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

 ¹⁾Mathematical Institute, University of Oxford, Oxford OX2 6GG, United Kingdom
 ²⁾CABDyN Complexity Centre, University of Oxford, Oxford OX1 1HP,

United Kingdom

(Dated: 3 May 2016)

We use topological data analysis to study "functional networks" that we construct from time-series data from both experimental and synthetic sources. Specifically, we use persistent homology in combination 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 consists of biological data in the form of functional magnetic resonance imaging (fMRI) data that was acquired from human subjects during a simple motor-learning task. Our second example uses time-series output from networks of coupled Kuramoto oscillators. 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 play a more vital role than persistence in the interpretation of topological features, even though the latter is used more commonly to distinguish between signal and noise^{1,2}. We find that in particular, persistent homology can detect differences in synchronisation patterns in our data sets over time giving insight on changes in community structure in the networks, and on increased synchronisation between brain regions forming loops in the functional network during motor-learning. For the motor-learning data we also observe that persistence landscapes reveal that the majority of changes in the loops of the network takes place on the second of three days of the learning process.

Keywords: Persistent homology, networks, time series, fMRI, persistence landscapes, functional networks, functional brain networks, nonlinear oscillators, Kuramoto model, dynamical systems on networks

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

I. INTRODUCTION

Analysis of neuronal networks is crucial for understanding the human brain^{11–15}. The human brain consists of approximately 100 billion neurons, whose major task is to receive, conduct, and transmit signals. Every neuron consists of a cell body and one long axon, which is responsible for propagating signals to other cells¹⁶. Neurons, or on a larger scale, different brain regions can be regarded as nodes of a network, whose edges represent either structural or functional connection between the nodes. Looking at neuronal data using a network approach allows one to use mathematical tools 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, or compare the structure of neuronal interactions to those of other complex systems. For example, data analysis using tools from graph theory has lead to the insight that the brain has an underlying modular structure, in that it is organized in small subunits, which are able to carry out specific functions while minimally influencing other parts of the network^{11,14,17}.

Standard methods from network theory are based on pairwise connections, which one can nevertheless use to study microscale, mesoscale, and macroscale structures¹⁸. An alternative approach to study networks¹⁹ is to use methods from *computational topology*, which explicitly incorporates "higher-order" structures beyond pairwise connections and which includes algorithmic methods for understanding topological invariants such as connectedness, loops, or holes in high-dimensional data structures^{3–5}. While one can represent such higher-order structures using other methods such as hypergraphs²⁰, see for example²¹, these do not give information about shape or scale of these structures in the network.

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 applications⁷ in diverse topics ranging from granular materials⁸ and contagions on networks²² to path $planning^{23}$ and collective behavior in animals²⁴. In particular, *persistent homology*, a mathematical formalism to explore the persistence of topological structures in data sets, has become increasingly prominent in neuroscience in the last few years^{9,10}. Among other applications, it has been used to determine differences in brain networks of children with hyperactivity disorders and autism spectrum²⁵, to study the effect of the psychoactive component of "magic mushrooms" (psilocybin mushrooms) on functional brain networks of humans²⁶, to analyze covariates that influence neural spike-train data²⁷, and to study structural and functional organization of neural microcircuits²⁸. Other neuronal applications have included the consideration of place cells in the hippocampus of rats during spatial navigation $^{29-31}$, the analysis of mathematical models of transient hippocampal networks³², and a demonstration that topological features of networks of brain arteries in humans are correlated with their age³³. Persistent homology is however not the only topological method that has been used to study the human brain. For example, Zeeman³⁴ has used tolerance spaces and Vietoris homology theory to study aspects of visual perception.

In our investigation, we use time-series data to construct so-called *functional net-works*^{11,12,35,36}. Functional brain networks consist of a set of nodes, usually brain regions, and a set of weighted edges between nodes whose time series are similar according to a chosen measure. The term is used in contrast to "structural networks", which refer to underlying

physical 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 their functional networks. We adopt the term "functional networks" 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 bevy of other applications (e.g., finance³⁷, voting among legislators³⁸, and climate³⁹). Importantly, the times series can come either from empirical data or from the output of a dynamical system, and the latter is helpful for validating methods for network analysis⁴⁰. In our paper, we will consider times 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 and j. There are many different ways to measure similarity of times series 12,41,42 , and that can be a major issue when it comes to interpreting results. Comparing the networks that are arise from different similarity measures is beyond the scope of our work, so we will simply use two common measures of time-series similarity (pairwise synchrony, wavelet coherence). 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 graphs¹². To do this, one typically applies a global threshold $\tau \in \mathbb{R}^+$ to the 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

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

The choice of threshold has a large influence on the resulting graph, and it thereby exerts a major influence on the structure of the associated graph¹². Some approaches to address this issue include determining a single "optimal" threshold, thresholding the weighted adjacency matrix at different values^{43,44} and examining the network properties as a function of thresh-

old, or not thresholding at all and considering the weighted adjacency matrix itself^{12,14}. If one is thresholding the 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 networks⁴⁵. 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.

Note that while we use functional networks and persistent homology for our examples, one can also study coupled time series using a variety of different approaches^{46–48}.

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, and persistence landscapes. In Section III, we introduce our first example, the Kuramoto model of nonlinearly coupled oscillators, and present the 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 task. We discuss our findings in Section V, and we provide a more mathematical introduction to persistent homology in Appendix A.

II. PERSISTENT HOMOLOGY

Persistent homology $(PH)^{3-5}$ 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 lense conveys topological features inside data at a different resolution. One then construes structures that persist over a range of different lenses to represent a significant feature of the data. Structures that are observed only through a small number of lenses are commonly construed as noise^{1,2}, in particular in settings where the data is sampled from a manifold. For real world data, this has so far not been shown to hold statistically and we will see later in the paper that there are cases where such structures represent important properties and possibly geometry of the data.

In this section, we provide an intuitive introduction to the mathematical concepts behind PH. In Appendix A, we give a mathematically more rigorous introduction (including precise definitions).

A. Simplicial complexes

One can study the properties of a topological space^{49,50} 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. 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 us to capture the topological properties of more complicated spaces, such as a 2-sphere or the surface of the earth. In this case, we need a simple object that carries the information that the space is connected but also encloses a hole. For example, we 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 contains k + 1 independent nodes: a point \bullet is a 0-simplex, an edge $\bullet \bullet \bullet$ is a 1-simplex, a triangle is a 2-simplex, and a tetrahedron is a 3-simplex. 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 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 the simplicial complex, we demand that all of its faces are also contained in the simplicial complex. We also demand that two simplices that are part of a simplicial complex only intersect in common faces. In Figure 1 we show examples of simplicial complexes.

We take the *dimension* of a simplicial complex to be the dimension of its highestdimensional simplex. One can use simplicial complexes to represent topological spaces if there exists a continuous deformation that can stretch and bend the simplicial complex into



FIG. 1. 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 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 again a simplicial complex.

the topological space and vice versa. 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, then 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 might be connected or even form a loop. However, one may have different levels of interest in different topological features (e.g., a collection of edges surrounding a hole or void could be more interesting than individual edges), and in particular we seek 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 that gives insight into the nature of a topological space. 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 $p \in \{0, \ldots, k\}$. For dimension 1 we call the elements of the homology group "loops". The elements of the homology group can be divided into different *homology classes*, which each represent a hole in the topological space. For example in dimension 1, loops in the same homology class all surround the same 1-dimensional hole. 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.

C. Filtrations

Although homology gives information about a single simplicial complex, it is typically more relevant to study topological features across sequences (called *filtrations*) of simplicial complexes. A filtration^{1,2,4} of a simplicial complex Σ is a sequence of embedded simplicial complexes,

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

starting with the empty complex and ending with the entire simplicial complex. One can use homology to study topological features such as 1-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 Σ_m is the first simplicial complex in the filtration to contain the feature. Similarly, a topological feature dies in Σ_n if it is present in Σ_{n-1} but not 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 lifetime³ within 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 the consideration of computational scaling.

1. Weight rank clique filtration

Although we focus on network data, note that PH has been applied much more often to data in the form of point clouds^{1,2}. The simplest way to create a sequence of embedded

graphs (e.g., a filtration) from a weighted network is to filter by weights⁵¹. To do this, one creates a sequence of embedded (binary) graphs by ranking all edge weights ν_t in descending order. In filtration step t, we only retain edges whose weight is greater than or equal to ν_t . To construct the filtration, one repeats this procedure until the graph is complete in the last step. Using this method, only 0-simplicies (i.e., nodes) and 1-simplices (i.e., edges) are present in the filtration. The weight rank clique filtration (WRCF) ⁵², which we will use in our analysis and which has been applied previously for examining weighted neuronal networks^{26,31,52}, 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 $\{\nu_1, \ldots, \nu_{\tau}\}$, where $\nu_1 = \nu_{\text{max}}$ and $\nu_{\tau} = \nu_{\text{min}}$, where τ is the number of distinct weights in the graph.
- 3. In filtration step t, threshold the graph at weight ν_t , where $t \in \{1, \ldots, \tau\}$, 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. Hence, 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*¹. A barcode for a given dimension is a collection of intervals $\{b_l, d_l\}_{i=1}^m$, where every interval represents a topological feature l of the given dimension (Examples of such features include connected components and loops), where b_l gives the birth time of feature l and d_l gives its death time with respect to the filtration step. The length of the bar measures the persistence of the feature. In Figure 2, we show an example of a weight rank clique filtration and the corresponding barcode.



FIG. 2. Example of a weight rank clique filtration and the corresponding 0-dimensional and 1dimensional 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 representing the original components dies in the barcode, while 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 into two separate components in filtration step 1, and they 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 that is born in filtration step 2 is first divided 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.

D. Persistence landscapes

As an alternative topological summary to barcodes, one can use *persistence land-scapes*^{53,54}, which consist of piecewise linear functions in a separable Banach space. For a given barcode interval (b, d), one defines the function

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

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

$$\lambda_q : \mathbb{R} \to \mathbb{R} ,$$

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

If the qth 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 the barcode. To obtain the 1st persistence landscape, one peels away this topmost layer of peaks and again 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 of the functions $\{\lambda_q\}$. Even though persistence landscapes visualise the same information as barcodes and one can construct a one to one correspondence between the two objects, they have distinct advantages over the latter. For example, one can calculate the unique average landscape for a set of persistence landscapes. This is not possible for barcodes as they are not elements of a Banach space. Note that for an average landscape it is therefore not possible to find a corresponding average barcode. One can also define L^p distances between two (average) landscapes and thereby use a range of statistical tools⁵³. This allows one to compare multiple groups of barcodes by calculating a measure of pairwise similarity between them. Such calculations have been used to study conformational changes in protein binding sites⁵⁵.

E. Computational tools

For our persistent-homology calculations, we use MATLAB codes that we construct using JAVAPLEX^{56,57}, a software package for persistent homology. For the WRCF, we also use a maximal clique-finding algorithm from the Mathworks library⁵⁸ based on the Bron–Kerbosch algorithm, which the most efficient algorithm known for this problem. For the analysis and interpretation of our barcodes we apply the PERSISTENCE LANDSCAPES TOOLBOX⁵⁴.

III. EXAMPLE I: COUPLED KURAMOTO OSCILLATORS

A. The Kuramoto model

The Kuramoto model^{59–63} is a well-studied model for a set of coupled limit-cycle oscillators with distinct natural frequencies that are drawn from a prescribed distribution. The model was developed in the 1970s to understand collective synchronization in a large system of oscillators. It has subsequently been used as a toy model by many neuroscientists (as well as many other scholars), as some of the characteristics of its synchronization patterns resemble some of the ones in neuronal communities⁶⁴⁻⁶⁶. The Kuramoto model and its generalizations have also been applied to numerous other applications in chemistry, biology, and other disciplines^{61,62,67}.

For the case of all-to-all coupling, the Kuramoto model is most commonly written as 60,62

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

where θ_i denotes the phase of oscillator *i*, the parameter ω_i is its natural frequency, $K \ge 0$ 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 equations are bounded as $N \to \infty$. 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 (5) is invariant under translation of θ_i]. The parameter ω_i then denotes the deviation from the mean frequency.

We simulate the basic Kuramoto model for all-to-all coupled oscillators using the Runge– Kutta MATLAB solver ODE45 (with integration time $T_{\text{max}} = 10$ and time step $\Delta t = 0.02$).

We also adapt Eq. (5) to create a network of N oscillators by introducing binary coupling between the oscillators^{40,47,61,62,68}. We consider the following generalized version of Eq. (5):

$$\frac{d\theta_i}{dt} = \omega_i + \sum_{j=1}^N \kappa A_{ij} \sin(\theta_j - \theta_i), \quad i \in \{1, \dots, N\},$$
(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 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^{62} separate communities⁶⁹ 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 oscillators outside the community. As in^{40} , we choose the coupling strength to be $\kappa = 0.2$, the number of oscillators to be N = 128, and the natural frequency $\omega_i \sim \mathcal{N}(0, 1)$ (i.e., it has a Gaussian distribution with mean 0 and standard deviation 1). However, our network architecture differs somewhat from that in^{40} , where every oscillator had at least 13 connections inside its community and at least 1 connection outside its community.

We observe the system for M = 500 time steps in total (including the initial time step) and obtain time series $\tau_i = (\theta_i(t_0), \ldots, \theta_i(t_{499}))$ as the output of the model for every oscillator θ_i . Kuramoto oscillators with a similar imposed community structure have previously been shown to initially synchronise rapidly within their communities, followed by a phase of global synchronisation within the whole network⁴⁰. To study the dynamics of the coupled Kuramoto oscillators, we therefore partition the time series into two separate time regimes $\hat{k} = 1, 2$ that consist of 250 time steps each.

To quantify the pairwise synchrony of two oscillators i and j, we use the local measure^{40,68}

$$\phi_{ij}^{\hat{k}} = \left\langle \left| \cos \left(\tau_i^{\hat{k}} - \tau_j^{\hat{k}} \right) \right| \right\rangle \,, \tag{7}$$

where the angular brackets indicate that we average over 20 simulations. In every simulation we choose the initial values for ω_i from a uniform distribution on $[0, 2\pi]$ and draw the natural frequencies 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 both time regimes. We also study the network based on one full time regime consisting of 500 time steps. In analogy to neuronal networks, we call these networks "functional networks."

In Figure 3, we illustrate our pipeline for creating a functional network from the Kuramoto model.

B. Null models for the Kuramoto data

To assess whether our observations are a direct result of the 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 reassign the order of the time series for every oscillator according to a uniform distribution before computing the similarity measure with Eq. (7). The second null model is based on creating surrogate data using a discrete Fourier transformation. This approach⁷⁰ has the advantage of preserving not only the mean and the variance of the original time



FIG. 3. We impose a structural network for the Kuramoto model by grouping the oscillators into 8 separate communities. Oscillators are coupled predominantly to other oscillators in their community and are coupled only very sparsely to oscillators outside their community. We then use the time-series output of the model to create a functional network based on the similarity of the time series.

series but also the linear autocorrelations and cross-correlations between the different time series.

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

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

of a time-series vector $\tau = (\theta(t_0), \dots, \theta(t_{\mu}))$ of length μ . In our case $\mu = 250$ and $\mu = 500$, depending on whether we are looking at two different time regimes or just one.

We then construct surrogate data by multiplying the Fourier transform by phases a_n chosen uniformly from the interval $[0, 2\pi)$ and satisfying a 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 $\sigma = (\sigma_1, \ldots, \sigma_\mu)$ are thus given by

$$\sigma_m = \frac{1}{\sqrt{\mu}} \sum_{n=0}^{\mu-1} e^{ia_n} \hat{\tau}_n e^{-\frac{2\pi i n m}{\mu}} \,. \tag{9}$$

We call this null model the "Fourier null model." Both null models have been used previously on time series output of coupled Kuramoto oscillators and haven been observed to display different dynamics than the Kuramoto oscillators^{40,71}.

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

We apply the weight rank clique filtration to functional networks created from the output of two time regimes of the Kuramoto model, one time regime for the Kuramoto model, the simple null model, and the Fourier null model. We run the filtrations up to filtration step 1800 and 2000 for the two time regimes and up to filtration step 1100 for the cases in which we only consider one time regime. The total number of edges in the network, and thus total number of possible filtration steps is 8128. Thus this corresponds to an edge density in the networks of 22%, 25% and 14%, respectively, and it thresholds them approximately in the middle of the range of the edge-weight values. In the thesis that presents a precursor of this work, Stolz⁷² applied PH to networks created from the Kuramoto model, such an example was subsequently studied using Betti curves⁷³.

In the first row of Figure 4, 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., the second half of the dynamics) for the weight rank clique filtration of the Kuramoto model. For both time regimes, the barcodes include several very short-lived bars in the beginning of the barcode 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. We find that the 1-loops corresponding to these short bars are all formed within the highly synchronized communities. In fact, in time regime 1 the first 44 bars in the barcodes represent in-community loops, while in time regime 2 only the first 28 bars represent in-community loops. As these strong intra-community edges are added to the simplicial complexes, they start to cover the 1-loop with triangles (i.e., 2-simplices), and the loops disappear from the filtration.

In the second row of Figure 4, we show the persistence landscapes that we construct from the 1-dimensional barcodes, although we ignore any infinitely-persisting bars in the barcode (we also studied persistence landscapes including the infinite bars as features with a death time that corresponds to the maximum filtration value but did not find that this lead to additional insights).

As we expect from the barcodes, the landscapes show a group of small peaks early in the filtration for both time regimes, which cover a longer filtration range in the second time regime before more persistent loops appear. We also observe that some of these short peaks



FIG. 4. Dimension-1 barcodes and persistence landscapes for the weight rank clique filtration for the two time regimes of the Kuramoto model time series. The horizontal axis in both the barcodes and landscapes represents the filtration steps. 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 infinitelypersisting bars in the barcodes). The short peaks in the beginning of the filtration in the persistence landscapes indicated by the red ellipses represent loops formed within communities.

become larger, with peaks reaching heights around 100, whereas the peaks are half as high (less persistent) in the first time regime. Other peaks in this group on the other hand stay at a similar hight as in the first time regime.

In general, the persistence landscapes show more persistent 1-loops for the second time regime, and there also appears to be a clearer separation between the group of the very early short peaks and a second group of medium sized peaks towards the end of the filtration. For this second group of medium sized peaks, we can see a larger absolute increase in persistence in the second time regime than for the shorter peaks in the beginning of the filtration. These observations reflect the dynamics of the two time regimes in the Kuramoto model⁴⁰: In time regime 1 we have strong synchronisations within the communities, which are reflected by the appearance of short-lived in-community 1-loops in the beginning of the filtration corresponding to the short peaks in the persistence landscapes. In the second time regime, the global synchronisation increases and in addition to in-community loops some of the peaks in the beginning of the filtration, the gap between the initial peak group and the group of medium sized peaks in the end

of the filtration increases. In general we see an increase in the persistence of the peaks in the landscapes due to the stronger synchronisation between the communities. We note that these observations are much easier visualised by the persistence landscapes than they were in the barcodes. We also compare the Kuramoto model to the two null models. To do this, we construct a functional network by considering a single time regime that consists of all 500 time steps. In Figure 5, 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 in-community synchronization for the Kuramoto times series than for the null models, as there is a very distinct group of short peaks in the beginning of the filtration similar to the analysis of the two separate time regimes.



FIG. 5. (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.

Again, the corresponding loops are within communities. The peaks in the Kuramoto landscape appear to be separated from a second group of short peaks further along in the filtration. In between the two peak groups we observe two significantly higher peaks corresponding to more persistent loops. These persistent loops appear to be formed by connections between different communities. For both null models, we also observe groups of short peaks at the beginning of the filtration, but these are less persistent and less clearly separated from other peaks than for the Kuramoto model. Indeed, we see that for the Fourier null model, which exhibits a much weaker in-community synchronisation than the simple null model, we don't see any separation at all. Moreover the persistence landscape for the Fourier null model appears to be more "noisy" as the majority of the peaks in the landscape are of similar persistence and appear in similar areas of the filtration.

For both null models we see more medium-size and long persisting features compared to the Kuramoto data. These features appear in parts of the filtration where the Kuramoto data exhibits 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.

For the Kuramoto model we find that PH can detect the dynamics of the system and the landscapes show a clear difference between the Kuramoto model and the null models. In fact, we can even distinguish between the two null models. In contrast to^{1,2}, we do not find that the persistence of the topological features in this case distinguishes between signal and noise: We find that the short bars in the beginning of the filtration of the Kuramoto model carry important information on the dynamics, while the medium persistent peaks in the Fourier null model are a symptom of the weaker in-community and stronger intercommunity synchronisation in the Fourier null model. We therefore suggest that the position of features in the barcode is equally important for their interpretation in this context as their persistence. The landscapes alone do however not provide enough information to assess the dynamics. It is only when we combine them with the information on the loops that are represented by certain groups of peaks that we can obtain information on in- and intercommunity synchronisation.

IV. EXAMPLE II: TASK-BASED FMRI DATA

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

We use a data set of time-dependent functional brain networks from experiments that were first analyzed in⁷⁴. It consists of functional magnetic resonance imaging (fMRI) timeseries data from 20 healthy subjects who are undertaking a motor-learning task on three days (during a five-day period). An "atlas" of 112 brain areas was monitored while the subjects were performing a simple motor-learning task (similar to a musical sequence), which they executed using four fingers of their non-dominant hand. The fMRI data consists of 2000 time points for each day and these were subsequently used to create functional networks based on the wavelet coherence between the activity of every pair of brain regions. 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 correlation amount that one expects to occur at random) were set to 0. The other matrix elements were retained.

These networks were previously studied using community detection⁷⁴. The results indicated that there is a significant segregation of the nodes in the functional networks into a small number of different communities with densely-weighted connections inside the communities and sparsely-weighted connections to nodes in other communities. Within these communities, certain nodes appeared to remain in the same community during the experiment, whereas others (the "flexible" ones) often switch between different communities.

Networks from a similar experiment but with medium-term learning including training sessions have also been studied^{75,76}. These networks exhibited a noticeable core–periphery organization, with the sensimotor and visual regions of the brain being grouped in a temporal stiff core of nodes, whose connectivity in contrast to the nodes in the flexible periphery region do not change over the course of the learning task⁷⁵. It was later also shown that the interaction between primary and secondary sensorimotor regions and the primary visual cortex decreases as the regions presumably become more autonomous with task practice⁷⁶.

As we observed short-lived 1-dimensional loops in the beginning of the filtrations for the Kuramoto model in a simulated community setting, we will explore whether the fMRI data exhibits similar features during the 3 observation days.

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

We run the weight rank clique filtration until filtration step 2600, where 42% of the edges are present in the network. We find that larger filtration steps lead to very long computational times. We construct persistence landscapes for dimension 1 (omitting infinitely persisting 1-loops). In Table 6, we summarize our results for one particular subject and the whole data set. The subject was chosen as an example representative for the particular landscape features which we are observing.

Similar to the Kuramoto case, we find a group of small peaks at the beginning of the filtration (between filtration steps 1 and 200). We can see this group very clearly both by zooming into the landscape of individual subjects and in the average landscape, where the hight of the peaks is only slightly lower than for the peaks in the individual landscapes; this indicates that there is a group of short peaks arising in the beginning of the filtration in



FIG. 6. Persistence landscapes for dimension 1 of the weight rank clique filtration 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.

the majority of the barcodes. We expect these peaks to be associated with the communities in the networks, which have been observed previously using other methods⁷⁴. We observe in particular, that these short peaks undergo changes on day 2: during filtration steps 20 to 60 we see that some of the peaks that are present in the landscapes for day 1 and 3 vanish while between filtration step 80 and 200 we observe more persistent peaks than on the other two days. We suggest that this implies a change of the community structure taking place on day 2 with either very strong synchronisation strengths in some of the communities leading to very short lived 1-loops or very strong individual differences between the subjects leading to the vanishing of peaks in the average landscapes for the first 50 filtration steps. The more persistent peaks on day 2 could either represent persistent loops between different communities or loops occurring due to sparse in-community connections.

We calculate the pairwise L_2 -distances between all dimension-1 persistence landscapes. The L_2 distance has been previously used to compare persistence landscapes in a biological context⁵⁵. We create distance vectors, which we use as an input for k-means clustering and average linkage clustering for k = 3. Using both methods, we obtain the same results: we find that 9 of the 20 distance vectors corresponding to persistence landscapes from day 1 cluster in one common group (together with very few landscapes from days 2 and 3), while 11 and 10 landscapes from days 2 and 3 cluster together into a separate group. We summarize

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

TABLE I. Results for k-means clustering and average linkage clustering of pairwise L_2 -distance vectors of persistence landscapes for k = 3.

our results in Table I.

We also consider the average landscapes based on dimension 1 of weight rank clique filtration steps 1–2600 and calculate the L_2 -distances between them. We show the results of these calculations in Figure 7.



FIG. 7. Visualization of average persistence landscapes for days 1, 2, and 3. We observe that the distance between the landscape for day 1 and the other two landscapes is larger than the distance between the landscapes for days 2 and 3. (The L_2 distances between them are 5200 between days 1 and 2, 5000 between days 1 and 3, and 3500 between days 2 and 3.) We also observe a shift to the left of the landscape peak during the three days, indicating that the more persistent 1-loops in these networks arise earlier in the filtration for the later days. In other words, they are formed by edges with a higher edge weight and indicate that there is stronger synchronization between the brain regions involved.

The large distances between day 1 and the subsequent days of the experiment indicate

that the weight rank clique filtrations are able to detect changes in the functional networks across the filtration range, most of these changes happening between the first and the second day. We also observe that the main peak of the average landscapes is shifting to the left over the course of the three days. This implies that the edge weights between the brain regions giving rise to persistent 1-loops increase over the three days due to stronger synchronisation. Brain regions synchronising in a 1-loop in a network could be an indication of an interesting neurobiological communication pattern that in this case also gets stronger over the course of the learning process. We suggest to study the brain regions involved in these loops in order to investigate their biological role in learning a motor-learning task and in particular to investigate changes in the loops occurring on day 2.

As in other biological contexts where PH has been applied successfully and lead to new insights, for example in^{26,28–31,33}, we find that PH and WRCF in particular lead to insights about the dynamics of our studied systems. We not only find that we can detect symptoms of previously observed community segregation, but we also find differences between a set up with a strong community structure such as in the Kuramoto model and weakly synchronised communities in the null models. For the task-based fMRI data we also find that we can detect changes in the community structure over the three days as well as changes in the 1-dimensional loops appearing in the functional network. The majority of these changes appear to take place on the second day of the learning task. We in particular observe effects of stronger synchronisation between brain areas forming loops in the network throughout the three days of the motor-learning task. The findings on the 1-dimensional loops in the network provide novel insight and would be difficult to study using other methods.

V. CONCLUSION AND DISCUSSION

We have illustrated applications of persistent homology to functional networks constructed from the 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 1-loops occur at the beginning of the filtration. Although this is commonly construed as noise in topological data analysis^{1,2}, 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 we even found that more persistent features appeared to be linked to a network with a weak intra-community synchronisation (Fourier null model) and represented noise. Our results match previous observations for synthetic examples where barcodes consisted of short intervals which would commonly be regarded as noise, but the differences between the corresponding persistence landscapes for the various spaces were statistically significant⁵⁴. For weighted networks, we therefore suggest that when using a filtration based on edge weights, the birth and death times with respect to the filtration of features such as 1-loops needs to be considered in addition to their persistence to be able to decide whether they should be considered as noise or signal. In particular, loops appearing early in a filtration could give important insights on the geometry of the studied data.

We also found, using both average persistence landscapes and when measuring the pairwise distances between separate landscapes, that persistence landscapes for dimension 1 of the weight rank clique filtration were able to capture the changes in the studied functional brain networks during the process of learning a 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 to the full filtration length) persistent 1-loops are able to capture changes in the network, so it is not always necessary to consider a full weigh rank clique filtration to study the dynamics of a system. This observation is similar to a finding in Bendich *et al.*³³, who observed that medium-scale barcode features were significant in distinguishing human brain artery networks from different age groups. This again suggests that the persistence should not be the only measure of signal versus noise when applying PH. Moreover, we found that the persistent features that dominate the middle part of the filtration appear in earlier filtration steps on days 2 and 3 of the experiment than they do on day 1, which further points to interesting dynamics in synchronisation patterns in the network being captured by medium persistent bars in the middle of the barcode.

Using methods from topological data analysis for network approaches has the advantage that these methods are mathematically well founded and generalisable. It is however necessary to include information on specific nodes forming part of topological features for biological interpretation. Moreover, the interpretation of the results and importance of persistence versus position of a topological feature in the barcode could change depending on the filtration that is used. Note that different topological features can also have different relevance for different dynamical systems. For example, the occurrence of many medium persistent features in the persistence landscape for the Fourier null model was a symptom of the weak synchronisation in the communities while for the task-based fMRI data the medium persistent bars captured increasing synchronisation in 1-loops. We suggest to apply WRCF to different synthetic networks with underlying communities, for example using plantedpartition models, to investigate this further and include both persistence and position of topological features in the analysis. It would also be beneficial to combine topological tools with additional methods, such as persistence images⁷⁷, to determine the exact topological features responsible for the detected differences between the persistence landscapes of the different network groups.

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

VI. ACKNOWLEDGEMENTS

The experimental data was originally collected by Nicholas F. Wymbs and Scott T. Grafton through funding from Public Health Service Grant NS44393, and we thank Nicholas and Scott for access to the data. We thank Danielle Bassett for help in providing use of the data, use of her MATLAB code when we were debugging our code, and helpful discussions. We also thank Pawel Dłotko for his help with the PERSISTENCE LANDSCAPES TOOLBOX and for providing us with new versions of the code during our work. BJS thanks the Berrow foundation for funding during her MSc degree. BJS further gratefully acknowledges the EPSRC, and MRC (EP/G037280/1) as well as F. Hoffmann-La Roche AG for funding her doctoral studies. HAH acknowledges funding from EPSRC Fellowship EP/K041096/1.

REFERENCES

- ¹R. Ghrist, "Barcodes: The persistent topology of data," Bulletin of the American Mathematical Society **45**, 61 75 (2008).
- ²G. Carlsson, "Topology and data," Bulletin of the American Mathematical Society **46**, 255 308 (2009).

- ³H. Edelsbrunner, D. Letscher, and A. Zomorodian, "Topological persistence and simplification," Discrete and Computational Geometry **28**, 511 533 (2002).
- ⁴H. Edelsbrunner and J. L. Harer, "Persistent homology a survey," Contemporary mathematics **453**, 257 282 (2008).
- ⁵H. Edelsbrunner and J. L. Harer, *Computational Topology* (American Mathematical Society, Providence R. I., 2010).
- ⁶R. Ghrist, *Elementary Applied Topology*, 1st ed. (Createspace, 2014) available at https://www.math.upenn.edu/ ghrist/notes.html.
- ⁷N. Otter, M. A. Porter, U. Tillmann, P. Grindrod, and H. A. Harrington, "A roadmap for the computation of persistent homology," (2015), arXiv: 1506.08903.
- ⁸M. Kramár, A. Goullet, L. Kondic, and K. Mischaikow, "Persistence of force networks in compressed granular media," Physical Review E **87**, 042207 (2013).
- ⁹C. Curto, "What can topology tell us about the neural code?" (2016), http://www.ams.org/journals/notices/201601/rnoti-CVR3.pdf.
- ¹⁰C. Giusti, R. Ghrist, and D. S. Bassett, "Two's company and three (or more) is a simplex: Algebraic-topological tools for understanding higher-order structure in neural data," (2015), arXiv:1601.01704v2.
- ¹¹E. T. Bullmore and O. Sporns, "Complex brain networks: graph theoretical analysis of structural and functional systems," Nature Reviews **10**, 186 198 (2009).
- ¹²E. T. Bullmore and D. Bassett, "Brain graphs: Graphical models of the human brain connectome," Annual Review of Clinical Psychology 7, 113 – 140 (2011).
- ¹³O. Sporns, "Contributions and challenges for network models in cognitive neuroscience," Nature Reviews Neuroscience **17**, 653 – 660 (2014).
- ¹⁴D. Papo, M. Zanin, J. A. Pineda-Pardo, S. Boccaletti, and J. M. Buldú, "Functional brain networks: Great expectations and hard times and the big leap forward," Philosophical Transactions of the Royal Society B **369**, 25180303 (2014).
- ¹⁵D. Papo, J. M. Buldú, S. Boccaletti, and E. T. Bullmore, "Complex network theory and the brain," Philosophical Transactions of the Royal Society B **371**, 20130520 (2014).
- ¹⁶B. 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).
- ¹⁷E. T. Bullmore and O. Sporns, "The economy of brain network organization," Nature Reviews Neuroscience **13**, 336 – 349 (2012).

- ¹⁸M. E. J. Newman, *Networks: An Introduction* (Oxford University Press, Oxford, 2010).
 ¹⁹In the present paper, we will 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.
- ²⁰B. Bollobás, *Modern Graph Theory* (Springer, New York, 1998).
- ²¹D. 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).
- ²²D. 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," Nature Communications 6, 7723 (2015).
- ²³S. Bhattacharya, R. Ghrist, and V. Kumar, "Persistent homology for path planning in uncertain environments," IEEE Transactions on Robotics **31**, 578 – 590 (2015).
- ²⁴C. M. Topaz, L. Ziegelmeier, and T. Halverson, "Topological data analysis of biological aggregation models," PLoS One **10**, e0126383 (2015).
- ²⁵Lee, Hyekyoung, 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.
- ²⁶G. Petri, P. Expert, F. Turkheimer, R. Carhart-Harris, D. Nutt, P. J. Hellyer, and F. Vaccarino, "Homological scaffolds of brain functional networks," Journal of the Royal Society Interface **11**, 20140873 (2014).
- ²⁷G. Spreemann, B. Dunn, M. B. Botnan, and N. A. Baas, "Using persistent homology to reveal hidden information in neural data," (2015), arXiv: 1510.06629v1.
- ²⁸P. 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," (2016), arXiv: 1601.01580v1.
- ²⁹C. Curto and V. Itskov, "Cell groups reveal structure of stimulus space," PLoS Computational Biology 4, e000205 (2008).
- ³⁰Y. Dabaghian, F. Mémoli, L. Frank, and G. E. Carlsson, "A topological paradigm for hippocampal spatial map formation using persistent homology," PLoS ONE 8, e1002581 (2012).
- ³¹C. Giusti, E. Pastalkova, C. Curto, and V. Itskov, "Clique topology reveals intrinsic geometric structure in neural correlations," arXiv: 1502.06172, 1 12 (2015).

- ³²A. Babuchev and Y. Dabaghian, "Persistent memories in transient networks," (2016), arXiv: 1602.00681v1.
- ³³P. Bendich, J. S. Marron, E. Miller, A. Pieloch, and S. Skwerer, "Persistent homology analysis of brain artery trees," (2014), arXiv: 1411.6652v1.
- ³⁴E. 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.
- ³⁵O. Sporns, "Graph-theoretical analysis of brain networks," in *Brain Mapping: An Encyclopedic Reference*, Vol. 1, edited by A. W. Toga (Academic Press: Elsevier, Cambridge, Massachusetts, 2015) pp. 629 – 633.
- ³⁶S. E. Petersen and O. Sporns, "Brain networks and cognitive architectures," Neuron **88**, 207–219 (2015).
- ³⁷D. 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).
- ³⁸A. S. Waugh, L. Pei, J. H. Fowler, P. J. Mucha, and M. A. Porter, "Party polarization in congress: A network science approach," (2009), arXiv: 0907.3509.
- ³⁹J. F. Donges, Y. Zou, N. Marwan, and J. Kurths, "The backbone of the climate network," Europhysics Letters (EPL) 87, 48007 (2009).
- ⁴⁰D. 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).
- ⁴¹S. 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).
- ⁴²D. Zhou, W. K. Thompson, and G. Siegle, "MATLAB toolbox for functional connectivity," NeuroImage 47, 1590 – 1607 (2009).
- ⁴³M. Á. Serrano, M. Boguná, and A. Vespignani, "Extracting the multiscale backbone of complex weighted networks," Proceedings of the national academy of sciences **106**, 6483– 6488 (2009).
- ⁴⁴A. F. Alexander-Bloch, N. Gogtay, D. Meunier, R. Birn, and L. Clasen, "Disrupted modularity and local connectivity of brain functional networks in childhood-onset schizophrenia," Frontiers in Systems Neuroscience 4, 1 – 16 (2010).

- ⁴⁵F. D. V. Fallani, J. Richiardi, M. Chavez, and S. Archard, "Graph analysis of functional brain networks: Practical issues in translational neuroscience," Philosophical Transactions of the Royal Society B **369** (2014).
- ⁴⁶C. Freitas, E. Macau, and A. Pikovsky, "Partial synchronization phenomena in networks of identical oscillators with non-linear coupling," Chaos 24, 024402 (2014).
- ⁴⁷G. Tirabassi, R. Sevilla-Escoboza, J. M. Buldú, and C. Masoller, "Inferring the connectivity of coupled oscillators from time-series statistical similarity analysis," Scientific Reports 5, 1 – 14 (2014).
- ⁴⁸X. Sun, M. Small, Y. Zhao, and X. Xue, "Characterizing system dynamics with a weighted and directed network constructed from time series data," Chaos: An Interdisciplinary Journal of Nonlinear Science **24**, 024402 (2014).
- ⁴⁹C. Kosniowski, A First Course in Algebraic Topology (Cambridge University Press, Cambridge, London, New York, New Rochelle, Melbourne, Sydney, 1980).
- ⁵⁰J. R. Munkres, *Topology* (Pearson Prentice Hall, New Jersey, 2000).
- ⁵¹H. 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).
- ⁵²G. Petri, M. Scolamiero, I. Donato, and F. Vaccarino, "Topological strata of weighted complex networks," PLoS ONE 8, e66505 (2013).
- ⁵³P. Bubenik, "Statistical topological data analysis using persistence landscapes," Journal of Machine Learning Research 16, 77 – 102 (2015).
- ⁵⁴P. Bubenik and P. Dłotko, "A persistence landscapes toolbox for topological statistics," arXiv: 1501.00179v3, 1 – 24 (2015).
- ⁵⁵V. Kovacev-Nikolic, P. Bubenik, D. Nikolic, and G. Heo, "Using persistent homology and dynamical distances to analyze protein binding," Statistical Applications in Genetics and Molecular Biology, 1 – 27 (2016), in press, arXiv: 1412.1394v2.
- ⁵⁶A. Tausz, M. Vejdemo-Johansson, and H. Adams, "JAVAPLEX: A research software package for persistent (co)homology," (2011).
- ⁵⁷H. Adams and A. Tausz, "JAVAPLEX tutorial," (2015), pDF version available online: http://javaplex.googlecode.com/svn/trunk/reports/javaplex_tutorial/ javaplex_tutorial.pdf.

- ⁵⁸J. Wildmann, "Bron-kerbosch maximal clique finding algorithm," (2011), code available at: http://www.mathworks.co.uk/matlabcentral/fileexchange/ 30413-bron-kerbosch-maximal-clique-finding-algorithm.
- ⁵⁹Y. Kuramoto, *Chemical Oscillations and Waves and Turbulence* (Springer, Berlin, 1984).
- ⁶⁰S. H. Strogatz, "From kuramoto to crawford: Exploring the onset of synchronisation in populations of coupled oscillators," Physica D 143, 1 20 (2000).
- ⁶¹A. Arenas, A. Díaz-Guilera, J. Kurths, Y. Moreno, and C. Zhou, "Synchronization in complex networks," Physics Reports 469, 93 – 153 (2008).
- ⁶²F. Rodrigues, T. K. D. Peron, P. Ji, and J. Kurths, "The kuramoto model in complex networks," Physics Reports 610, 1 – 98 (2016).
- ⁶³S. Gupta, A. Campa, and S. Ruffo, "Kuramoto model of synchronization: Equilibrium and nonequilibrium aspects," Journal of Statistical Mechanics Theory and Experiment **2014**, R08001 (2014).
- ⁶⁴J. L. P. Velazquez, "Brain research: A perspective from the coupled oscillators field," NeuroQuantology 4, 55 – 165 (2006).
- ⁶⁵M. Breakspear, S. Heitmann, and A. Daffertshofer, "Generative models of cortical oscillations: Neurobiological implications of the kuramoto model," Frontiers in Human Neuroscience 4, 190 (2010).
- ⁶⁶P. Ashwin, S. Coombes, and R. Nicks, "Mathematical framework for oscillatory network dynamics in neuroscience," The Journal of Mathematical Neuroscience **6**, 1 92 (2016).
- ⁶⁷M. A. Porter and J. P. Gleeson, "Dynamical systems on networks: A tutorial," Frontiers in Applied Dynamical Systems: Reviews and Tutorials and in press (arXiv:1403.7663) (2016).
- ⁶⁸A. Arenas, A. Díaz-Guilera, and C. Pérez-Vicente, "Synchronization reveals topological scales in complex networks," Physical Review Letters **96**, 114102 (2006).
- ⁶⁹In 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 science⁷⁸.
- ⁷⁰D. Prichard and J. Theiler, "Generating surrogate data for time series with several simultaneously measured variables," Physical Review Letters **73**, 951 – 954 (1994).
- ⁷¹D. S. Bassett, N. F. Wymbs, M. A. Porter, P. J. Mucha, and S. T. Grafton, "Cross-linked structure of network evolution," Chaos: An Interdisciplinary Journal of Nonlinear Science

24, 013112 (2014).

- ⁷²B. Stolz, Computational Topology in Neuroscience, Master's thesis, University of Oxford, http://people.maths.ox.ac.uk/porterm/research/Dissertation-stolz2014-Corr. pdf (2014).
- ⁷³A. Sizemore, C. Giusti, and D. Bassett, "Classification of weighted networks through mesoscale homological features," (2015), arXiv: 1512.06457v2.
- ⁷⁴D. 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," Proceedings of the National Academy of Sciences **108**, 7641 – 7646 (2011).
- ⁷⁵D. 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 Computational Biology 9, e1003171 (2013).
- ⁷⁶D. S. Bassett, M. Yang, N. F. Wymbs, and S. T. Grafton, "Learning-induced autonomy of sensorimotor systems," Nature Neuroscience 18, 744 – 751 (2015).
- ⁷⁷H. 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," (2015), arXiv: 1507.06217.
- ⁷⁸M. A. Porter, J.-P. Onnela, and P. J. Mucha, "Communities in networks," Notices of the American Mathematical Society 56, 1082 – 1097, 1164 – 1166 (2009).
- ⁷⁹F. H. Croom, Basic Concepts of Algebraic Topology (Springer, New York, Heidelberg, Berlin, 1978).

Appendix A: Topological background and definitions

We give a brief introduction to the mathematical concepts behind persistent homology. (We adapt and summarize the discussion from B. Stolz's masters thesis⁷².)

1. Simplicial complexes

One can represent the underlying structures of topological spaces 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 to use a small number of complicated pieces or to use a large number of simple pieces. From a computational point of view, the latter is preferable⁵.

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 A.1 (affine combination and affine hull). Let $\mathcal{U} = \{u_0, u_1, \ldots, u_k\}$ 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 $i \in \{0, \ldots, k\}$, 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 A.2 (affinely independent). Let $\mathcal{U} = \{u_0, u_1, \ldots, u_k\}$ be points in \mathbb{R}^d . The k+1 points in \mathcal{U} are said to be affinely independent if the vectors $\{u_i - u_0 : 0 \leq i \leq k\}$ 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 A.3 (convex combination and convex hull). An affine combination $x = \sum_{i=0}^{k} \lambda_i u_i$ is a convex combination if $\lambda_i \ge 0$ for all $i \in \{0, \ldots, k\}$. The set of all convex combinations of the points in \mathcal{U} is called the convex hull of \mathcal{U} .

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

Example A.2. In Figure 8, 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 the tetrahedron is a 3-simplex.



FIG. 8. Examples of (left) a 0-simplex, (second) a 1-simplex, (third) a 2-simplex, and a (right)
3-simplex. The examples and figure are adapted from⁵.

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

Definition A.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 .

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 appendix, we can ask whether it is only possible to build shapes using 2-simplices (i.e., triangles) or whether we 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*. In order to construct a simplicial complex one needs to follow a set of minimal rules:

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

i. if $\sigma \in \Sigma$ and $\tau \leq \sigma$, it follows that $\tau \in \Sigma$;



FIG. 9. Panels (a), (b), and (c) gives examples of simplicial complexes. The collection of simplices in panel (d) is not a simplicial complex. We use colors to indicate 2-simplices.

ii. if $\sigma, \tilde{\sigma} \in \Sigma$, it follows that the intersection of both simplices is either the empty set or a face of both.

In Figure 9, we show examples of simplicial complexes. 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 again 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.

2. Homology and Betti numbers

Homology is a formal way of quantitatively detecting holes in various dimensions to give insight into the connectivity of a topological space. These holes are quantified by classifying the space, e.g. loops, surrounding them. For example, one can distinguish a 2-sphere from a torus by capturing the fact it is possible to contract any 1-dimensional loop on the sphere to a point, whereas there are two distinct loops on the torus surface that cannot be continuously deformed into each other. These loops also cannot be contracted to a point, because they surround different holes.

Although homology is not the only and most detailed formalism that can be used for the distinction between two shapes, thus far it has the fastest algorithms⁵. *Homology groups*, which are topological invariants of a space, and *Betti numbers* (which are derived from them) play a key role in this endeavour. Homology groups detect holes in a topological space, whereas Betti numbers give a way to count the number of holes or distinct loops in that space. We start constructing the homology groups by looking at formal sums of simplices.

Definition A.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

$$c = \sum_{i \in I} a_i \sigma_i \tag{A1}$$

is mathematically a formal sum 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 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 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 $C_p = C_p(\Sigma)$ 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 component-wise manner:

$$c + c' = \sum_{i \in I} (a_i + b_i)\sigma_i.$$
(A2)

It then follows that *p*-chains form an Abelian group. Observe when working with coefficients from $\mathbb{Z}/2\mathbb{Z}$ that the sum of two *p*-chains results in the sum of all *p*-simplices in which the two original *p*-chains differ. The *p*-simplices that the two *p*-chains have in common is present in the sum twice, and it therefore vanishes 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 A.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:

$$\partial_p \sigma = \sum_{j=0}^p [u_0, \dots, \hat{u}_j, \dots, u_p], \qquad (A3)$$

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

We can naturally extend this definition to *p*-chains 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:

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

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^{5,79} that the following theorem holds for boundary homomorphisms in a chain complex:

Theorem A.1. Let $d \in C_{p+1}$. It follows that

$$\partial_p \partial_{p+1} d = 0. \tag{A4}$$

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 boundary homomorphisms and their property stated in Theorem A.1, now form the main ingredients in constructing the homology group of a simplicial complex.

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

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

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

Denote the set of *p*-boundaries as \mathcal{B}_p , and notice that $(\mathcal{B}_p, +)$ is a subgroup of $(\mathcal{C}_p, +)$. Using Theorem A.1, we can now relate the subgroups to each other: From Theorem A.1, it follows that $\partial_p(\operatorname{Im} \partial_{p+1}) = 0$, so $\mathcal{B}_p \subseteq \mathcal{Z}_p$. One can then show that \mathcal{B}_p is indeed a subgroup of \mathcal{Z}_p . Note that 1-dimensional loops behave differently from other 1-edges. The latter are mapped to their end nodes by ∂_1 , but every node in a 1-loop occurs as the boundary of two edges and thus sums to 0 over $\mathbb{Z}/2\mathbb{Z}$. We have now come very close to our goal of being able to count holes and loops. We have thus far identified the subgroup to which such loops will belong, but this subgroup still also contains the boundaries of higher-dimensional chains. To distinguish the two subgroups, we need to define the *p*th homology group of a simplex.

Definition A.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,

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

Two *p*-cycles in the *p*th homology group are regarded 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*-cycle in the homology group. Counting the number of classes in \mathcal{H}_p thus gives an estimate of the number of *p*-dimensional loops of a simplicial complex. However, loops that surround the same hole are counted separately. A solution to this problem is to count the minimal number of elements that are needed to generate the group. This leads to the definition of *p*th Betti number.

Definition A.12 (*pth Betti number*). The *pth Betti number* β_p of a simplicial complex is

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

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, respectively, the number of connected components, the number of 1-dimensional loops, and the number of 2-dimensional holes in a simplicial complex.

3. Filtrations

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

Definition A.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 A.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:

$$\emptyset = \Sigma_0 \subseteq \Sigma_1 \subseteq \Sigma_2 \subseteq \dots \subseteq \Sigma_k = \Sigma.$$
(A5)

Observe that one can define natural inclusion maps $i_j : \Sigma_j \hookrightarrow \Sigma_{j+1}$ along the filtration.

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

Definition A.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 inclusion map $i_{m-1} : \Sigma_{m-1} \hookrightarrow \Sigma_m$.

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 inclusion map $i_{n-1}: \Sigma_{n-1} \hookrightarrow \Sigma_n$.

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

$$p_h = n_h - m_h \, .$$