


FIG. 4. Example of a weight rank clique filtration (WRCF) and the corresponding 0-dimensional and 1-dimensional barcodes. The barcode of dimension 0 indicates the connected components in every filtration step. When two components merge into one connected component, one of the bars that represent the original components dies in the barcode; the other continues to the next filtration step and now represents the newly-formed component. In filtration step 0, every node is a separate component, resulting in 12 bars in the barcode. The nodes are joined to become two components in filtration step 1, and they then become a single component in step 2. In dimension 1, we observe that as more edges are added to the filtration, the loop surrounding the blue hole born in filtration step 2 is divided first into two holes and subsequently into three holes before it is completely covered by 2-simplices and dies in filtration step 7. The colors of the bars indicate which loop they represent.

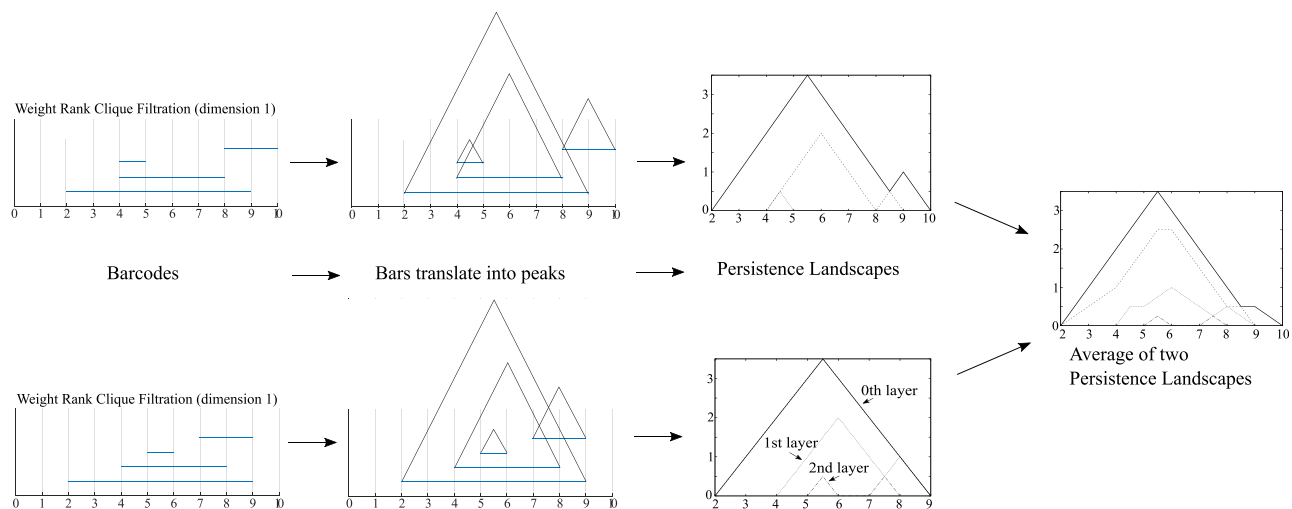


FIG. 5. Visualization of the relationship between barcodes and an average persistence landscape. To obtain a landscape from a barcode, one replaces every bar of the barcode by a peak, whose height is proportional to the persistence of the bar. In the landscape, we translate all peaks so that they touch the horizontal axis. The persistence landscape consists of different layers, where the q th layer corresponds to the q th-largest function values in the collection of peak functions. One creates an average of two landscapes by taking the mean over the function values in every layer.

barcode. We show a schematic illustration of how to obtain an average persistence landscape in Fig. 5.

One can also define an L^p distance between two landscapes (either individual ones or average ones) and thereby use a variety of statistical tools.⁶⁰ This allows one to compare multiple groups of landscapes (and thereby also compare groups of barcodes) by calculating a measure of pairwise similarity between landscapes. Persistence landscapes have been used to study conformational changes in protein binding sites,⁶² phase separation in binary metal alloys,⁶³ and music audio signals.⁶⁴

E. Computational tools

For our PH calculations, we use MATLAB code that we construct using Javaplex,^{65,66} a software package for persistent homology. For the WRCFs, we also use a maximal clique-finding algorithm from the Mathworks library⁶⁷ that is based on the Bron–Kerbosch algorithm, which is the most efficient algorithm known for this problem. For statistical analysis and interpretation of our barcodes, we apply the Persistence Landscapes Toolbox.⁶¹

III. EXAMPLE I: COUPLED KURAMOTO OSCILLATORS

A. The Kuramoto model

The Kuramoto model^{68–72} is a well-studied model for a set of coupled phase oscillators whose natural frequencies are drawn from a prescribed distribution. The model was developed in the 1970s to understand synchronization in a large system of oscillators. It has subsequently been used as a toy model by many neuroscientists (as well as scholars in many other areas), as some of the characteristics of its synchronization patterns resemble some of the ones in neuronal communities.^{73–76} The Kuramoto model and its generalizations have

also been applied to numerous other applications in chemistry, biology, and other disciplines.^{70,71,77}

When all oscillators are coupled to each other, the Kuramoto model is most commonly written as^{69,71}

$$\frac{d\theta_i}{dt} = \omega_i + \frac{K}{N} \sum_{j=1}^N \sin(\theta_j - \theta_i), \quad i \in \{1, \dots, N\}, \quad (5)$$

where θ_i denotes the phase of oscillator i , the parameter ω_i is its natural frequency, $K \geq 0$ parametrizes the coupling strength between different oscillators, and N is the number of oscillators in the model. The normalization factor $\frac{1}{N}$ ensures that the right-hand side of Eq. (5) is bounded as $N \rightarrow \infty$. The distribution from which the frequencies ω_i are drawn is usually assumed to be unimodal and symmetric about its mean frequency, which can be set to 0 due to the rotational symmetry of the model (because Eq. (5) is invariant under translation of θ_i). The parameter ω_i then denotes the deviation from the mean frequency.

We also adapt Eq. (5) to examine a network of N oscillators with uniform coupling between the oscillators.^{40,46,70,71,78} We consider the following generalized version of Eq. (5):

$$\frac{d\theta_i}{dt} = \omega_i + \sum_{j=1}^N \kappa A_{ij} \sin(\theta_j - \theta_i), \quad i \in \{1, \dots, N\}, \quad (6)$$

where $\kappa \geq 0$ denotes the normalized coupling strength and the entries of the coupling matrix $A = (A_{ij})_{i,j=1}^N$ indicate whether oscillators i and j are coupled. That is, A is an unweighted adjacency matrix, and $A_{ij} = 1$ for coupled oscillators and $A_{ij} = 0$ for uncoupled oscillators. The coupling matrix A thereby imposes a “structural network” between the oscillators. One can further generalize Eq. (6) by using

heterogeneous coupling strengths κ_{ij} or by considering functions other than sine on the right-hand side.

We divide the oscillators into 8 separate communities⁷⁹ of 16 distinct oscillators each, and we suppose that every oscillator has exactly 14 connections, 13 of which are with oscillators in the same community and 1 of which is to an oscillator outside the community. As in Bassett *et al.*,⁴⁶ we choose a coupling strength of $\kappa = 0.2$, consider a network with $N = 128$ oscillators, and suppose that the i th natural frequency is $\omega_i \sim \mathcal{N}(0, 1)$. (That is, we draw natural frequencies from a Gaussian distribution with mean 0 and standard deviation 1.) However, our network architecture differs somewhat from that in Bassett *et al.*,⁴⁶ where every oscillator had at least 13 connections inside its community and at least 1 connection outside its community.

We simulate the Kuramoto model [Eq. (6)] using the Runge–Kutta MATLAB solver ODE45 (with an integration time interval of $[0, T_{\max}]$, where $T_{\max} = 10$).⁸⁰ We observe the system for $M = 500$ time steps in total (including the initial time step) and obtain time series $\mathcal{T}_{(i)} = (\theta_i(t_0), \dots, \theta_i(t_{499}))$ as the output of the model for each oscillator θ_i . Kuramoto oscillators with a similar imposed community structure were demonstrated previously to initially synchronize rapidly within their communities, followed by a phase of global synchronization in an entire network.⁴⁶ (There have also been other studies of community structure via synchronization of Kuramoto oscillators.^{78,81}) To study the dynamics of the coupled Kuramoto oscillators, we follow the approach of Bassett *et al.*⁴⁶ and partition the time series into two time regimes, which we denote by $\hat{k} = 1$ and $\hat{k} = 2$. In our example, these time regimes each consist of 250 time steps. (The authors of Ref. 46 also split their time series into two equal parts, but their time series consist of 100 time steps in total rather than 500.)

To quantify the pairwise synchrony of two oscillators i and j , we use the local measure^{46,78}

$$\phi_{ij}^{\hat{k}} = \langle |\cos(\mathcal{T}_{(i)}^{\hat{k}} - \mathcal{T}_{(j)}^{\hat{k}})| \rangle, \quad (7)$$

where the angular brackets indicate that we take a mean over 20 simulations. We use the absolute value both to facilitate comparison with Arenas *et al.*⁷⁸ and Bassett *et al.*⁴⁶ (by making the same choice that they made) and to avoid negative values, which can complicate interpretation and pose other difficulties in network analysis.^{43,47,82}

In each simulation, we choose the initial values for the phases θ_i from a uniform distribution on $(0, 2\pi)$ and draw the natural frequencies ω_i from $\mathcal{N}(0, 1)$. We apply the same underlying coupling matrix $A = (A_{ij})_{i,j=1}^N$ for all 20 simulations and then use the values ϕ_{ij} to define the edge weights in the fully connected, weighted network of Kuramoto oscillators for each time regime. We also study a network based on one full time regime that consists of 500 time steps. In analogy to neuronal networks, we call these networks “functional networks.” In Fig. 6, we illustrate our pipeline for creating a functional network from the output of a simulation of the Kuramoto model.

B. Null models for the Kuramoto data

To assess whether our observations illustrate meaningful dynamics of the Kuramoto model or whether they can be explained by a random process, we consider two different null models based on the time-series output. In the first null model, which we call the “simple null model,” we independently reassign the order of the time series of each oscillator according to a uniform distribution before computing the similarity measure with Eq. (7). The second null model, which we call the “Fourier null model,” is based on creating surrogate data using a discrete Fourier transformation. This approach⁸³ has the advantage of preserving not only the mean and the variance of the original time series but also the linear autocorrelations and cross correlations between the different time series.

To construct the Fourier null model, we start by taking the discrete Fourier transform

$$\hat{\mathcal{T}}_n = \frac{1}{\sqrt{\mu}} \sum_{m=0}^{\mu-1} \mathcal{T}_m e^{\frac{2\pi i m n}{\mu}} \quad (8)$$

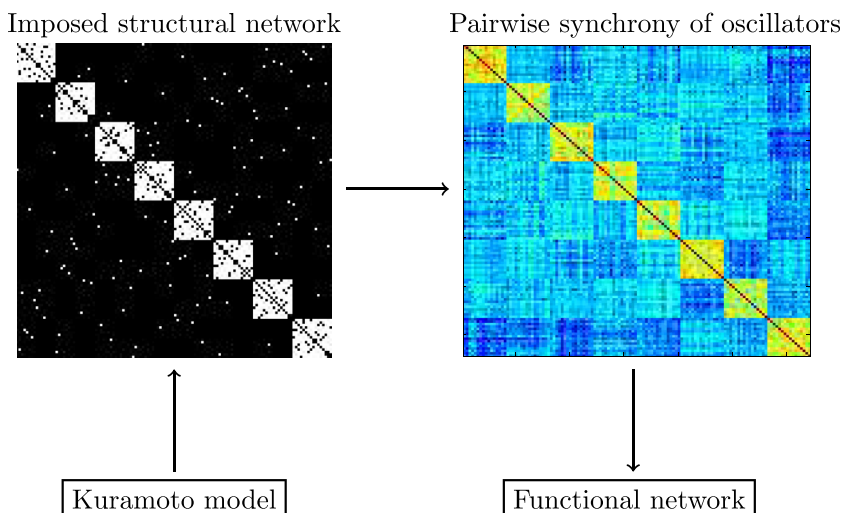


FIG. 6. We construct a structural network for coupled Kuramoto oscillators by grouping the oscillators into 8 separate communities. Oscillators are coupled predominantly to other oscillators in their community, and they are coupled only very sparsely to oscillators outside their community. We use the time-series output of a simulation of the Kuramoto model to create a functional network based on the similarity of the time series of individual oscillators. We use the measure of similarity in Eq. (7).

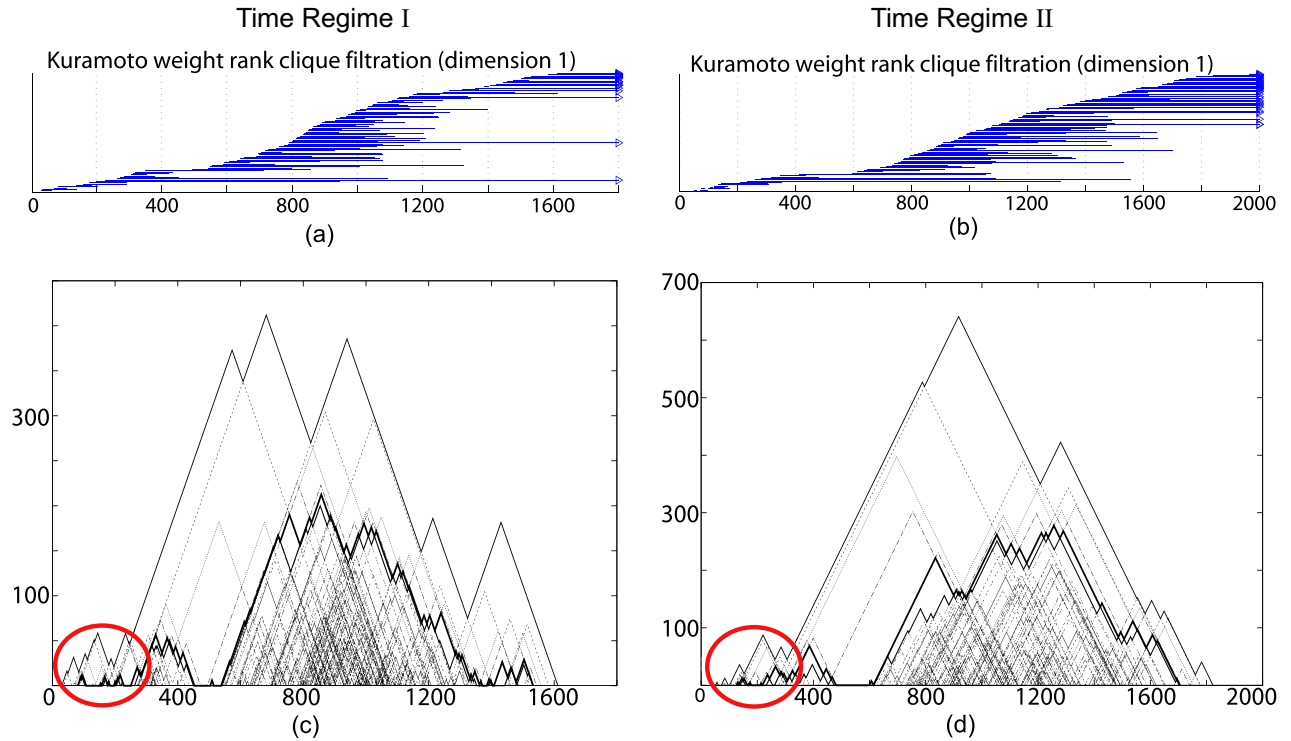


FIG. 7. Dimension-1 barcodes and persistence landscapes for the WRCF for the two time regimes (time steps 1–250 and time steps 251–500) of time-series output of the Kuramoto model. The horizontal axis represents the filtration steps in both the barcodes and the landscapes. The vertical axis in the persistence landscape captures the persistence of the features in the barcode. In the first row, we show the barcodes for dimension 1. In the second row, we show persistence landscapes (although we ignore infinitely-persisting bars in the barcodes). The short peaks early in the filtration in the persistence landscapes that are indicated by the red ellipses represent loops formed within communities. The most prominent difference between the two landscapes is the occurrence of high peaks in the second time regime; these peaks correspond to persistent inter-community loops in the network that are formed between communities.

of a time-series vector \mathcal{T} (with components \mathcal{T}_m) of length μ . In our case, $\mu = 250$ or $\mu = 500$, depending on whether we are examining two different time regimes or just one. We then construct surrogate data by multiplying the Fourier transform $\hat{\mathcal{T}}_n$ by phases a_n chosen uniformly at random from the interval $(0, 2\pi)$, aside from the constraint that they must satisfy the following symmetry property: for every $n \leq \mu$, there exists \tilde{n} such that $a_n = -a_{\tilde{n}}$. This symmetry ensures that the inverse Fourier transform yields real values. The surrogate data $\sigma = (\sigma_1, \dots, \sigma_\mu)$ are thus given by

$$\sigma_m = \frac{1}{\sqrt{\mu}} \sum_{n=0}^{\mu-1} e^{ia_n} \hat{\mathcal{T}}_n e^{-\frac{2\pi i n m}{\mu}}. \quad (9)$$

Both the simple null model and the Fourier null model were used previously on time-series output of coupled Kuramoto oscillators, and they exhibit different dynamics from those of the oscillators.^{22,46}

C. Persistent homology applied to the Kuramoto model and null models

We apply the WRCF to functional networks created from the output of two time regimes of the Kuramoto model, one time regime (i.e., the full time-series output) for the Kuramoto model, the simple null model, and the Fourier null

model. We run the filtrations up to filtration step 1800 for the first time regime and up to 2000 for the second; we go up to filtration step 1100 for cases in which we only consider one time regime. The total number of edges in the network, and thus the total number of possible filtration steps, is 8128. The number of filtration steps thereby corresponds to respective edge densities of 0.22, 0.25, and 0.14 for the three examples above; in each case, this amounts to a threshold that is approximately in the middle of the range of the edge-weight values. The Masters thesis of Stolz,⁸⁴ which is a precursor to the present paper, also applied PH to networks created from the Kuramoto model, and such an example was subsequently also studied using Betti curves by other authors.⁸⁵

As we described in Section I, we focus our analysis on topological features in dimension 1, so we examine loops in the network. In the first row of Fig. 7, we show the 1-dimensional barcodes for the networks constructed from time regime 1 (i.e., the first 250 time steps of the dynamics) and time regime 2 (i.e., time steps 251–500 of the dynamics) for the WRCF of the Kuramoto model. The barcode for each time regime includes several very short-lived bars between filtration steps 50 and 300. For the second time regime, we find more short bars for a longer filtration range at the beginning of the barcode. The loops that correspond to these short bars are all formed within the strongly synchronized communities. In fact, in time regime 1, the first

44 bars in the barcodes represent intra-community loops; in time regime 2, only 2 of the first 28 bars represent intra-community loops. As strong intra-community edges are added to the simplicial complexes, they start to cover the loops with triangles (i.e., 2-simplices), and the loops disappear from the filtration.

In the second row of Fig. 7, we show the persistence landscapes that we construct from the 1-dimensional barcodes. We ignore infinitely-persisting bars in the barcode. (We also studied persistence landscapes that include the infinite bars as features with a death time that corresponds to the maximum filtration value but did not obtain any additional insights that way.)

As expected, the landscapes have a group of small peaks early in the filtration for both time regimes. This feature occurs in a longer filtration range in the second time regime before more-persistent loops appear. In the second time regime, some of the peaks that occur early in the filtration appear to almost double their heights to values of about 100. In contrast, in the first time regime, peaks at a similar location are about half as high (i.e., they are less persistent).

The persistence landscapes reveal more persistent loops in the second time regime (i.e., between time steps 251 and 500) than in the first (i.e., between time steps 1 and 250), and the second time regime also appears to reveal a clearer separation between a group of very early short peaks and a group of medium-sized peaks towards the end of the filtration. For this second group of medium-sized peaks, we observe a larger absolute increase in persistence in the second time regime than for the shorter peaks early in the filtration. These observations reflect the dynamics of the two time regimes in the Kuramoto model.⁴⁶ In time regime 1, there is strong synchronization within the communities, and such dynamics are reflected by the appearance of short-lived intra-community loops (corresponding to the short peaks in the persistence landscapes) early in the filtration. In the second time regime, the amount of global synchronization is more prominent than in the first time regime. Moreover, in addition to intra-community loops, some of the peaks early in the filtration now represent inter-community loops, which are more persistent than the loops within communities. Additionally, as some of the peaks that correspond to inter-community loops have shifted to earlier parts of the filtration, there is an increase in the gap between the initial group of peaks and the group of medium-sized peaks at the end of the filtration. In general, we observe increased persistence of the peaks in the landscapes due to the stronger synchronization between the communities. These observations are much easier to visualize using persistence landscapes than using barcodes.

We calculate pairwise L^2 distances between all dimension-1 persistence landscapes, and we note that L^2 distance has been used previously to compare persistence landscapes in an application to protein binding.⁶² The L^2 distance between the two time regimes is 27078. (Here, and in subsequent discussions, we round the L^2 distances to the nearest integer.) Given the length of the support of the landscapes and the function values that the landscapes attain, this is a

large distance, which captures the aforementioned visible differences between the landscapes. The L^2 distance is unable to capture the fact that the peaks that appear early in the filtration in the first time regime correspond to loops between nodes within one community, whereas they correspond to loops that form between nodes of different communities in the second time regime. Consequently, this feature does not contribute to the value of the distance.

We also compare the Kuramoto model with the two null models that we discussed in Section III B. To do this, we construct a functional network by considering a single time regime that consists of 500 time steps. In Fig. 8, we show the weighted adjacency matrices of the three functional networks, and we also show their corresponding persistence landscapes based on WRCFs of the functional networks. One can observe clearly that there is stronger intra-community synchronization for the Kuramoto time series than for the null models, as there is a very distinct group of short peaks early in the filtration (which, as we discussed above, is also the case for the Kuramoto model when performing separate calculations in the two time regimes).

Again, the corresponding loops occur within communities. The peaks in the Kuramoto landscape appear to be separated from a second group of short peaks further along in the filtration. Between the two groups of peaks, there are two strikingly higher peaks that correspond to persistent loops that appear to be formed by connections between different communities. For both null models, we also observe groups of short peaks early in the filtration, but these are less persistent and separated less clearly from other peaks than for the Kuramoto model. Indeed, we do not see any separation at all for the Fourier null model, which exhibits a much weaker intra-community synchronization than the simple null model. Moreover, the persistence landscape for the Fourier null model appears to be “noisier,” as the majority of the peaks in the landscape have similar persistences and appear in similar regions of the filtration.

The peaks in the landscapes of the null models appear to have a very different distribution along the filtration than is the case for the Kuramoto model. They also possess more medium-sized and long persisting features than we observe in the Kuramoto data. These features occur in parts of the filtration in which the Kuramoto data have a smaller number of peaks. They consist of inter-community loops and are a symptom of the weaker intra-community and stronger inter-community synchronization. The null models thus appear to have more topological features in the form of loops than is the case for the Kuramoto data. This is consistent with previous examinations of null models in other studies.^{33,53,59} The fact that there are fewer persistent loops in the Kuramoto model than in the null models implies that there are more high-dimensional simplices (e.g., triangles and tetrahedra) in the corresponding network than in the networks constructed from the null models.

To distinguish between the three landscapes, we calculate the L^2 distances between them. The L^2 distance between the Kuramoto landscape and the Fourier null-model landscape is 13540; the L^2 distance between the two null-model

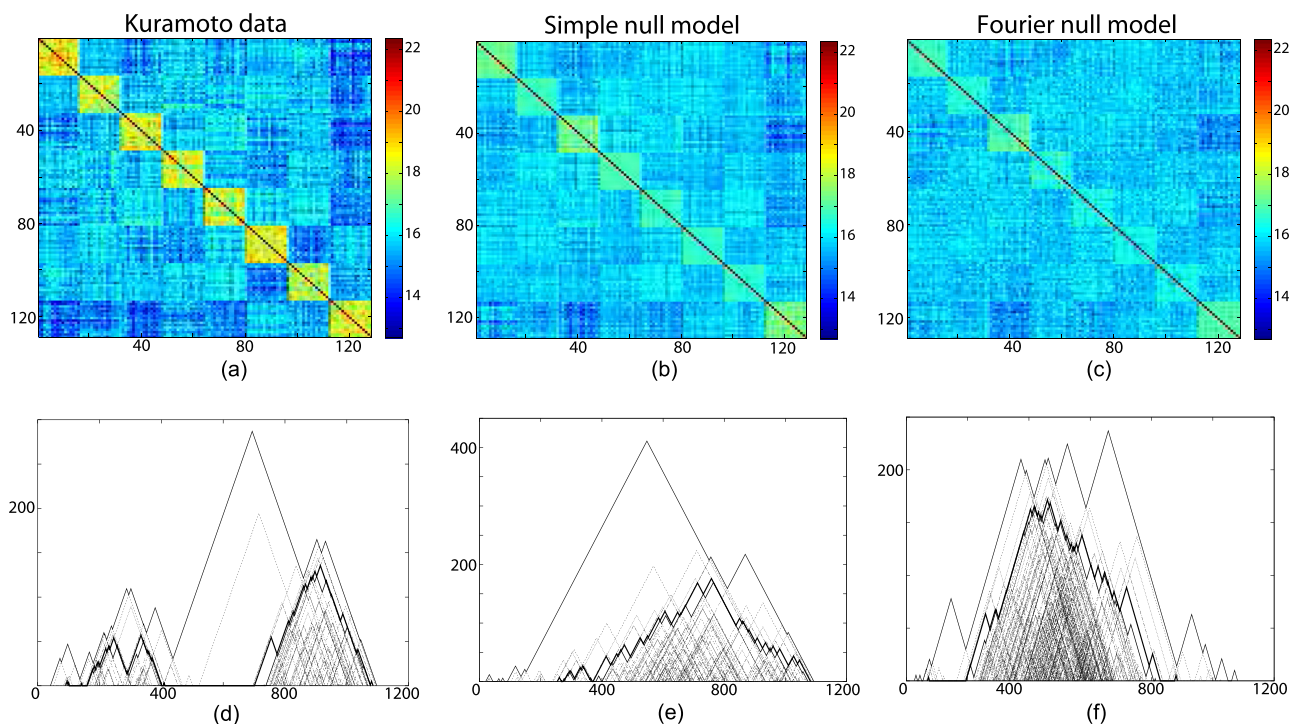


FIG. 8. (Top row) Functional networks for (left) the Kuramoto model, (center) the simple null model, and (right) the Fourier null model. (Bottom row) Dimension-1 persistence landscapes for the WRCF of (left) the Kuramoto model, (center) the simple null model, and (right) the Fourier null model using one time regime and ignoring infinitely-persisting bars. The persistence landscapes illustrate differences in the occurrence of loops in the three different networks. Most prominently, these differences manifest in the heights and distributions of the peaks in the landscapes, which appear to exhibit a stronger separation along the filtration between groups of peaks of different heights for the Kuramoto model than in the two null models.

landscapes is 13263; and the L^2 distance between the Kuramoto landscape and the simple null-model landscape is 11703. Again considering the support of the landscapes and the magnitudes of the function values, we see that three distances can be construed as large.

We find that PH can detect Kuramoto-model dynamics and that the persistence landscapes are rather different for the Kuramoto model and the null models. The L^2 distances between landscapes underscore these differences. We are also able to distinguish between the two null models using persistence landscapes. In contrast to conventional wisdom,^{54,55} we do not find for our examples that only the persistence of topological features distinguishes signal from noise. In fact, the short bars early in the filtration of the Kuramoto model carry important information about the dynamics, and the medium-sized persistent peaks in the Fourier null model are a symptom of the weaker intra-community and stronger inter-community synchronization in that model. We therefore assert that the position of features in the barcode is as important as persistence length for their interpretation in our examples; this provides an important point to consider for future studies. Note that persistence landscapes alone do not provide enough information to assess a system's dynamics. It is only by combining them with information about nodes that form loops (which are represented by certain groups of peaks) that we are able to obtain conclusions about intra-community and inter-community synchronization.

IV. EXAMPLE II: TASK-BASED fMRI DATA

A. Human brain networks during learning of a simple motor task

We use a data set of functional brain networks from experiments that were first analyzed by Bassett *et al.*⁸⁶ The data set was collected to study human subjects during learning of a simple motor task, and a full description of the experiments conducted is available in Ref. 86. We apply a WRCF to the functional networks, and we compare our findings with previous studies on these and similar networks.^{86–88} The functional networks are based on functional magnetic resonance imaging (fMRI) time series^{89,90} from 20 healthy subjects who undertook a motor-learning task on three days (during a five-day period). During the imaging of the subjects, an “atlas” of 112 brain areas was monitored while they were performing a simple motor-learning task (similar to a musical sequence), which they executed using four fingers of their non-dominant hand. For each subject and for each day of the study, the fMRI images are interpreted as 2000 time points for each monitored brain region. The brain regions and their time series were used subsequently to construct functional networks based on a functional connectivity measure known as the coherence of the wavelet scale-2 coefficients. This measure was applied to the time series to determine edge weights between each pair of brain regions in the network. The weighted adjacency

matrices for the functional networks were then corrected for a false-discovery rate, as matrix elements under a certain threshold (which represents a coherence amount that one expects to occur at random) were set to 0. The other matrix elements were retained.

The functional networks that we just described were studied previously using community detection by Bassett *et al.*,⁸⁶ whose results suggest that there is a significant segregation of the nodes in the functional networks into a small number of different communities with strongly-weighted connections inside the communities and weakly-weighted connections to nodes in other communities. Within these communities, certain nodes appeared to remain in the same community, whereas others (the “flexible” ones) often switched between different communities.

There have also been studies of networks from a similar experiment that examined medium-term learning (instead of short-term learning) and included training sessions.^{87,88} These networks have a noticeable core–periphery organization, with the sensorimotor and visual regions of the brain grouped into a temporally “stiff” core of nodes, whose community memberships (in contrast to flexible, peripheral nodes) do not change much over the course of the learning task.⁸⁷ It was also shown subsequently that the interaction of the primary and secondary sensorimotor regions with the

primary visual cortex decreases as the regions (presumably) become more autonomous with task practice.⁸⁸

Because we observed short-lived loops early in the filtrations for the Kuramoto model in a simulated setting with community structure in oscillator connectivity, we will explore whether the fMRI data exhibit similar features during the three observation days.

B. Persistent homology applied to the task-based fMRI data

We run the WRCF until filtration step 2600, which is when 42% of the edges are present in the network. (Note that using more filtration steps leads to very long computational times.) We again focus our analysis on topological features in dimension 1. We construct persistence landscapes for dimension 1 (omitting infinitely-persisting loops). In Fig. 9, we summarize our results for one particular subject and for the whole data set. We use this subject to illustrate a representative example of the particular landscape features that we observe in the data.

Similar to the Kuramoto oscillators in Section III, we find a group of small peaks early in the filtration (between filtration steps 1 and 200). We can see this group very clearly by magnifying both the landscape of individual subjects and the average landscape, whose peak heights are slightly smaller than the heights of the peaks in the

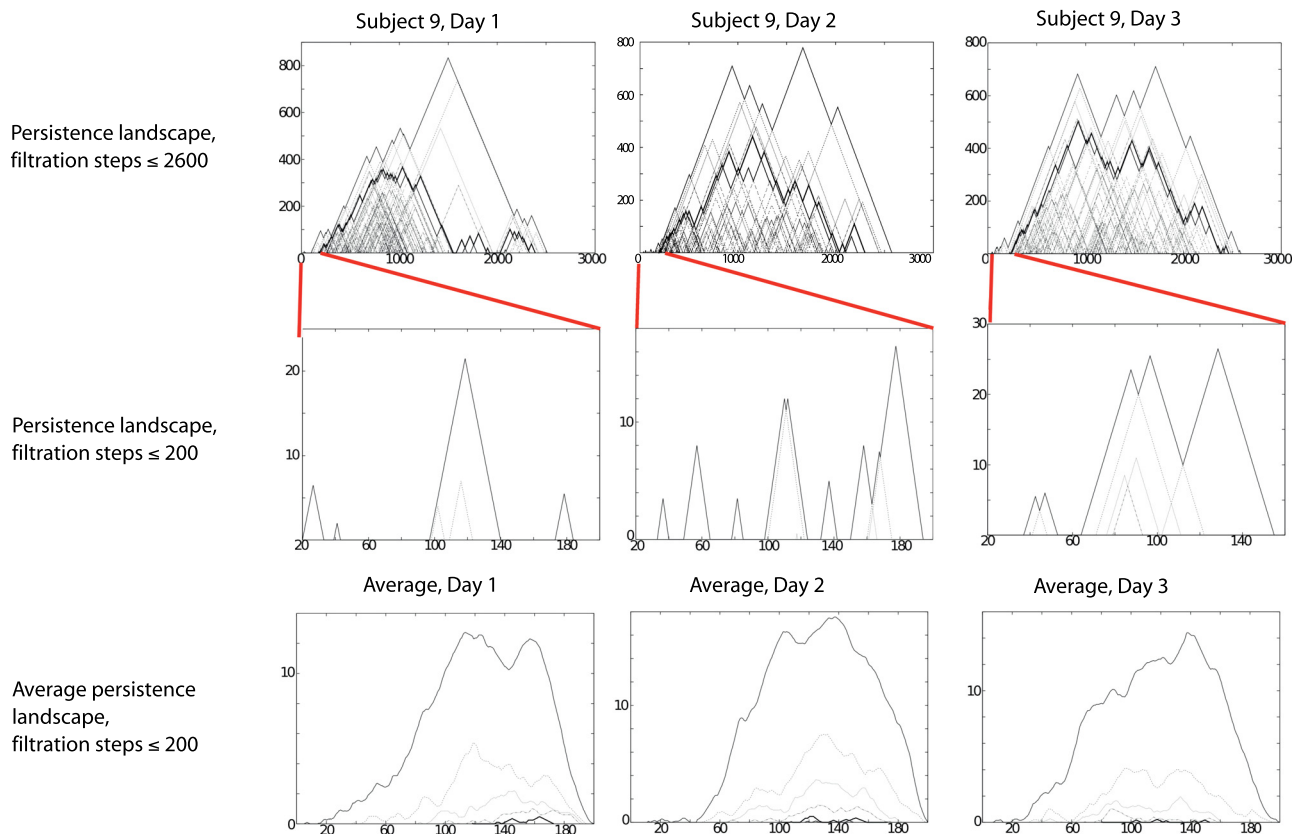


FIG. 9. Persistence landscapes for dimension 1 of the WRCF applied to the human brain networks. (First row) Persistence landscapes for subject 9 based on filtration steps 1–2600 for days 1, 2, and 3. (Second row) Persistence landscapes for subject 9 based on filtration steps 1–200 for days 1, 2, and 3. (Third row) Average persistence landscapes over all subjects for days 1, 2, and 3. We observe on average that short peaks occur in the first 200 filtration steps of the landscapes.

individual landscape that we show. This feature of the heights indicates that a group of short peaks arises early in the filtration in the majority of the barcodes. We also consider the standard deviation from the average landscapes in the first 200 filtration steps. For all three days, it is very small: it is 127 for the first day, 167 for the second day, and 126 for the third day.

We expect the observed short peaks early in the filtration to be associated with network communities, which have been observed previously using other methods.⁸⁶ We observe, in particular, that these short peaks undergo changes on day 2: during filtration steps 20 to 60, some of the peaks that are present in the landscapes for days 1 and 3 vanish, and a larger number of persistent peaks occur for day 3 than for the other two days between filtration steps 80 and 200. This appears to suggest that there is a change in community structure that takes place on day 2, with either (1) very strong synchronization in some of the communities, leading to very short-lived loops; or (2) very strong individual differences between the subjects, leading to the vanishing of peaks in the average landscapes for the first 50 filtration steps. The particularly persistent peaks on day 2 could represent either persistent inter-community loops or loops that occur due to sparse intra-community connections.

We calculate pairwise L^2 distances between each pair of dimension-1 persistence landscapes. We create distance vectors, which we use as an input for k -means clustering (with $k=3$) and average linkage clustering (until we obtain 3 groups), and we obtain the same qualitative result for both methods. We find that 9 of the 20 distance vectors that correspond to persistence landscapes from day 1 are assigned to a common group (together with a small number of landscapes from days 2 and 3), whereas 11 and 10 landscapes from days 2 and 3, respectively, are assigned together to a separate group. We summarize our results in Table I.

We also consider the average dimension-1 landscapes for WRCF steps 1–2600 and calculate the L^2 distances between them. We show the results of these calculations in Fig. 10.

The distances between the average landscape for day 1 and the subsequent days of the experiment indicate that the WRCFs on average are able to detect changes in the functional networks across the filtration range. Based on the distances, we observe that most of these changes occur between the first and second days. However, the standard deviations from the average landscapes are a factor of about 4 larger than the distances between the landscapes, and one therefore needs to be cautious about interpreting

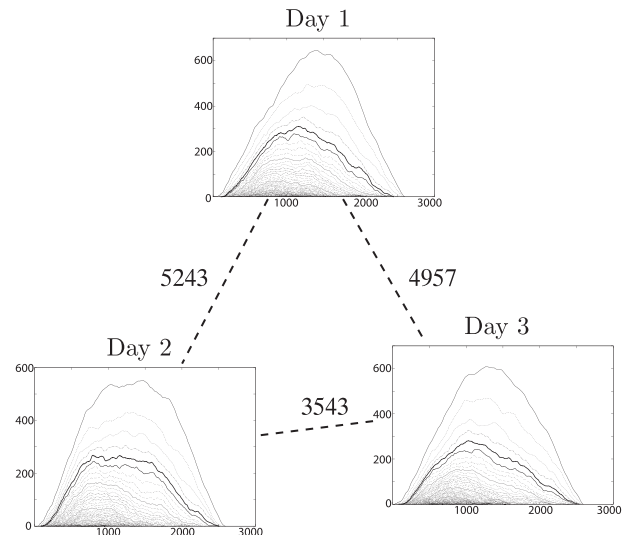


FIG. 10. Visualization of average persistence landscapes for days 1, 2, and 3 of task-based fMRI networks. The distance between the landscape for day 1 and the other two landscapes is larger than that between the landscapes for days 2 and 3. (The L^2 distances between them are 5243 between days 1 and 2, 4957 between days 1 and 3, and 3543 between days 2 and 3.) The standard deviations from the average landscapes are larger than the calculated distances, so these values need to be interpreted cautiously. We also observe a shift to the left of the landscape peak during the three days, indicating that the particularly persistent loops in these networks arise earlier in the filtration for the later days. In other words, they are formed by edges with higher edge weights, indicating that there is stronger synchronization between the associated brain regions.

the results of these calculations. In a permutation test with 10000 regroupings of the landscapes, we do not find the distances to be statistically significant. We obtain p -values of about 0.4 for the distance between the average landscapes of day 1 and day 2, about 0.85 for the distance between the average landscapes of day 2 and day 3, and about 0.6 for the distance between the average landscapes of day 1 and day 3.

We also find that the primary peak in the average landscapes in Fig. 10 shifts to the left over the course of the three days. This implies that the edge weights (between the brain regions) that give rise to persistent loops increase on average over the three days (presumably due to stronger synchronization). This can imply either that loops present on the first day consist of edges with higher edge weights on days 2 and 3 (i.e., the same brain regions synchronize more on the latter two days) or that new loops that appear on days 2 and 3 consist of edges with higher weights than those of day 1 but involve different brain regions than on day 1 (i.e., different brain regions exhibit stronger synchronization than what occurs on day 1). Brain regions that synchronize in a loop in a network may be an indication of an interesting neurobiological communication pattern that in this case also becomes stronger over the course of the learning process. To analyze the most frequently occurring edges involved in these loops, we extract “representatives” for all loops in dimension 1 across all subjects and days. (See Fig. 3 for an illustration of two different representatives of a loop in a network.) For each day, we construct a network, which we call an “occurrence network,” using the

TABLE I. Results for k -means clustering and average linkage clustering of pairwise L^2 distance vectors of persistence landscapes for $k=3$. Each value in the table indicates the number of landscapes in a cluster.

	Cluster 1	Cluster 2	Cluster 3
Day 1	9	6	5
Day 2	5	4	11
Day 3	5	5	10

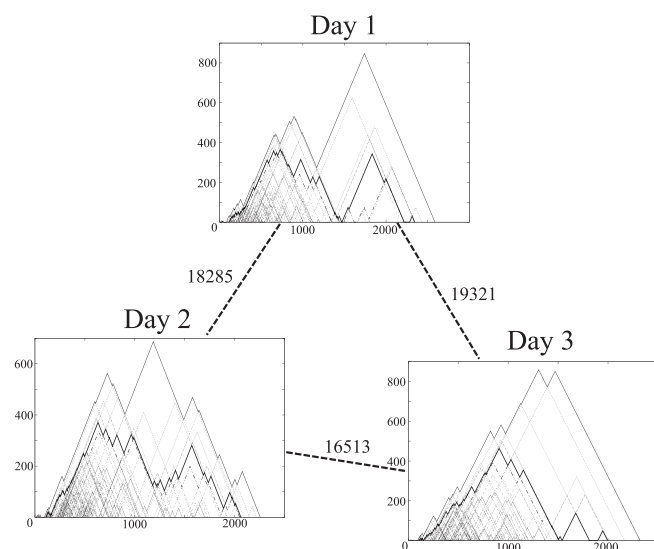


FIG. 11. Visualization of persistence landscapes based on average functional networks on days 1, 2, and 3 of the motor-learning task. The distance between the landscape for day 1 and the other two landscapes is larger than that between the landscapes for days 2 and 3. (The L^2 distances between them are 18285 between the first and second days, 16513 between the first and third days, and 19321 between the second and third days.) We find short peaks early in the filtration for all three landscapes, and larger peaks begin earlier in the filtration on day 3 than on day 1.

same nodes (i.e., brain regions) that we used before and assign every edge an edge weight that is equal to the number of occurrences of that edge in loops in the subjects on the given day. We then perform a WRCF on the three occurrence networks and study the representative loops that we obtain. In Table II in the Appendix, we list the brain regions that we find in loops consisting of edges that occur at least 50 times in functional networks in the subjects. We now examine loops in the occurrence networks. These particular loops may not correspond exactly to loops in associated functional networks. For example, individual edges with high edge weights that are part of a loop in the occurrence network may be part of a variety of different loops in associated functional networks, rather than part of one specific loop that occurs in many of the functional networks. Nevertheless, it is very likely that such loops are also loops in a functional network. One also needs to consider that the representative loops given by the software JAVAPLEX are not necessarily chosen optimally or are “geometrically nice”⁹¹ representatives of the loop.⁶⁶ We address the issue of JAVAPLEX’s choice of representatives to some extent by using PH on occurrence networks, but even then we cannot rule out possible artifacts. There exist loops in the occurrence networks that remain stable across the three days, although other loops occur on only one or two days. There also seem to be more loops that occur at least 50 times in the functional networks on days 2 and 3 than on day 1. It would be useful to study the brain regions (see Table II in the Appendix) involved in these loops to investigate their biological role(s) in motor-learning tasks.

Finally, we also apply WRCF to the average networks for each of the three days. To create an average network, we take the mean of the edge-weight values over all 20 subjects

for each day separately and study the resulting network. We show the corresponding landscapes in Fig. 11.

As with the average landscapes, we find that the landscapes for the average networks have very short peaks early in the filtration. There are more-persistent features (e.g., larger peaks) on day 1 and day 3 than on day 2, and we even find (as in the average landscapes) that the larger peaks appear earlier (at about filtration step 400) in the filtration on day 3 than on day 1 (where they appear at about step 900). Additionally, on day 2, we observe many short peaks, especially in the later stages of the filtration. This is not the case for day 1 and day 3, so the day-2 landscape is strikingly different visually from the other two landscapes. When calculating L^2 distances, we again find that the landscape distance between days 1 and 2 and that between days 1 and 3 are larger than the landscape distance between days 2 and 3. From visual inspection, we see that this arises from the fact that the day-1 landscape appears to have a clearer separation of short and high peaks than the landscapes for the later days. Taken together, the results for the landscapes of the average networks mirror our prior results for the average landscapes.

V. CONCLUSION AND DISCUSSION

We have illustrated applications of persistent homology to functional networks constructed from time-series output of the Kuramoto model, null models constructed from the Kuramoto time series, and task-based fMRI data from human subjects. In all cases, we observed that non-persistent loops occur early in the filtrations. Although such non-persistent features are commonly construed as noise in topological data analysis,^{54,55} we observed that these features appear to be consistent with prior segregations of the studied networks into communities of densely-connected nodes. In one case (the Fourier null model), we even found that particularly persistent features appear to be linked to a network with a weak intra-community synchronization. These very persistent features in the null model may thus represent noise. In other studies of PH using (different) null models,^{33,53,59} it was also observed that the null models often exhibit a richer topological structure than the empirical data. One could thus perhaps interpret the persistent features in the Fourier null model as features of the null model rather than as noise. Our results on the importance of non-persistent features resemble previous observations for synthetic examples with barcodes that consist of short intervals (which are commonly construed as noise), but the differences between the corresponding persistence landscapes for the various spaces are nevertheless statistically significant.⁹² Our results are also consistent with the findings of a study on protein structure using PH that reported that bars of any length in the barcodes were equally important.⁹³ For weighted networks, we suggest that when using a filtration based on edge weights, one needs to consider the actual birth and death times of filtration features (such as loops) in addition to their persistence to be able to determine whether they should be construed as noise or part of a signal. In particular, in the present paper, we observed that the early

appearance of loops in a filtration is an important distinguishing feature of these data. They may also yield important insights on the geometry⁹² of data.⁹⁴

We also found—both by calculating average persistence landscapes and studying landscapes of average networks—that persistence landscapes for dimension 1 of the weight rank clique filtration (WRCF) are able to capture changes in the studied functional brain networks during the process of learning a simple motor task. Because we did not consider infinitely-persisting features and only included filtration steps 1–2600 when creating the landscapes, our result also suggests that the medium-lived (when compared with the full filtration length) persistent loops are able to capture changes in the networks, so it is not always necessary to consider a full WRCF to study the dynamics of a system. This observation is similar to a finding of Bendich *et al.*,³⁵ who observed in their study that medium-scale barcode features were able to distinguish human brain artery networks from different age groups. This again suggests that persistence length should not be the only measure of signal versus noise when applying PH. We also found that the persistent features that dominate the middle part of the filtrations appear in earlier filtration steps on days 2 and 3 of the experiment than they do on day 1, which suggests that interesting dynamics in synchronization patterns are captured by medium-lived bars in the middle of a barcode.

As in other biological contexts, where PH has been applied successfully and has led to new insights,^{28,30–33,35} we find that PH can lead to fascinating insights about the dynamics of a system. We were able not only to detect symptoms of previously observed community segregation, but we also found notable differences between a setup with strong community structure (in the coupled Kuramoto oscillators) and weakly synchronized communities (in the associated null models). For the task-based fMRI data, we found that we can detect symptoms of community structure over the three days (in the short peaks early in the landscapes) of the data as well as changes in the loops that we observe in the average persistence landscapes of the functional networks. On average, most of these changes appear to take place on day 2 of the learning task. In particular, brain regions that yield loops in the functional networks on days 2 and 3 seem to exhibit stronger synchronization on average than those that yield loops on day 1. We obtained this observation both by calculating average persistence landscapes of the WRCF performed on individual functional networks and by calculating persistence landscapes based on the WRCF performed on average networks for each day. Although the landscape distances between the average landscapes are not statistically significant, our similar observations in both of our approaches suggest that our observations indeed reflect the average dynamics of the system. Our findings on loops thereby provide novel insights that complement previous studies of synchronization in functional brain networks. Of course, it is desirable to repeat our study using larger data sets.

There is a known relation between homology and graph Laplacians,⁹⁵ and an interesting possible direction for future research would be to study the possible connections between

graph Laplacians (and, more generally, spectral graph theory) and our results on barcodes and persistence landscapes.

Using methods from topological data analysis for studying networks has the important benefit of being both mathematically principled and generalizable. However, for biological interpretation, it is necessary to include information on the specific nodes that are part of topological features such as loops. Moreover, the interpretation of the results and importance of persistence versus position of a topological feature in a barcode can differ depending on which type of filtration is employed. Different topological features can also have different levels of relevance for different dynamical systems. For example, the occurrence of many medium-sized persistent features in the persistence landscape for the Fourier null model is a symptom of the weak synchronization in the communities, whereas the medium-sized persistent bars capture increasing synchronization in loops for the task-based fMRI data. It would be interesting to apply WRCF (and other types of filtrations) to different synthetic networks with underlying communities (e.g., using stochastic block models) to investigate such ideas further. Importantly, one should include both the persistence and the position of topological features in analysis of PH. It would also be beneficial to combine topological tools with additional methods, such as persistence images,⁹⁶ to determine the exact topological features that are responsible for the detected differences between the persistence landscapes of the different networks.

In conclusion, we have shown that persistent homology and persistence landscapes can be applied successfully to functional networks (from either experimental data or time-series output of models), and that they can lead to fascinating insights, such as segregation of a network into communities and changes of network structure over time.

SUPPLEMENTARY MATERIAL

See [Supplementary Material](#) for additional mathematical background, definitions, examples, and theorems.

ACKNOWLEDGMENTS

The experimental data were collected originally by Nicholas F. Wymbs and Scott T. Grafton through funding from the Public Health Service Grant No. NS44393, and we thank Nicholas and Scott for access to the data. We thank Danielle S. Bassett for help in providing the data, use of her MATLAB code when we were debugging our code, and helpful discussions. We also thank Pawel Dłotko for useful discussions, his help with the Persistence Landscapes toolbox, and providing us with new versions of his code during our work. We also thank Alex Arenas and Nina Otter for helpful comments. B.J.S. thanks the Berrow foundation for funding during her M.Sc. degree, and she also gratefully acknowledges the EPSRC, MRC (Grant No. EP/G037280/1), and F. Hoffmann–La Roche AG for funding her doctoral studies. H.A.H. acknowledges funding from EPSRC Fellowship No. EP/K041096/1.

APPENDIX: TABLE OF OFTEN-OCCURRING BRAIN REGIONS IN LOOPS

In Table II, we indicate the brain regions that occur often in loops.

TABLE II. Loops (in occurrence networks) that consist of edges that occur in loops of functional networks at least 50 times over all subjects. We list the loops that we find in the left column. (We start with one of the nodes, which we choose arbitrarily, and end with the node that is adjacent to the starting node in the loop.) We denote an occurrence of a loop on a specific day with the symbol “x” in the table and present variations of the loop that we interpret as representing the same loop. We use the following abbreviations for the brain regions: l: left; r: right; ant: anterior; post: posterior; AnGy: Angular gyrus; CinGy: Cingulate gyrus; COC: Central opercular cortex; FOC: Frontal operculum cortex; FMedC: Frontal medial cortex; FP: Frontal pole; HG: Heschl’s gyrus; IC: Insular cortex; InfFGyPT: Inferior frontal gyrus pars triangularis; IntCalC: Intracalcrine cortex; LinGy: Lingual gyrus; OFG Occipital fusiform gyrus; OFC: Orbital frontal cortex; OP: Occipital pole; PaCinGy: Paracingulate gyrus; ParOpC: Parietal operculum cortex; PHGy: Parahippocampal gyrus; PostGy: Postcentral gyrus; PP: Planum polare; PreGy: Precentral gyrus; PT: Planum temporale; Put: Putamen; SupCalC: Supercalcrine Cortex; SuppMA: Supplemental motor area; SupMargGy: Supramarginal gyrus; SupPL: Superior parietal lobule; SupTempGy: Superior temporal gyrus; InfFGyPO: Inferior frontal gyrus pars opercularis; MTGy: Middle temporal gyrus.

Loop	Day 1	Day 2	Day 3
–lSuppMA–rSuppMA–rPreGy–lPreGy–	x	x	x
–lOFG–lOP–rOP–rOFG–	x	x	x
–lSupTempGy ant–lPP–lHG–lPT–lSupTemGy post–	x	x	x
–lIC–rIC–rPP–lPP–	x	Variant: –rIC–rPP–lPP– lHG–lCOC–lIC–	Variant: –rIC–rPP–rHG– lPP–lIC–
–rIC–lIC–lPut–rPut–	x	x	x
–lIntCalC–lLinGy–lOFG–rOFG–rLinGy–rIntCalC–rSupCalC–lSupCalC–	x		Variant: –lIntCalC–lLinGy– rLinGy–rIntCalC–
–lFP–lPaCinGy–rPaCinGy–rFP–		x	x
–rPP–lPP–lHG–lCOC–lIC–lFOC–lInfFGyPO–lInfFGyPT–lFP–lSuppMA– lPreGy–RPreGy–rPostGy–rSupMargGy ant–rParOpC–rPT– rSupTempGy post–rSupTempGy ant–		x	Variant: –rPP–rHG–rPT–rSup TempGy post–rSupTempGy ant–
–rOFG–lOFG–lLinGy–rLinGy–	x		
–lPaCinGy–rPaCinGy–rCinGy ant–lCinGy ant–		x	
–lFP–lFMedC–rFMedC–rFP–		x	
–lIC–lCOC–lHG–lPP–		x	
–lPHGy ant–lPHGy–rPHGy–rPHGy ant–		x	
–lInfFGyPT–lInfFGyPO–lFOC–lIC–lPP–lSupTempGy ant–lSupTemGy post– lMTGy post–lMTGy ant–lFP–lOFC–		x	
–rSupPL–rSupMargGy post–rSupMargGy ant–rPostGy–RPreGy–			x
–rPostGy–rSupPL–lSupPL–lSupMargGy ant–rSupMargGy ant–			x
–lIntCalC–lLinGy–rLinGy–rIntCalC–			x
–lSupMargGy post–lAnGy–rAnGy–rSupMargGy post–rSupMargGy ant– lSupMargGy ant–			x
–lPT–lHG–lPP–lSupTempGy ant–lSupTemGy post–			x
Total number of loops	7	12	13

¹H. Edelsbrunner, D. Letscher, and A. Zomorodian, “Topological persistence and simplification,” *Discrete Comput. Geom.* **28**, 511–533 (2002).
²H. Edelsbrunner and J. L. Harer, “Persistent homology—A survey,” in *Surveys on Discrete and Computational Geometry. Twenty Years Later*, Contemporary Mathematics Vol. 453, edited by J. E. Goodman, J. Pach, and R. Pollak (American Mathematical Society, 2008), pp. 257–282.
³H. Edelsbrunner and J. L. Harer, *Computational Topology* (American Mathematical Society, Providence, RI, 2010).
⁴R. Ghrist, *Elementary Applied Topology*, 1st ed. (Createspace, 2014); see <https://www.math.upenn.edu/~ghrist/notes.html>.
⁵N. Otter, M. A. Porter, U. Tillmann, P. Grindrod, and H. A. Harrington, “A roadmap for the computation of persistent homology,” e-print [arXiv:1506.08903v6](https://arxiv.org/abs/1506.08903v6).
⁶M. Kramár, A. Goulet, L. Kondic, and K. Mischaikow, “Persistence of force networks in compressed granular media,” *Phys. Rev. E* **87**, 042207 (2013).
⁷C. Curto, “What can topology tell us about the neural code?,” *American Mathematical Society* **54**(1), 63–78 (2017).
⁸C. Giusti, R. Ghrist, and D. S. Bassett, “Two’s company and three (or more) is a simplex,” *J. Comput. Neurosci.* **41**, 1–14 (2016).
⁹M. A. Porter, J.-P. Onnela, and P. J. Mucha, “Communities in networks,” *Not. Am. Math. Soc.* **56**, 1082–1097, 1164–1166 (2009); available at <http://www.ams.org/notices/200909/rtx090901082p.pdf>.

¹⁰S. Fortunato and D. Hric, “Community detection in networks: A user guide,” *Phys. Rep.* **659**, 1–44 (2016).
¹¹E. T. Bullmore and O. Sporns, “Complex brain networks: Graph theoretical analysis of structural and functional systems,” *Nat. Rev.* **10**, 186–198 (2009).
¹²E. T. Bullmore and D. Bassett, “Brain graphs: Graphical models of the human brain connectome,” *Annu. Rev. Clin. Psychol.* **7**, 113–140 (2011).
¹³O. Sporns, “Contributions and challenges for network models in cognitive neuroscience,” *Nat. Rev. Neurosci.* **17**, 652–660 (2014).
¹⁴D. Papo, M. Zanin, J. A. Pineda-Pardo, S. Boccaletti, and J. M. Buldú, “Functional brain networks: Great expectations, hard times and the big leap forward,” *Philos. Trans. R. Soc. B* **369**, 20130525 (2014).
¹⁵D. Papo, J. M. Buldú, S. Boccaletti, and E. T. Bullmore, “Complex network theory and the brain,” *Philos. Trans. R. Soc. B* **369**, 20130520 (2014).
¹⁶R. F. Betzel and D. S. Bassett, “Multi-scale brain networks,” *NeuroImage* (in press); available at <http://dx.doi.org/10.1016/j.neuroimage.2016.11.006>.
¹⁷B. Alberts, D. Bray, K. Hopkin, A. Johnson, J. Lewis, M. Raff, K. Roberts, and P. Walter, *Essential Cell Biology* (Garland Science, New York and London, 2014).
¹⁸E. T. Bullmore and O. Sporns, “The economy of brain network organization,” *Nat. Rev. Neurosci.* **13**, 336–349 (2012).

- ¹⁹M. E. J. Newman, *Networks: An Introduction* (Oxford University Press, Oxford, 2010).
- ²⁰In the present paper, we use the terms “network” and “graphs” synonymously, although the former is often used in a way that includes structures that are more complicated than ordinary graphs.
- ²¹B. Bollobás, *Modern Graph Theory* (Springer, New York, 1998).
- ²²D. S. Bassett, N. F. Wymbs, M. A. Porter, P. J. Mucha, and S. T. Grafton, “Cross-linked structure of network evolution,” *Chaos* **24**, 013112 (2014).
- ²³A. R. Benson, D. F. Gleich, and J. Leskovec, “Higher-order organization of complex networks,” *Science* **353**, 163–166 (2016).
- ²⁴D. Taylor, F. Klimm, H. A. Harrington, M. Kramár, K. Mishchaikow, M. A. Porter, and P. J. Mucha, “Topological data analysis of contagion maps for examining spreading processes on networks,” *Nat. Commun.* **6**, 7723 (2015).
- ²⁵S. Bhattacharya, R. Ghrist, and V. Kumar, “Persistent homology for path planning in uncertain environments,” *IEEE Trans. Rob.* **31**, 578–590 (2015).
- ²⁶C. M. Topaz, L. Ziegelmeier, and T. Halverson, “Topological data analysis of biological aggregation models,” *PLoS ONE* **10**, e0126383 (2015).
- ²⁷H. Lee, M. K. Chung, H. Kang, B.-N. Kim, and D. S. Lee, “Discriminative persistent homology of brain networks,” in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro* (2011), pp. 841–844.
- ²⁸G. Petri, P. Expert, F. Turkheimer, R. Carhart-Harris, D. Nutt, P. J. Hellyer, and F. Vaccarino, “Homological scaffolds of brain functional networks,” *J. R. Soc. Interface* **11**, 20140873 (2014).
- ²⁹G. Spreemann, B. Dunn, M. B. Botnan, and N. A. Baas, “Using persistent homology to reveal hidden information in neural data,” e-print [arXiv:1510.06629](https://arxiv.org/abs/1510.06629).
- ³⁰P. Dłotko, K. Hess, R. Lavi, M. Nolte, M. Reimann, M. Scholamiero, K. Turner, E. Muller, and H. Markram, “Topological analysis of the connectome of digital reconstructions of neural microcircuits,” e-print [arXiv:1601.01580](https://arxiv.org/abs/1601.01580).
- ³¹C. Curto and V. Itskov, “Cell groups reveal structure of stimulus space,” *PLoS Comput. Biol.* **4**, e1000205 (2008).
- ³²Y. Dabaghian, F. Mémoli, L. Frank, and G. E. Carlsson, “A topological paradigm for hippocampal spatial map formation using persistent homology,” *PLoS Comput Biol* **8**(8), e1002581 (2012).
- ³³C. Giusti, E. Pastalkova, C. Curto, and V. Itskov, “Clique topology reveals intrinsic geometric structure in neural correlations,” *Proc. Natl. Acad. Sci. U.S.A.* **112**, 13455–13460 (2015).
- ³⁴A. Babichev and Y. Dabaghian, “Persistent memories in transient networks,” e-print [arXiv:1602.00681](https://arxiv.org/abs/1602.00681).
- ³⁵P. Bendich, J. S. Marron, E. Miller, A. Pieloch, and S. Skwerer, “Persistent homology analysis of brain artery trees,” *Ann. Appl. Stat.* **10**, 198–218 (2016).
- ³⁶E. C. Zeeman, “The topology of the brain and visual perception,” in *The Topology of 3-Manifolds*, edited by M. K. Fort (Prentice Hall, Englewood Cliffs, NJ, 1962), pp. 240–256.
- ³⁷M. R. Muldoon, R. S. MacKay, J. P. Huke, and D. S. Broomhead, “Topology from time series,” *Physica D* **65**, 1–16 (1993).
- ³⁸O. Sporns, “Graph-theoretical analysis of brain networks,” in *Brain Mapping: An Encyclopedic Reference*, edited by A. W. Toga (Academic Press, Elsevier, Cambridge, Massachusetts, 2015), Vol. 1, pp. 629–633.
- ³⁹S. E. Petersen and O. Sporns, “Brain networks and cognitive architectures,” *Neuron* **88**, 207–219 (2015).
- ⁴⁰G. Tirabassi, R. Sevilla-Escoboza, J. M. Buldú, and C. Masoller, “Inferring the connectivity of coupled oscillators from time-series statistical similarity analysis,” *Sci. Rep.* **5**, 1–14 (2015).
- ⁴¹X. Sun, M. Small, Y. Zhao, and X. Xue, “Characterizing system dynamics with a weighted and directed network constructed from time series data,” *Chaos* **24**, 024402 (2014).
- ⁴²T. Nakamura, T. Tanizawa, and M. Small, “Constructing networks from a dynamical system perspective for multivariate nonlinear time series,” *Phys. Rev. E* **93**, 032323 (2016).
- ⁴³D. J. Fenn, M. A. Porter, M. McDonald, S. Williams, N. F. Johnson, and N. S. Jones, “Dynamic communities in multichannel data: An application to the foreign exchange market during the 2007–2008 credit crisis,” *Chaos* **19**, 033119 (2009).
- ⁴⁴A. S. Waugh, L. Pei, J. H. Fowler, P. J. Mucha, and M. A. Porter, “Party polarization in congress: A network science approach,” e-print [arXiv:0907.3509](https://arxiv.org/abs/0907.3509).
- ⁴⁵J. F. Donges, Y. Zou, N. Marwan, and J. Kurths, “The backbone of the climate network,” *Europhys. Lett. (EPL)* **87**, 48007 (2009).
- ⁴⁶D. S. Bassett, M. A. Porter, N. F. Wymbs, S. T. Grafton, J. M. Carlson, and P. J. Mucha, “Robust detection of dynamic community structure in networks,” *Chaos* **23**, 013142 (2013).
- ⁴⁷S. M. Smith, K. L. Miller, G. Salimi-Khorshidi, M. Webster, C. F. Beckmann, T. E. Nichols, J. D. Ramsay, and M. W. Woolrich, “Network modelling methods for fMRI,” *NeuroImage* **54**, 875–891 (2011).
- ⁴⁸D. Zhou, W. K. Thompson, and G. Siegle, “MATLAB toolbox for functional connectivity,” *NeuroImage* **47**, 1590–1607 (2009).
- ⁴⁹M. Á. Serrano, M. Boguná, and A. Vespignani, “Extracting the multiscale backbone of complex weighted networks,” *Proc. Natl. Acad. Sci. U.S.A.* **106**, 6483–6488 (2009).
- ⁵⁰A. F. Alexander-Bloch, N. Gogtay, D. Meunier, R. Birn, L. Clasen, F. Lalonde, R. Lenroot, J. Giedd, and E. T. Bullmore, “Disrupted modularity and local connectivity of brain functional networks in childhood-onset schizophrenia,” *Front. Syst. Neurosci.* **4**, 1–16 (2010).
- ⁵¹F. D. V. Fallani, J. Richiardi, M. Chavez, and S. Archard, “Graph analysis of functional brain networks: Practical issues in translational neuroscience,” *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* **369**, 0130521 (2014).
- ⁵²Conventionally, loops (other than self-loops) in undirected graphs must have at least 3 edges, and loops in directed graphs must have at least 2 edges. In our paper, we adapt the terminology to represent the topological features that we detect in our simplicial complexes.
- ⁵³A. Sizemore, C. Giusti, R. F. Betzel, and D. S. Bassett, “Closures and cavities in the human connectome,” e-print [arXiv:1608.03520v2](https://arxiv.org/abs/1608.03520v2).
- ⁵⁴R. Ghrist, “Barcodes: The persistent topology of data,” *Bull. Am. Math.* **45**, 61–75 (2008).
- ⁵⁵G. Carlsson, “Topology and data,” *Bull. Am. Math. Soc.* **46**, 255–308 (2009).
- ⁵⁶C. Kosniowski, *A First Course in Algebraic Topology* (Cambridge University Press, Cambridge, London, New York, New Rochelle, Melbourne, Sydney, 1980).
- ⁵⁷J. R. Munkres, *Topology* (Pearson Prentice Hall, New Jersey, 2000).
- ⁵⁸H. Lee, H. Kang, M. K. Chung, B.-N. Kim, and D. S. Lee, “Weighted functional brain network modeling via network filtration,” in NIPS Workshop on Algebraic Topology and Machine Learning (2012).
- ⁵⁹G. Petri, M. Scolamiero, I. Donato, and F. Vaccarino, “Topological strata of weighted complex networks,” *PLoS ONE* **8**, e66505 (2013).
- ⁶⁰P. Bubenik, “Statistical topological data analysis using persistence landscapes,” *J. Mach. Learn. Res.* **16**, 77–102 (2015), available at <http://www.jmlr.org/papers/volume16/bubenik15a/bubenik15a.pdf>.
- ⁶¹P. Bubenik and P. Dłotko, “A persistence landscapes toolbox for topological statistics,” *J. Symbolic Comput.* **78**, 91–114 (2017).
- ⁶²V. Kovacev-Nikolic, P. Bubenik, D. Nikolic, and G. Heo, “Using persistent homology and dynamical distances to analyze protein binding,” *Stat. Appl. Genet. Mol. Biol.* **15**, 1–27 (2016).
- ⁶³P. Dłotko and T. Wanner, “Topological microstructure analysis using persistence landscapes,” *Physica D* **334**, 60–81 (2016).
- ⁶⁴J.-Y. Liu, S.-K. Jeng, and Y.-H. Yang, “Applying topological persistence in convolutional neural network for music audio signals,” e-print [arXiv:1608.07373](https://arxiv.org/abs/1608.07373).
- ⁶⁵H. Adams, A. Tausz, and M. Vejdemo-Johansson, “JAVAPLEX: A research software package for persistent (co)homology (2011),” in *Mathematical Software - ICMS 2014*, edited by H. Hong and C. Yap (2014), Vol. 8592, pp. 129–136, see <http://javaplex.github.io/>.
- ⁶⁶H. Adams and A. Tausz, JAVAPLEX tutorial, 2015; http://www.math.colostate.edu/~adams/research/javaplex_tutorial.pdf.
- ⁶⁷J. Wildmann, Bron–Kerbosch maximal clique finding algorithm, 2011.
- ⁶⁸Y. Kuramoto, *Chemical Oscillations and Waves and Turbulence* (Springer, Berlin, 1984).
- ⁶⁹S. H. Strogatz, “From Kuramoto to Crawford: Exploring the onset of synchronization in populations of coupled oscillators,” *Physica D* **143**, 1–20 (2000).
- ⁷⁰A. Arenas, A. Díaz-Guilera, J. Kurths, Y. Moreno, and C. Zhou, “Synchronization in complex networks,” *Phys. Rep.* **469**, 93–153 (2008).
- ⁷¹F. Rodrigues, T. K. D. Peron, P. Ji, and J. Kurths, “The Kuramoto model in complex networks,” *Phys. Rep.* **610**, 1–98 (2016).
- ⁷²S. Gupta, A. Campa, and S. Ruffo, “Kuramoto model of synchronization: Equilibrium and nonequilibrium aspects,” *J. Stat. Mech.: Theory Exp.* **2014**, R08001 (2014).
- ⁷³J. L. P. Velazquez, “Brain research: A perspective from the coupled oscillators field,” *NeuroQuantology* **4**, 155–165 (2006).
- ⁷⁴M. Breakspear, S. Heitmann, and A. Daffertshofer, “Generative models of cortical oscillations: Neurobiological implications of the Kuramoto model,” *Front. Human Neurosci.* **4**, 1–14 (2010).

- ⁷⁵P. Ashwin, S. Coombes, and R. Nicks, “Mathematical framework for oscillatory network dynamics in neuroscience,” *J. Math. Neurosci.* **6**, 1–92 (2016).
- ⁷⁶W. H. Lee, E. Bullmore, and S. Frangou, “Quantitative evaluation of simulated functional brain networks in graph theoretical analysis,” *NeuroImage* **146**, 724 (2017).
- ⁷⁷M. A. Porter and J. P. Gleeson, “Dynamical systems on networks: A tutorial,” *Frontiers in Applied Dynamical Systems: Reviews and Tutorials* (Springer, 2016), Vol. 4.
- ⁷⁸A. Arenas, A. Díaz-Guilera, and C. Pérez-Vicente, “Synchronization reveals topological scales in complex networks,” *Phys. Rev. Lett.* **96**, 114102 (2006).
- ⁷⁹In this context, we use the term “community” to indicate a set of densely-connected nodes with sparse connections to other nodes outside of this set. There are also other uses of the term, and community structure is a popular subject in network science (see Refs. 9 and 10).
- ⁸⁰We use an input time step of $\Delta t = 0.02$, but we note that ODE45 uses an adaptive step size.
- ⁸¹J. Stout, M. Whiteway, E. Ott, M. Girvan, and T. M. Antonsen, “Local synchronization in complex networks of coupled oscillators,” *Chaos* **21**, 025109 (2011).
- ⁸²V. A. Traag and J. Bruggeman, “Community detection in networks with positive and negative links,” *Phys. Rev. E* **80**, 036115 (2009).
- ⁸³D. Prichard and J. Theiler, “Generating surrogate data for time series with several simultaneously measured variables,” *Phys. Rev. Lett.* **73**, 951–954 (1994).
- ⁸⁴B. Stolz, “Computational Topology in Neuroscience,” Master’s thesis (University of Oxford, 2014).
- ⁸⁵A. Sizemore, C. Giusti, and D. S. Bassett, “Classification of weighted networks through mesoscopic homological features,” *J. Complex Networks* **5**(2), 2450–273 (2017).
- ⁸⁶D. S. Bassett, N. F. Wymbs, M. A. Porter, P. J. Mucha, J. M. Carlson, and S. T. Grafton, “Dynamic reconfiguration of human brain networks during learning,” *Proc. Natl. Acad. Sci. U.S.A.* **108**, 7641–7646 (2011).
- ⁸⁷D. S. Bassett, N. F. Wymbs, M. P. Rombach, M. A. Porter, P. J. Mucha, and S. T. Grafton, “Task-based core–periphery organization of human brain dynamics,” *PLoS Comput. Biol.* **10**, e1003617 (2013).
- ⁸⁸D. S. Bassett, M. Yang, N. F. Wymbs, and S. T. Grafton, “Learning-induced autonomy of sensorimotor systems,” *Nat. Neurosci.* **18**, 744–751 (2015).
- ⁸⁹See Ref. 90 for a recent discussion of fMRI inference and potential perils in the statistical methods in use in neuroimaging.
- ⁹⁰A. Eklund, T. E. Nichols, and H. Knutsson, “Cluster failure: Why fMRI inferences for spatial extent have inflated false-positive rates,” *Proc. Natl. Acad. Sci. U.S.A.* **113**, 7900–7905 (2016).
- ⁹¹For example, a loop can be represented by a double loop.
- ⁹²P. Bubenik, personal communication (2016).
- ⁹³K. Xia and G.-W. Wei, “Persistent homology analysis of protein structure, flexibility, and folding,” *Int. J. Numer. Methods Biomed. Eng.* **30**, 814–844 (2014).
- ⁹⁴Note that we use the term “geometry” for properties that are called “shape” in other contexts (see, e.g., Ref. R. MacPherson and B. Schweinhart, “Measuring shape with topology,” *J. Math. Phys.* **53**, 073516 (2012)) to avoid confusion with our previous usage of the term “shape.”
- ⁹⁵F. R. K. Chung, *Spectral Graph Theory*, Regional Conference Series in Mathematics No. 92 (AMS and CBMS, 1997).
- ⁹⁶H. Adams, S. Chepushtanova, T. Emerson, E. Hanson, M. Kirby, F. Motta, R. Neville, C. Peterson, P. Shipman, and L. Ziegelmeier, “Persistent images: A stable vector representation of persistent homology,” *J. Mach. Learn. Res.* **18**(8), 1–35 (2017); available at <http://www.jmlr.org/papers/volume18/16-337/16-337.pdf>.

Persistent homology of time-dependent functional networks constructed from coupled time series: Supplementary Information

Bernadette J. Stolz,¹ Heather A. Harrington,¹ and Mason A. Porter^{2,1,3}

¹*Mathematical Institute, University of Oxford, Oxford OX2 6GG,
UK*

²*Department of Mathematics, University of California Los Angeles, Los Angeles,
USA*

³*CABDyN Complexity Centre, University of Oxford, Oxford OX1 1HP,
UK*

(Dated: 9 April 2017)

I. TOPOLOGICAL BACKGROUND AND DEFINITIONS

We give a brief introduction to the mathematical concepts behind persistent homology (PH). For our presentation, we adapt and summarize the discussion from B. Stolz’s masters thesis¹.

A. Simplicial complexes

One can represent the underlying structures of a topological space by partitioning the space into smaller and topologically simpler pieces, which carry the same aggregate topological information as the original space when they are assembled back together. One can choose either either a small number of complicated pieces or a large number of simple pieces. From a computational point of view, the latter is preferable².

A simple example for such a construction is the tetrahedron in Euclidian space. The tetrahedron consists of four triangular faces that are each bounded by three edges (which each connect two points). One can view the tetrahedron as a simplified version of a 2-sphere, as it carries the same topological properties (e.g., connectedness and the enclosure of a hole) as the sphere. Similarly, one can imagine using triangles as building blocks to build more complicated constructions (e.g., ones that resemble a torus or some other manifold).

To mathematically grasp these concepts, we need a few definitions. For concreteness, we frame our discussion using the space \mathbb{R}^d with dimension $d \in \mathbb{N}$.

Definition I.1 (*affine combination and affine hull*). Let $\mathcal{U} = \{u_0, u_1, \dots, u_k\}$ be points in \mathbb{R}^d . A point $x \in \mathbb{R}^d$ is an *affine combination* of the points $u_i \in \mathcal{U}$, with $i \in \{0, \dots, k\}$, if there exist $\lambda_i \in \mathbb{R}$ such that

- i. $x = \sum_{i=0}^k \lambda_i u_i$;
- ii. $\sum_{i=0}^k \lambda_i = 1$.

The set of all affine combinations of \mathcal{U} is called the *affine hull* of \mathcal{U} .

To ensure uniqueness of the affine combination, we introduce the following definition.

Definition I.2 (*affinely independent*). Let $\mathcal{U} = \{u_0, u_1, \dots, u_k\}$ be points in \mathbb{R}^d . The $k + 1$ points in \mathcal{U} are said to be *affinely independent* if the vectors $\{u_i - u_0 : i \in \{0, \dots, k\}\}$ are linearly independent.

For example, any two distinct points in \mathbb{R}^2 are affinely independent. Similarly, any three points in \mathbb{R}^2 are affinely independent as long as they do not lie on the same straight line.

Convex combinations and hulls are a special case of affine combinations.

Definition I.3 (*convex combination* and *convex hull*). An affine combination $x = \sum_{i=0}^k \lambda_i u_i$ is a *convex combination* if $\lambda_i \geq 0$ for all $i \in \{0, \dots, k\}$. The set of all convex combinations of the points in \mathcal{U} is called the *convex hull* of \mathcal{U} .

Example I.1. A triangle spanned by three points $u_0, u_1, u_2 \in \mathbb{R}^2$ is the convex hull of these points.

We can now define a k -simplex.

Definition I.4 (*k-simplex*). A *k-simplex* $\sigma = [u_0, u_1, \dots, u_k]$ is the convex hull of the $k + 1$ affinely independent points $u_0, u_1, \dots, u_k \in \mathbb{R}^d$. One calls k the *dimension* of the simplex.

Example I.2. In Fig. 1, we show examples of simplices for the first few dimensions: a point is a 0-simplex, an edge is a 1-simplex, a triangle is a 2-simplex, and a tetrahedron is a 3-simplex.




FIG. 1. From left to right, we show examples of a 0-simplex, a 1-simplex, a 2-simplex, and a 3-simplex. [We adapt these examples and the figure from [2].]

The lower-dimensional simplices from example I.2 are contained in the higher-dimensional simplices, because subsets of affinely independent points are also affinely independent. The lower-dimensional simplices form so-called *faces* of the higher-dimensional objects.

Definition I.5 (*(proper) faces* and *cofaces*). A *face* τ of a k -simplex σ is the convex hull of a subset $\mathcal{V} \subseteq \mathcal{U}$. Additionally, the face is *proper* if the subset relationship is a proper one. If τ is a (proper) face, then σ is called a *(proper) coface* of τ .




FIG. 2. Panels (a), (b), and (c) give examples of simplicial complexes. The collection of simplices in panel (d) is not a simplicial complex. We use colors to indicate 2-simplices.

Remark 1. We use the notation $\tau \leq \sigma$ to denote a face of σ , and we use $\tau < \sigma$ to denote a proper face of σ .

Recalling the building blocks that we described at the beginning of this Supplementary Information, we can ask whether it is only possible to build shapes using 2-simplices (i.e., triangles) or whether one can also combine these simplices with higher-dimensional or lower-dimensional simplices. A (permissible) shape built from a combination of simplices is called a *simplicial complex*. To construct a simplicial complex, one needs to follow a set of minimal rules:

Definition I.6 (*simplicial complex*). A *simplicial complex* is a finite collection of simplices Σ such that

- i. if $\sigma \in \Sigma$ and $\tau \leq \sigma$, then $\tau \in \Sigma$;
- ii. if $\sigma, \tilde{\sigma} \in \Sigma$, then the intersection of both simplices is either the empty set or a face of both.

In Fig. 2, we show several examples of simplicial complexes and one example that is not a simplicial complex. Example (a) illustrates that simplicial complexes are not necessarily the same as simplices. The three edges do not form a 2-simplex, but they do form a simplicial complex that consists of 1-simplices. In examples (b) and (c), all 1-simplices and 2-simplices are connected by 0-simplices. Example (d) is a collection of simplices that violates the definition of a simplicial complex, because the intersection between the two triangles does not consist of a complete edge. Note that any combination of the three simplicial complexes (a), (b), and (c) is also a simplicial complex.

We take the *dimension* of Σ to be the dimension of its highest-dimensional simplex. One can use simplicial complexes to represent topological spaces if there exists a homeomorphism

between the simplicial complex and the topological space. Only then can one be sure that topological properties such as connectedness are preserved.

B. Homology and Betti numbers

Homology is a formal way of quantitatively detecting holes in topological spaces. These holes are quantified by classifying the space that surrounds them. For example, one measures 1-dimensional holes in a torus by considering loops on its surface. One then classifies them into different types according to whether or not they can be deformed into each other by bending and stretching. In this way, one can distinguish a 2-sphere from a torus by capturing the fact it that is possible to contract any loop on the sphere to a point, whereas there are two distinct loops on the torus surface that cannot be deformed continuously into each other. These loops also cannot be contracted to a point, because they surround different holes.

Although homology is not the only formalism that can be used for distinguishing two shapes, it currently has the fastest algorithms for computing it². *Homology groups*, which are topological invariants of a space, and *Betti numbers* (which are derived from them) play a key role in computing homology. Homology groups detect holes in a topological space, and Betti numbers give a way to count the number of holes or distinct loops in that space.

We start constructing homology groups by looking at formal sums of simplices.

Definition I.7 (*p-chain*). Let Σ be a simplicial complex, let p be a given dimension, and let G be an Abelian group. A *p-chain*

$$c = \sum_{i \in I} a_i \sigma_i \tag{1}$$

is a so-called “formal sum”³ of p -simplices in Σ , where $a_i \in G$ are coefficients, σ_i are p -simplices, and I is an index set.

In computational topology, the employed commutative group G is usually $\mathbb{Z}/2\mathbb{Z}$, which has the advantage that one can regard p -chains as subsets of the set of all p -simplices in Σ by assigning the coefficient 1 to simplices that form part of the subset and the coefficient 0 to those that are not in the subset. Moreover, because $\mathbb{Z}/2\mathbb{Z}$ is also a field, one can also think of p -chains as elements of a vector space. We use $\mathcal{C}_p = \mathcal{C}_p(\Sigma)$ to denote the set of all p -chains of a simplicial complex Σ .

One defines the summation of two p -chains, $c = \sum_{i \in I} a_i \sigma_i$ and $c' = \sum_{i \in I} b_i \sigma_i$, on Σ in a componentwise manner:

$$c + c' = \sum_{i \in I} (a_i + b_i) \sigma_i. \quad (2)$$

It then follows that p -chains form an Abelian group. When working with coefficients from $\mathbb{Z}/2\mathbb{Z}$, the sum of two p -chains results in summing all p -simplices in which the two original p -chains differ. The p -simplices that the two p -chains have in common are present in the sum twice, and these contributions vanish by the properties of addition on $\mathbb{Z}/2\mathbb{Z}$.

The following definition will help relate the different p -chain groups of a simplicial complex.

Definition I.8 (*boundary of a p -simplex*). The *boundary* $\partial_p \sigma$ of a p -simplex $\sigma = [u_0, u_1, \dots, u_p]$ is the formal sum of its $(p - 1)$ -dimensional faces:

$$\partial_p \sigma = \sum_{j=0}^p [u_0, \dots, \hat{u}_j, \dots, u_p], \quad (3)$$

where \hat{u}_j denotes the point that is not included when spanning the simplex.

We can extend this definition to p -chains in a natural way by defining the boundary of a p -chain $c = \sum_{i \in I} a_i \sigma_i$ as $\partial c = \sum_{i \in I} a_i \partial \sigma_i$.

We can now construct a family of boundary homomorphisms ∂_p between the different groups of p -chains of a simplicial complex by mapping p -simplices to their boundaries:

$$\begin{aligned} \dots &\xrightarrow{\partial_{p+2}} \mathcal{C}_{p+1} \xrightarrow{\partial_{p+1}} \mathcal{C}_p \xrightarrow{\partial_p} \mathcal{C}_{p-1} \xrightarrow{\partial_{p-1}} \dots \xrightarrow{\partial_1} \mathcal{C}_0, \\ c &\longmapsto \partial c. \end{aligned}$$

By construction, taking the boundary of a p -chain satisfies the property $\partial_p(c + c') = \partial_p c + \partial_p c'$. Therefore, ∂_p is a homomorphism. Such a sequence of chains and homomorphisms is called a *chain complex*. One can show^{2,4} that the following theorem holds for boundary homomorphisms in a chain complex:

Theorem I.1. *Let $d \in \mathcal{C}_{p+1}$. It follows that*

$$\partial_p \partial_{p+1} d = 0. \quad (4)$$

For simplicity, we often denote the boundary homomorphism by ∂ . In other words, we omit the specification of p . Two subgroups of $(\mathcal{C}_p, +)$, together with boundary homomorphisms and their property from Theorem I.1, form the main ingredients in constructing the homology group of a simplicial complex.

Definition I.9 (*p-cycle*). A *p-cycle* is an element of $\mathcal{Z}_p = \ker \partial_p$, where $\ker \partial_p$ denotes the kernel of ∂_p .

We denote the set of *p-cycles* as \mathcal{Z}_p , and we observe that $(\mathcal{Z}_p, +)$ is a subgroup of $(\mathcal{C}_p, +)$.

Definition I.10 (*p-boundary*). A *p-boundary* is an element of $\mathcal{B}_p = \text{Im } \partial_{p+1}$, where $\text{Im } \partial_{p+1}$ denotes the image of ∂_{p+1} .

We denote the set of *p-boundaries* as \mathcal{B}_p , and we observe that $(\mathcal{B}_p, +)$ is a subgroup of $(\mathcal{C}_p, +)$.

Using Theorem I.1, one can now relate the cycle and boundary subgroups to each other. From Theorem I.1, it follows that $\partial_p(\text{Im } \partial_{p+1}) = 0$, so $\mathcal{B}_p \subseteq \mathcal{Z}_p$. One can then show that \mathcal{B}_p is indeed a subgroup of \mathcal{Z}_p .

We start with $p = 1$ and note that 1-dimensional loops (i.e., 1-loops, which are called simply “loops” in the main text) behave differently from other edges. Edges are mapped to their end nodes by ∂_1 , but every node in a a a loop occurs as the boundary of two edges and thus sums to 0 over $\mathbb{Z}/2\mathbb{Z}$.

We now have come very close to our goal of being able to count holes of a topological space using *p-dimensional loops* (i.e., *p-loops*). Thus far, we have identified that the boundary subgroup \mathcal{B}_1 has 1-loops, but more generally \mathcal{B}_p can also include the boundaries of higher-dimensional chains. To isolate *p-loops* from *p-boundaries* for any dimension *p*, we define the *p*th homology group of a simplex.

Definition I.11 (*pth homology group*). The *pth homology group* \mathcal{H}_p of a simplicial complex Σ is the quotient group of the group of *p-cycles* (\mathcal{Z}_p) modulo the group of boundaries (\mathcal{B}_p). That is,

$$\mathcal{H}_p = \mathcal{Z}_p / \mathcal{B}_p.$$

An element of \mathcal{H}_p is called a *p-dimensional loop* (i.e., a *p-loop*).

Two *p-cycles* in the *p*th homology group are construed as different if they differ by more than just a boundary. Otherwise, the quotient group treats them as belonging to the same homology class. Every hole of dimension *p* in a simplicial complex is surrounded by at least one *p-loop* in the homology group. Counting the number of classes in \mathcal{H}_p thus gives an estimate of the number of *p-loops* of a simplicial complex. However, *p-loops* that surround

the same hole are counted separately. A solution is to count the minimal number of elements that are needed to generate the group. This motivates the definition of p th Betti number.

Definition I.12 (*p th Betti number*). The p th *Betti number* β_p of a simplicial complex is

$$\beta_p = \text{rank } \mathcal{H}_p.$$

Recall that we are working with coefficients from $\mathbb{Z}/2\mathbb{Z}$. This turns the set of p -cycles into a vector space, so we can think of the homology group \mathcal{H}_p as a quotient vector space. The p th Betti number is then given by the dimension of this vector space. One can interpret the first three Betti numbers (β_0 , β_1 , and β_2) to represent, in turn, the number of connected components, the number of 1-dimensional holes, and the number of 2-dimensional holes in a simplicial complex.

C. Filtrations

We first define what we mean by a “subcomplex” of a simplicial complex Σ .

Definition I.13 (*subcomplex* of a simplicial complex). A *subcomplex* of a simplicial complex is a subset of simplices that satisfy the properties of a simplicial complex.

We can now build sequences of simplicial complexes that form subcomplexes of each other.

Definition I.14 (*filtration*). A *filtration* of a simplicial complex Σ is a nested sequence of subcomplexes starting with the empty complex \emptyset and ending with the entire simplicial complex:

$$\emptyset = \Sigma_0 \subseteq \Sigma_1 \subseteq \Sigma_2 \subseteq \cdots \subseteq \Sigma_k = \Sigma. \quad (5)$$

Observe that one can define natural inclusion maps $i_j : \Sigma_j \hookrightarrow \Sigma_{j+1}$ along the filtration. The inclusion maps induce maps $i_j^* : \mathcal{H}_p(\Sigma_j) \rightarrow \mathcal{H}_p(\Sigma_{j+1})$.

In a filtration, one is interested in determining (1) when prominent features (e.g., a homology class) first appear and (2) if and when those features disappear.

Definition I.15 (*birth and death* of a homology class, *persistence*). A homology class $h \in \mathcal{H}_p(\Sigma)$ is *born* at Σ_m if h is an element of $\mathcal{H}_p(\Sigma_m)$ but is not in the image of the map $i_{m-1}^* : \mathcal{H}_p(\Sigma_{m-1}) \rightarrow \mathcal{H}_p(\Sigma_m)$.

A homology class $g \in \mathcal{H}_p(\Sigma)$ *dies* entering Σ_n if g is an element of $\mathcal{H}_p(\Sigma_{n-1})$ but is not in the image of the map $i_{n-1}^* : \mathcal{H}_p(\Sigma_{n-1}) \rightarrow \mathcal{H}_p(\Sigma_n)$ (i.e., $i_{n-1}^*(g) = 0$).

Let m_h denote the filtration step at which h is born, and let n_h denote the filtration step at which h dies. One then defines the *persistence* of a homology class $h \in \mathcal{H}_p(\Sigma)$ as

$$p_h = n_h - m_h .$$

REFERENCES

- ¹B. Stolz, *Computational Topology in Neuroscience*, Master's thesis, University of Oxford, <http://www.math.ucla.edu/~mason/research/Dissertation-stolz2014-Corr.pdf> (2014).
- ²H. Edelsbrunner and J. L. Harer, *Computational Topology* (American Mathematical Society, Providence R. I., 2010).
- ³Reddit, "Formal definition of the term *formal sum*," (2013), https://www.reddit.com/r/math/comments/18p6rj/formal_definition_of_the_term_formal_sum/.
- ⁴F. H. Croom, *Basic Concepts of Algebraic Topology* (Springer, New York, Heidelberg, Berlin, 1978).