Topological Segmentation of Time-Varying Functional Connectivity Highlights the Role of Preferred Cortical Circuits

Jacob Billings*
Mathematics and Complex Systems Research Area, ISI Foundation, Turin, Italy

Manish Saggar
Department of Psychiatry & Behavioral Sciences,
Stanford University School of Medicine, Stanford, California, USA

Shella Keilholz
Coulter Department of Biomedical Engineering, Emory University and Georgia Institute of Technology, Atlanta, Georgia, USA

Giovanni Petri†
Mathematics and Complex Systems Research Area,
ISI Foundation, Turin, Italy
ISI Global Science Foundation, New York, USA
(Dated: September 6, 2020)

Functional connectivity (FC) and its time-varying analogue (TVFC) leverage brain imaging data to interpret brain function as patterns of coordinating activity among brain regions. While many questions remain regarding the organizing principles through which brain function emerges from multi-regional interactions, advances in the mathematics of Topological Data Analysis (TDA) may provide new insights into the brain's spontaneous self-organization. One tool from TDA, "persistent homology", observes the occurrence and the persistence of n-dimensional holes presented in the metric space over a dataset. The occurrence of n-dimensional holes within the TVFC point cloud may denote conserved and preferred routes of information flow among brain regions. In the present study, we compare the use of persistence homology versus more traditional TVFC metrics at the task of segmenting brain states that differ across a common time-series of experimental conditions. We find that the structures identified by persistence homology more accurately segment the stimuli, more accurately segment volunteer performance during experimentally defined tasks, and generalize better across volunteers. Finally, we present empirical and theoretical observations that interpret brain function as a topological space defined by cyclic and interlinked motifs among distributed brain regions, especially, the attention networks.

Keywords: functional connectivity, time-varying functional connectivity, topological data analysis

I. Introduction

One of the perennial questions in neuroscience concerns how signals in the brain self-assemble to generate time-varying experiences. One foundation from which to address this question asserts that brain function emerges from neuronal communication within the context of multiscale neuronal networks. Thus, to better understand the configuration of brain networks over time is to better understand how the brain functions.

The field of Time-Varying Functional Connectivity (TVFC) offers an empirical approach to understand the brain's functional dynamics across time [27, 29]. TVFC research is rooted in the theoretical foundations of Functional Connectivity (FC), and is its time-varying analogue. FC measures the long-term correlation among the spontaneous neuroimaging signals from pairs of brain regions. The method delimits the same brain networks as encountered in more traditional task-based Blood Oxygen Level-Dependent (BOLD) functional Magnetic Resonance Imaging (fMRI) methods, wherein the mean within-condition BOLD signal is differenced across conditions [7]. These results from fMRI were replicated and extended using a wide variety of other data sources, such as, EEG [31, 33, 43], MEG [16, 37], PET [65], and calcium imaging [15]. Additionally, their stability has been measured against a variety of manipulations, such as, lesion studies [36], drug interventions [42], and across behaviors [28, 64]. Similarly, a host of alternative quantitative methods broaden the set of perspectives from which to view intrinsic patterns formed by the brain's spontaneous activity [3, 22, 38, 56].

Conventional approaches to calculating FC and TVFC focus on studying the overlap between properties of the network of brain regions. The most conventional overlap property is across the normalized set of edge distances between nodal brain regions. Other properties include the mean BOLD activity at each node. And also the mean edge weight centered at each node, called the "weighted-degree" or "strength."

Conventional FC and TVFC methods tend to blur pat-

* jacob.billings@isi.it
† giovanni.petri@isi.it

terns among the properties over brain regions as they acquire a summation over overlap distances. A different perspective that has emerged over the last few years focuses on the occurrence of conserved topological features [4, 48]. These features provide for the quantification of neural data as ensembles of n-dimensional structures [10, 12, 55]. To assess the relative utility of the topological model, we investigate the relative sensitivity of TDA features vs network features in the task of segmenting experimentally defined brain states.

The topological approach to data analysis sees data as an approximate sampling of an underlying space describing the system under study. These spaces have shapes that may be quantified and compared across datasets. Indeed, TDA has been described as "exactly that branch of mathematics which deals with qualitative geometric information" [8]. Like graph theory, TDA facilitates the investigation of a network's 1-dimensional skeleton. Additionally, the TDA approach has the capability to capture patterns of coordination beyond pairs of nodes and local structures. TDA accomplishes this by promoting elementary objects from pairwise interactions to their group interactions: Beyond collecting the edges between nodes into a network, TDA focuses on simplices—groups of nodes of any size—collected together into "simplicial complexes." For example, simplicial complexes between triplets of nodes may be connected to form a 3-dimensional "clique," or, equivalently, the 2-dimensional surface that is the convex hull established between the node triplet. Complexes of surfaces may connect such that they are the boundary of a 4-dimensional "homology," or, more commonly, a "hole" or "void." These and other n-dimensional complexes constitute genuine mesoscopic features of datasets that are invisible to standard graph-theoretic metrics.

Given these foundations, the present study seeks to identify if a topological interpretation of brain dynamics improves the capacity to distinguish states of the brain; relative to more traditional network approaches. And, if so, what novel interpretations of brain function may we draw?

II. Methods

The study's methods comprise 4 operations:

1. data acquisition and preprocessing,
2. time-varying functional connectivity (TVFC),
3. topological data analysis (TDA), and
4. embedded visualization and inspection.

A graphical depiction of the study methods may be found in figure 1.

A. Data acquisition and preprocessing

To discern the relative capacities of a range of graph-graph distance metrics to disambiguate the brain-states induced by stimuli, for the present study, we adopted a dataset acquired during the presentation of multiple experimentally defined tasks. The present study benefited from scans acquired continuously over relatively long time spans as the process of filtration requires complete overlap between the signal and the filtration kernel to avoid affects at the undefined edges of the time series.

The data acquired by Gonzalez-Castillo et al. (2015) meet these criteria [24]. These data have been publicized as an open-access dataset through the XNAT neuroimaging database (https://central.xnat.org; Project ID: FC-StateClassif). Here, we briefly summarize the dataset as follows: 18 volunteers were scanned continuously over 25.5 minutes (7 Tesla, 32-element coil, gre-EPI, TR=1.5s, TE=25ms, 2mm isotropic). Preprocessing was performed to transform individual datasets into a common MNI space and to remove artifacts from slice timing, motion, linear trends, quadratic trends, white matter signals, and csf signals. Data were spatially smoothed to 4mm isotropic, and temporally band-pass filtered to between 0.009 Hz and 0.08 Hz. For complete preprocessing details, please refer to [52]. In addition to the aforementioned steps, voxelwise data were spatially aggregated onto an atlas of 333 brain regions [25]. Up to 5 brain regions contained no information from some volunteers, and were excluded from all datasets for the remainder of the analysis. Thus, the finest granularity of study results are over 333-5=328 brain regions. During the scan, volunteers interacted with 3 block-design tasks and one rest stimulus. Each task was presented twice. Each task presentation lasted 3 min, and was proceeded by a 12s instruction block. The tasks included 1) 'video': watching videos of a fish tank while responding to a visual target, 2) 'math': computing algebra problems, and 3) 'memory': a 2-back memory task with abstract shapes. A 'rest' stimulus was also included, and entailed the presentation of a fixation cross for 3 minutes. Stimuli were randomly ordered in a fixed sequence for all volunteers. For each task block, performance metrics were collected, including the percentage of correct responses.

B. Time-varying connectivity

Considering that individual frequency bands develop significantly different FC parcellations [5] and support different connectivity hubs [59], and that neuroelectric activity is intrinsically rate coded, the delayed and (hemodynamic response function) band-pass filtered version of neuroelectric activity that is the BOLD signal is likely to retain some rate-coded information. Given these observations, the present study recasts the BOLD signal from each brain parcel in terms of time-frequency spectrograms generated through the use of the Continuous

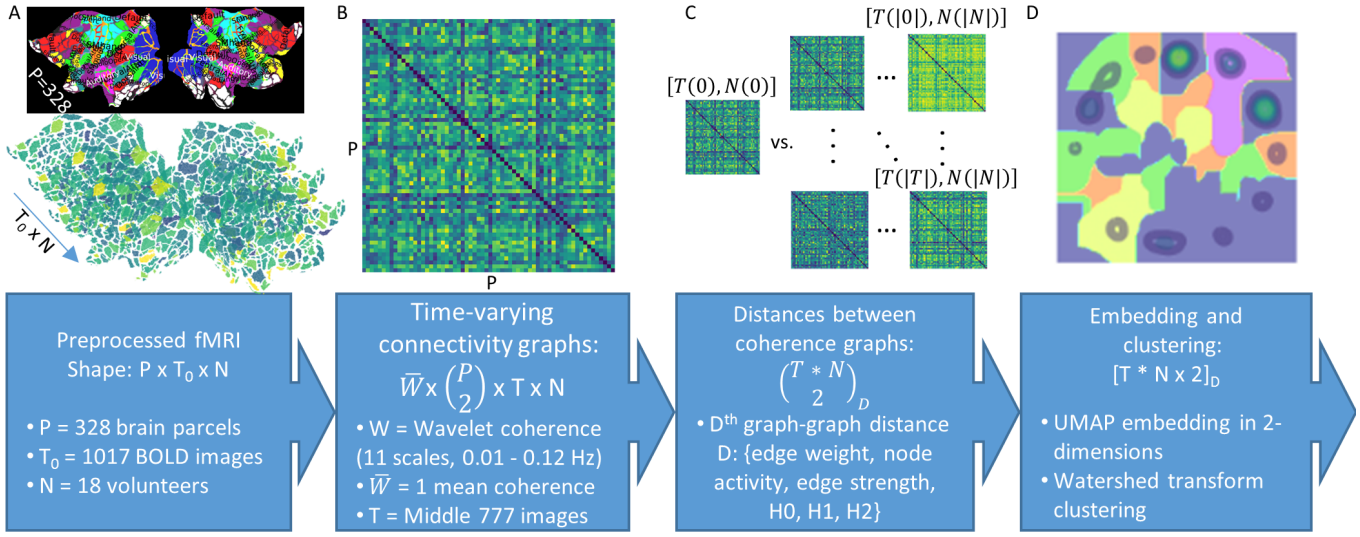

FIG. 1: **Analysis pipeline.** We present the analysis pipeline as a flow diagram in four steps. First, the pipeline accepts preprocessed and spatially segmented BOLD fMRI data as inputs. Then, for each scan, we compute time-varying functional connectivity (TVFC) matrices as the weighted mean of the wavelet coherence between all brain regions, and at each moment in time. Because the wavelet kernel operates over a portion of the time-frequency domain, we remove the outside temporal and spectral edges of the coherence matrix where data padding is required. Next, we compare TVFC matrices using one of six graph-graph matrics. Three metrics compare the overlap among simplex weights, while the other three compare the lifetimes of persistent homologies in dimensions 0, 1, and 2. Finally, we embed graph-graph metric spaces onto two dimensions where we may visualize and cluster point distributions.

Wavelet Transform (CWT)

$$W_t(s) = \sum_{t=0}^{T-1} f(t) \frac{1}{s} \psi^* \left(\frac{t-u}{s} \right) \delta t,$$

where $\cdot^*$ indicates the complex conjugate. By adjusting the time localization parameter u and the scale parameter s for the wavelet kernel ψ, the CWT affects a multiscale decomposition of input signal $f(t)$ for all times T. For the present study, the filterbank comprised 15 scales log-distributed between 0.007 and 0.15 Hz.

We adopt the complex Morlet wavelet as the CWT kernel, $\psi = e^{i\omega_0 t/s} e^{-t^2/(2s^2)}$, with base frequency $\omega_0 = 6$. The availability of instantaneous amplitude and phase information provides a basis from which to compute the coherence between signals. For a pair of spectrograms, W^X and W^Y, the quantity $W_t^{XY}(s) = W_t^X(s) W_t^{Y*}(s)$ is the cross-wavelet spectrum. Its absolute value, $|W_t^{XY}(s)|$, is the cross wavelet power which represents the shared power between signals at scale s and time t. Coordinated changes in amplitude may be computed in terms of the wavelet squared coherence,

$$R_t^2(s) = \frac{|\langle s^{-1} W_t^{XY}(s) \rangle|^2}{\langle s^{-1} |W_t^X(s)|^2 \rangle \langle s^{-1} |W_t^Y(s)|^2 \rangle}.$$

The functional $\langle \cdot \rangle$ indicates smoothing in both time and scale. The factor s^{-1} is used to convert to scale-dependent energy densities. The wavelet squared coherence is an instantaneous and multispectral analogue of the Pearson correlation [39, 60, 61]. Its values range between 0 (completely incoherent) and 1 (completely coherent). While it is theoretical possible to treat TVFC as a multilayer graph having as many layers as spectral scales, practical computational concerns prompt us to concatenate multispectral coherence into a single broadband average by taking the weighted mean of the wavelet squared coherence with respect to the normalized cross wavelet power:

$$TVFC_t^{XY} = 1 - \sum^{s} \frac{|W_t^{XY}(s)|}{\max_s |W_t^{XY}(s)|} R_t^2(s). \quad (1)$$

Normalizing the cross wavelet power ensures that the mean coherence remains bounded between 0 and 1. The peak of the mean cross wavelet power occurs in the frequency range between 0.02 Hz and 0.01 Hz (see part A of supplemental figure A.1). TVFC graph edges are 1 minus the power-weighted coherence to represent coherence distances between brain regions.

To account for the cone of influence at the temporal edges of the wavelet filtration, as well as the loss of precision at the temporal and spectral edges of the smoothed coherence data, the outside 120 time points and the outside 2 scales are dropped before taking the summation in equation 1. The removed time points include one whole "rest" block, and one whole "video" block. Coherence graphs thus represent the middle 777 images of the scan, and the 11 spectral scales between 0.01 and 0.095 Hz.

C. Simplicial and homological graph-graph distances

Having constructed TVFC graphs for all included time points and for all volunteers, we pursue two complementary classes of metrics to describe and compare graphs. The first class of metrics are the traditional network connectivity-based metrics associated with the weights of graph nodes and/or edges. As we will see, these metrics refer directly to 0-dimensional and/or 1-dimensional *simplices*. The second class of metrics takes an explicitly topological view of the data and describes *homologies* embedded within the graph. To describe in detail these two views, it is useful to address some definitions from topology.

A graph $G = (V, E)$ represents a set of V nodes interconnected by E edges. Nodes and edges may be decorated with properties such as weight, directionality, sign, and layer. A collection of k completely interconnected nodes forms a clique, C. In the following, we identify cliques with simplices in standard fashion [45, 47]; that is, to a clique of $k+1$ nodes we associate the corresponding k-simplex, σ_k. For instance, 2 connected nodes forms a 2-clique, a 1-simplex, or an 'edge'. A 2-simplex formed by a clique of 3 connected nodes is a 'filled triangle', and so forth for higher-order simplices. Thus the simplicial approach directly compares the decorations applied to sets of simplices.

Traditionally, metrics measuring the difference between datasets compute a normalized summation of the total difference between edge weights. By contrast, the topological approach continues from the identification of discrete simplexes to interpreting complexes of simplices. Formally, a simplicial complex is a topological space, $\mathcal{K}$, composed of all σ_k and their subfaces. Along the same lines, a clique complex, $Cl(G)$, is a simplicial complex formed from an unweighted graph G by promoting every k-clique into a $(k-1)$-simplex.

Holes in dimension k may develop within the boundaries established by closed chains of $(k-1)$-simplices. Such holes are called 'homologies.' The boundary of a homology is termed a, 'homological cycle' or 'generator.' To illustrate the concept, consider the case of four nodes connected in a cycle such that each node has exactly two edges. The nodes form neither a 4-clique nor a 3-dimensional simplex because there are two missing edges. Rather, these nodes form a connected cycle that is the boundary of a 2-dimensional hole. This void space is also called a *homology* in dimension 1 (i.e., formed by a set of 1-d edges). The kth homology group, $H_k(\mathcal{K})$, describes the $(k+1)$-dimensional holes bounded by chains of k-simplices. For example, the H_1 homology group are the holes bounded by edges in $\mathcal{K}$; H_2 are the voids bounded by filled triangles; etc.

The term 'homology' follows from the Greek 'homo,' the same, and 'logos,' relation, to indicate that the hole belongs to an equivalence class that is categorically the same across continuous deformations that neither break the boundary nor create new simplices spanning the boundary: e.g., inflation, compression, rotation, and translation. Different representative cycles may therefore exist that describe the same homological cycle. For instance, a very elastic coffee cup could be continuously contracted into the shape of a donut, as they share the same toroidal topology. For the sake of convenience, a homological cycle is often represented as the minimal representative cycle [26, 45].

Simplicial graph-graph distances

We first compare distances between graphs using a measure based on the overlap between network properties. This approach, which we will denote as 'simplicial,' quantifies the total difference between the values associated with each k-simplex in the complex. For example, the graph-graph distance between 1-dimensional simplices compares the values of each graph's *edges*. In the present study, we utilize the weighted Jaccard overlap distance to characterize simplicial overlap distances in dimension 1 as a single number:

$$D_J(G1, G2) = 1 - \frac{\sum_{e \in E} \min(e^w_{G1}, e^w_{G2})}{\sum_{e \in E} \max(e^w_{G1}, e^w_{G2})},$$

where e^w_G is the weight of the eth edge in graph G.

Further we compute graph-graph distances between the explicit 0-dimensional information of each time point; e.g. with respect to the signal activity of each *node*. Specifically, for each point in time, we treat the absolute values of multispectral wavelet coefficients from all brain regions as an ordered vector. We then compute the Euclidean distance between vectors from different points in time.

The third distance is inspired by previous work on relations between node connectivity and homological cycles. They demonstrate that the nodes' weighted degree (also called *strength*) is significantly correlated with the frequency and the intensity with which nodes participate in the shortest representatives of homological cycles [34]. The third distance is thus the weighted Jaccard distance, across time, of node-wise weighted degree, or *strength*.

Homological graph-graph distance

While many TVFC studies regard only the graph's connectivity as the feature of primary import, TDA provides a suite of tools to further develop network properties into conserved higher-order structures in point-cloud data [8, 17, 44] and in weighted networks [11, 47, 53].

Homology is defined on simplicial complexes. In the case of persistent homology of weighted graphs, simplices are added to the complex incrementally, and appear at and beyond some threshold. Varying this threshold allows to track how homological features appear and persist across thresholds [47]. In particular, by observing

topological features over multiple thresholds, "persistent homology" allows to take a multiscale view of the data which accounts for both the explicit connectivity structure of the system, as well as the relative importance of ensembles of connections that emerge over some range of scales. Formally, we define the Vietoris-Rips simplicial complex $\mathcal{K}_r = Rips(G(E < r))$ as the clique-complex of the weighted graph G composed after removing all edges longer than r. From this, we may recover the complex's k-dimensional homology group, $H_k(\mathcal{K}_r)$. Let $r_0, \ldots, r_k$ be the longest series wherein any $Rips(G(E < r))$ and $Rips(G(E < r - \epsilon))$ are not identical. The ordered sequence of this series defines a "Rips-filtration" over G. Letting $\mathcal{K}_{r=0} = \emptyset$, and, letting $\mathcal{K}_{r=\inf}$ be the maximal simplicial complex defined over G, a filtration over G defines a sequence of homology groups:

$$H_k(|\mathcal{K}_0|) \rightarrow H_k(|\mathcal{K}_1|) \rightarrow \ldots \rightarrow H_k(|\mathcal{K}_n|).$$

A homology class $[\alpha]$ is said to be *born* at radius s if the homology is not supported in $\mathcal{K}_r$ for any $r < s$. The homology class $[\alpha]$ is said to *die* going into $\mathcal{K}_t$ if t is the lowest index wherein $[\alpha]$ is no longer an homological cycle of the simplicial complex.

The persistent homology of a filtration over G is summarized by collecting the birth/death pairs of k-dimensional homology classes as points (s, t) in a "persistence diagram". It is naturally possible to compute a persistence diagram for each simplicial dimension up to their maximum dimension in the simplicial complex. The case of 0-dimensional persistence diagrams—corresponding to 0-dimensional holes, that is, disjoint sets of connected nodes—is particularly interesting as the homological classes are slices through an agglomerative clustering among nodes when using the "simple" linkage distance.

Persistent homology was computed using version 0.4.1 of the Ripser package as bundled with the Scikit-TDA toolbox for python [62]. Because the computational load to calculate persistence homology increases exponentially with the homology dimension, we limit the present study to the investigation of persistence homology in dimensions 0, 1, and 2.

Persistence diagrams can, themselves, be endowed with a metric structure. This means that it is possible to measure distances between persistence diagrams. Such distances encode how different the homological structures of two TVFC graphs are. One such distance is a multi-dimensional analogue of the earth-mover distance, known as the sliced-Wasserstein distance [9]. The sliced-Wasserstein distance between persistence diagrams is bounded from above by the total distance between the associated topological spaces [41]. In the present study, for each pair of persistence diagrams of a given dimension, we calculate homologically-informed TVFC graph-graph distances as the average Wasserstein distance over 20 slices (see [9] for details). In each respective dimension, we refer to these metrics as $H0$, $H1$, and $H2$.

D. Visualization/Output

Having developed TVFC metric spaces for simplicial and homological distances, we want to leverage them to investigate their relative capacities to represent statefulness. To this end, we embed each TVFC metric space onto a 2-dimensional manifold using the Uniform Manifold Approximation and Projection (UMAP) algorithm [40]. As illustrated in figure 2, the embedding process facilitates state-space visualization and segmentation. UMAP approximates a metric space's n-dimensional manifold in three steps. First, the algorithm calculates the k-nearest neighbors of each point. Second, each neighborhood is promoted to a local simplicial complex. Third, the algorithm searches for the n-dimensional distribution of points that best approximates the original simplicial complex. This search is conducted over successive iterations, with the initial position of low-dimensional points derived from a random distribution.

The UMAP algorithm possesses many features that surpass previous state of the art embedding algorithms, such as tSNE [63]. For example, whereas tSNE's low-dimensional representation becomes unstable in dimensions above 3, the simplicial approximation implemented by UMAP will, in theory, produce stable embeddings in any dimension. Additionally, UMAP harbors a suite of parameters that may be tuned to resolve different qualities of the input point cloud.

In the present study, we were concerned with resolving 2-dimensional embeddings that generalize across volunteers while simultaneously segmenting experimental stimuli. One challenge in the way of resolving this ideal embedding is that, through the lens of some TVFC metric spaces (see Results), brain states change slowly and smoothly through time. The resulting temporal autocorrelation makes the distance between adjacent time points much shorter than the distance between any two volunteers experiencing the same stimuli. This is especially the case for graph-graph metrics that compare weights of 1-d simplexes (see supplemental figure A.2, and Discussion). For dimensionality-reduction algorithms like UMAP and tSNE that leverage nearest-neighbor approximations, the attractive force between temporally-adjacent time points can force the embedding to over-emphasize information about the order of the scanning sessions when attempting to resolve population-wise brain states.

To disentangle graphs representing intrinsically similar brain states from those that are simply autocorrelated, we trained each embedding with a selected subsample of 13% of the data from each volunteer (100 time points from 777). The 13% subsampling reduced the dataset in two steps to ensure randomness and temporal separation. Step one randomly chooses 50% of points; and second step selects the final subsample based upon maximal temporal separation among the randomly halved dataset. After subsampling, we forced UMAP to observe similarities between points from different volunteers by setting

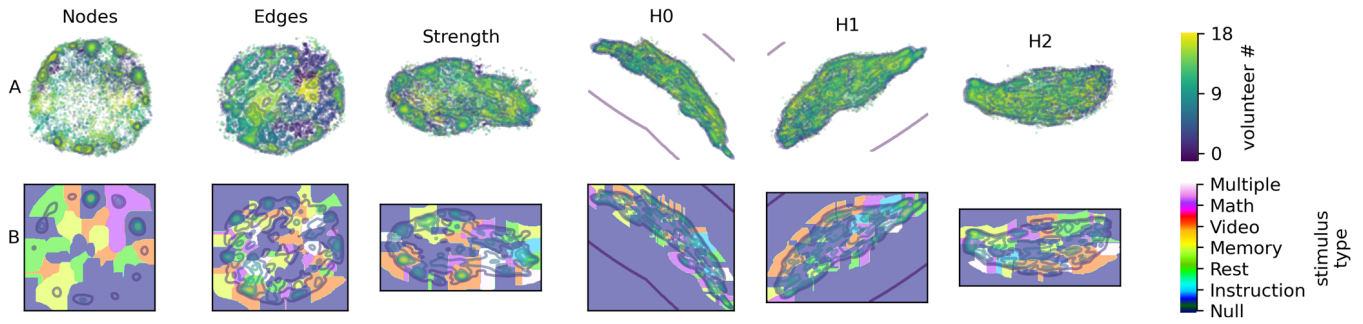

FIG. 2: **TVFC embeddings for different underlying metrics.** We display pointwise (part A) and cluster-wise (part B) time-varying functional connectivity embeddings using each of six graph-graph metrics. Part A of the figure color codes points with respect to volunteer labels. While Part B color codes regions significantly populated with points from one or more of the 5 experimental stimuli. The regions in part B are identified via a watershed transform of the underlying point density. Note: because the random initialization of embedded point locations affects the final distribution, the displayed embeddings should be viewed as one sample drawn from a central tendency.

the nearest neighborhood parameter to twice the number of points from each volunteer's subsample (k nearest neighbors = 200). In addition, we restricted the minimum distance between pairs of embedded points to encourage dispersion (min dist=0.3); and accelerated point dispersion by increasing the value of the repulsive force (repulsion=5). After training an embedding on a subsample, we fit the remaining out-of-sample data points to the trained embedding.

To better understand the distribution of points in the resulting embedding spaces, we transformed point clouds into a Gaussian distribution and estimated clusters via a watershed transform. An illustration of watershed clustering is found in part B of figure 2. The Gaussian grid size was initially set to 256x256. The number of grid points in the dimension having the smaller range was trimmed to maintain the aspect ratio of the embedding. The Gaussian kernel bandwidth factor was set to 0.08. The watershed transform marks the local densities as cluster centers, then grows clusters by adding adjacent pixels whose directed gradient is maximal in the direction of the cluster center.

E. Statistical Analysis

Watershed clusters provide a data-driven basis for hypothesis testing over the likelihood that certain metadata labels—that is, volunteer number, stimulus type, and performance—were more or less likely to be found in a given embedding region. For all statistical tests, we generated the null distribution mean and standard deviations by randomly permuting the relevant labels of cluster points 300 times. The significance threshold was always set to an $\alpha = 0.05$. Bonferroni correction was applied relative to the number of simultaneous tests performed. For reference, the total number of volunteers was 18. Tests related to volunteer co-localization calculated significant volunteer-wise under-representation in each cluster (left-tail test, Bonferroni correction equal to the number of volunteers times the number of clusters per embedding ($O100$)). Tests related to stimulus co-localization identified clusters that were more than likely to contain time periods during each stimulus condition (right-tail test, Bonferroni correction equal to the number of stimulus conditions (5) times the number of clusters in each embedding ($O100$)). Tests related to task performance were conducted for each task condition independently, and were confined only to the clusters that were significantly more likely to contain points from the task being tested (two-tailed test, with Bonferroni correction equal to the number of clusters with a significant per-stimulus population ($O10$)).

F. Secondary Statistics over Mean Graphs

It is possible to generate mean FC matrices from select time points of TVFC graphs. For instance, the mean TVFC graph over all time points reveals the average coherence between regions. Condition-dependent mean graphs such as that over all rest conditions may also be calculated. In the present study, we were particularly interested in mean graphs calculated across task performance.

Given the identification of clusters significantly associated with task performance, for each task, and for each cluster associated to the task, we tested whether the task-specific points within that cluster contained significantly more or fewer correct responses than the mean percentage of correct response for all of that task's time points (no Bonferroni correction). For each task, every time point from clusters having significantly more correct responses is stored into a task-specific list. The same process occurs for clusters showing fewer correct responses. The mean TVFC graph from each list constitutes a "mean performance graph."

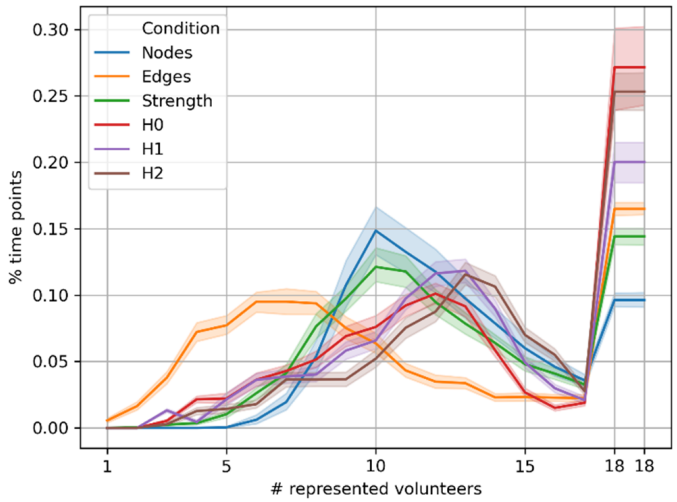

FIG. 3: **Volunteer specificity of watershed regions.** We plot the percentage of TVFC graphs points lying within each of 18 bins. Each bin considers the points within clusters wherein between 1 and 18 volunteers had not an insignificant number of points. Data are presented as a central tendency and 95% confidence interval over 128 independent samples, each from a randomly initialized embedding.

G. Spatial Coarsening

To facilitate visualization, it is worthwhile to reduce the dimensionality of the dataset rather than attempt to threshold the results. Thus we spatially coarsen the number of nodes by taking the Euclidean distance of a 2-dimensional UMAP projection of the mean graph over all time points. The UMAP embedding was initializing with the spatial location of each brain region on a flat brain map. Parcel locations in 2-dimensions were encouraged to remain adjacent to their neighbors by specifying the minimum point-to-point distance in UMAP as min_dist=0. The default repulsion_strength=1 was multiplied by a factor of 10 to compensate for arbitrary point overlap.

Given, thus, a Euclidean distance over spatially aggregated parcels, the similarity space among parcels is modeled as an agglomerative clustering tree using the average linkage distance. Coarsened parcellations are then generated across a range of resolutions by thresholding the linkage distance.

III. Results

A. Volunteer-wise representation

As an initial test of the quality of each embedding space, we ascertain how well the clusters in each embedding generalize across volunteers. In part A of figure 2, we report an instance of the embedding algorithm output from each of six different underlying metrics: simplicial overlap between nodes, edges, and edge strength, as well as homological distances for $H0$, $H1$ and $H2$. In this visualization, groups of dots of the same color imply that the embedding separates volunteers. This is especially observed in the simplicial overlap embeddings. Since the precise structure of each instance of an embedding is influenced by the random points sampled for its initialization, we quantify volunteer-wise generalizability over 128 embeddings independently reinitialized. In each embedding instance, we determined the number of volunteers having at least an average number of points in each watershed region.

Figure 3 displays the bootstrapped distribution of the number of points falling into each region having 1 through 18 well-represented volunteers. Bold lines display mean tendencies, while shaded regions show the 95% confidence interval. A right-skewed distribution indicates increased generalizability, because it means that more watershed regions are significantly populated with more volunteers. A left-skewed distribution indicates that most watershed regions are specific to one or few volunteer, making the embedding apt at separating volunteers, but less able to generalize states (Figure 3).

Overall, the homological metric spaces offer embeddings that generalize better across the volunteers than the simplicial metric spaces. Not only are the homological distributions skewed more to the right, they also aggregate significantly more points in clusters that are general for all 18 volunteers.

B. Stimulus segmentation

A central indicator of embedding quality is the degree to which time points co-localize when belonging to different stimulus conditions. Part B of figure 2 shows an example result of testing watershed clusters against the hypothesis that a significant number of within-cluster points corresponds to one or more of the 5 experimental conditions. For each stimulus type, figure 4 shows the percentage of points from that stimulus residing in clusters significantly associated with that stimulus (blue boxes). Here again, we report the result as a distribution after 128 independently reinitialized embeddings. Larger percentages of significantly co-localizing points indicate increased capacity to identify brain-states associated with experimental stimuli.

For comparison, we offer two null models computed from randomly permuted point labels. The first null distribution (yellow boxes) permutes point labels within the significant clusters defined previously. It reflects the expected number of points that would randomly collect into the predefined set of significant clusters. The inclusion of this null model is motivated by the fact that some embeddings clump more points than others into the same watershed region, and would thus hold a larger percent-

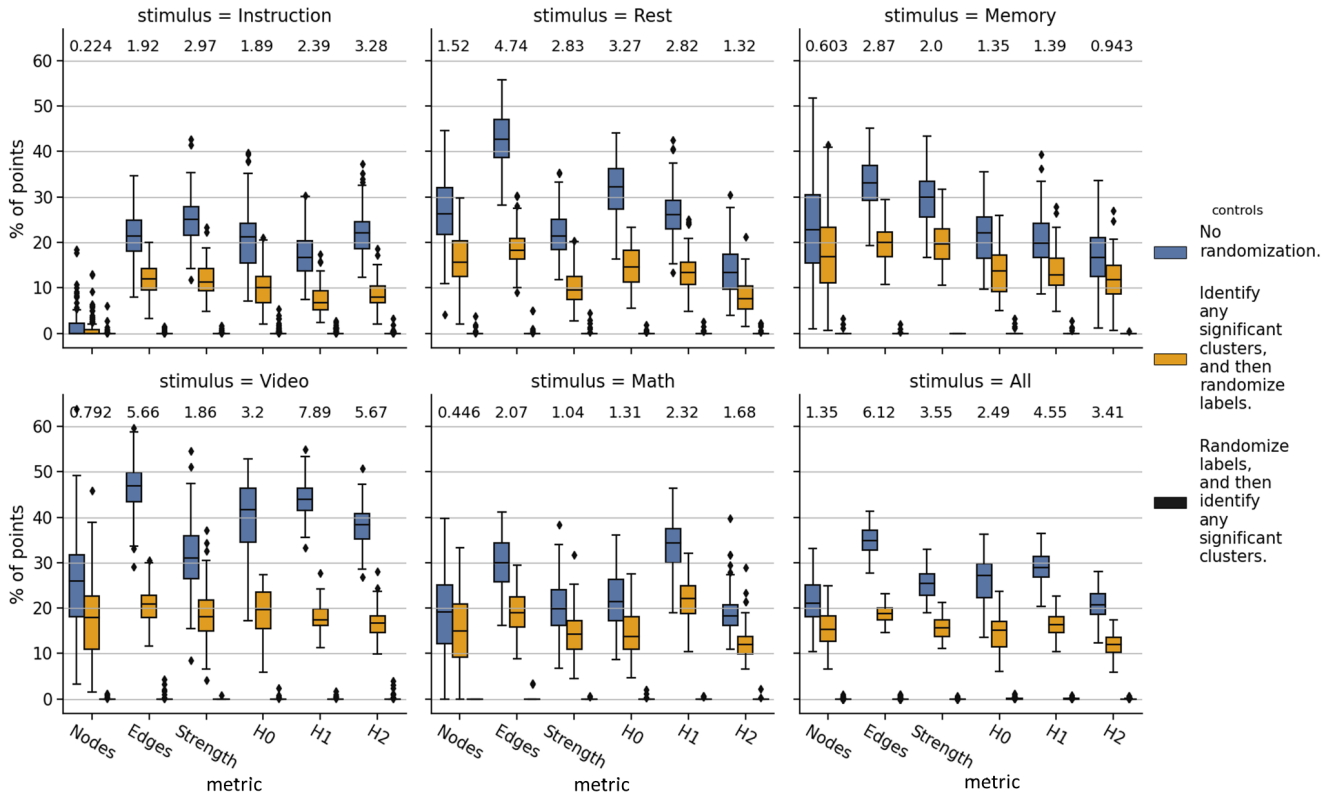

FIG. 4: **Comparison of task specificity for watershed regions across different metrics.** We report the percentage of time points from a given clusters having a significant amount of points from each experimental condition (blue boxplots). For those same clusters, we report the percentage of points from each experimental condition after randomly permuting point labels (yellow boxplots). Additionally, we report the effect size (Cohen's D) between these two distributions (values above boxes). A second null distribution (black boxplots) shows the false positive rate for identifying significant clusters. The condition labeled "Totals" are a time-weighted average over the condition-wise tests.

age of points from any experimental condition by default. The effect size (Cohen's d) between this null distribution and the real distribution provides an indication of how well each embedding isolates brain states induced by distinct experimental stimuli. The second null distribution simply permutes point labels before attempting to find watershed clusters having a significant number of points from any of the 5 experimental conditions (black boxes). This second null distribution provides a good check on the rate of false positives.

Here again, the homology-based embeddings perform very well compared to embeddings constructed from simplicial overlap. This is especially the case for the "H1" metric space which presents, over all stimuli, the second highest percentage of all points within significant clusters, as well as the second highest effect size.

It should be here noted that, while the "edge" metric space presents the highest percentage of points and the highest effect size, this metric space performs most poorly at generalizing across volunteers, being skewed furthest to the left in (figure 3). One might infer that the enhanced segmentation among stimuli in the edge embedding is an effect of edge-overlap distances co-localizing temporally adjacent points, and not of improved differentiation among brain states (see Discussion).

By contrast, the $H1$ metric space recovers distinct and broadly generalizable brain states. It is interesting to note that, of all the homology-based metrics, the embeddings using Wasserstein distances in $H2$ provide the worst segmentation over stimuli. While this may indicate that the topological homology of TVFC is restricted to very low dimensions, the computationally-motivated coarsening of voxelwise information into 328 brain regions also limits the appearance of higher-dimensional homologies.

Across the various stimuli, the video stimulus is, in all cases, the easiest condition to segment into unique clusters. This fact may owe itself to the uniformity of the brain-state as individuals watch simple visual stimuli (fish idling in a tank). However, it is also possible that this increased performance may be due to the reduction of the total number of points from the video task relative to the other tasks. This hypothesis is not supported by the level of segmentation obtained for the "instruction" stimulus, for which a lower than average percentage of points are found.

It is interesting to note that homology-based metrics perform differently across the various stimuli. The $H0$ metric performs better than the $H1$ metric when volunteers are asked to rest and fixate on a cross. The $H1$ metric performs better on the mathematics task. And while the improved performance for the $H1$ metric during the math task may be due to the selection of large clusters (i.e., a large percentage of points in the null distribution), the large effect size of real vs random results from the $H1$ metric suggests that the math task is easy for the $H1$ metric to identify.

Regarding the memory task: it is unclear why the homology based metric co-localizes so few data points from this task into unique clusters. Perhaps, looking into volunteer performance in the memory task will offer some possible elaborations.

C. Task performance

"Mean performance graphs" were computed for each task, and each valenced performance level as specified in the Methods. Within each task, performance was valenced as having either more correct responses, or fewer correct responses than the mean for the task. The results are displayed in figure 5.

Of particular note are the distances computed across the valenced performance levels, but within the same category of metric space (figure 5, white annotations). These comparisons directly observe how sensitive each TVFC metric is to volunteer-wise internal variations in task engagement under common external stimuli.

Across the board, mean graphs computed from homology-based embeddings show decreased similarity between valenced, task-specific performance. This is especially true when comparing the percentage of correct responses in the memory task. This observation suggests that homology observes differences in performance where they should be expected.

D. Visualize homology

Finally, having identified the increased utility of TVFC graph-graph metrics developed from homology to disambiguate group-general brain states, we wanted to gain some insights into what features of TVFC homology resolves. Owing to the good performance of the $H1$ metric space, we present a visualization of this metric space's mean performance graphs in figure 6. Whereas visualization of the minimal cycles associated with multiscale persistent homology is non-trivial, we opt to coarsen the original 328 brain regions, and reevaluate persistence homology over the coarsened graphs. The coarsening procedure is described in the Methods, as well as in the row labeled "Coarse Parcellation" in figure 6.

We chose two levels of coarsening from which to visualize homological cycles. Choosing 46 spatial brain parcels finds the fewest number of parcels for which homologies in $H1$ are still present across all selected mean graphs. Choosing 64 brain parcels is simply the next power of 2, and provides an indication of how homology changes when increasing the granularity by splitting nodes into their component parts. Using these clusters as starting points, and the time points associated with each mean performance graph. We recalculated TVFC matrices for each coarsened parcellation. And recomputed mean performance graphs using the original statistics over clustered UMAP embeddings. For each coarsened mean graph, minimally compressed homological cycles were calculated in $H1$ using the JavaPlex toolbox for Matlab [58].

By definition, a minimally compressed homological cycle in $H1$ is the boundary of a through-hole. It is observed as a set of edges that describe the shortest path around a disconnected 2-d surface. For the present case of weighted graphs, homologies persist between lower and upper limits of connectivity thresholds. Within those thresholds a bounding homological cycle traces the path having the minimum number of connections between a set of nodes. The brain regions included in each path tend to form a chain of relatively strong coherence linking the included nodes into a cycle.

We display the persistence diagrams for each mean graph in the last column of the figure. Most cycles are born in the second quarter of the range of coherence distances (0.2 to 0.4), and die around the middle of the range of coherence distances (0.4 to 0.5). Interestingly, these birth and death times span the threshold for significant wavelet coherence, 0.4, as defined relative to an AR1 model of the input data (see part B of supplemental figure A.1).

Thus, the cycles we observe might be interpreted as a functional architecture that routes information along a coordinated processing path. And, that the coordinating cycle only dies after inter-parcel coherence is no longer distinct from noise.

In the coarsened graphs, the minimally compressed homological cycles having the longest lifetimes (black dashed lines) express several trends. Notably, most conditions involve connected cycles that pass through the bilateral cingulate cortex (#1 in figure 6). The Gordon 333 atlas specifies this region as belonging to the cingulo-opercular task control network. Few direct connections are made between the bilateral cingulate cortex. Rather, these connections are mediated by the portion of the cingulo-opercular network in the left insula (#2). Other regions including the right visual cortex (#3), and the right right medial posterior node of the default mode network (#4). Anterior portions of the default mode network (#5) after increasing the granularity to 64 spatial regions. Additionally, the "more correct" mean graph of the mathematics task also recruits the dorsal attention network (#6).

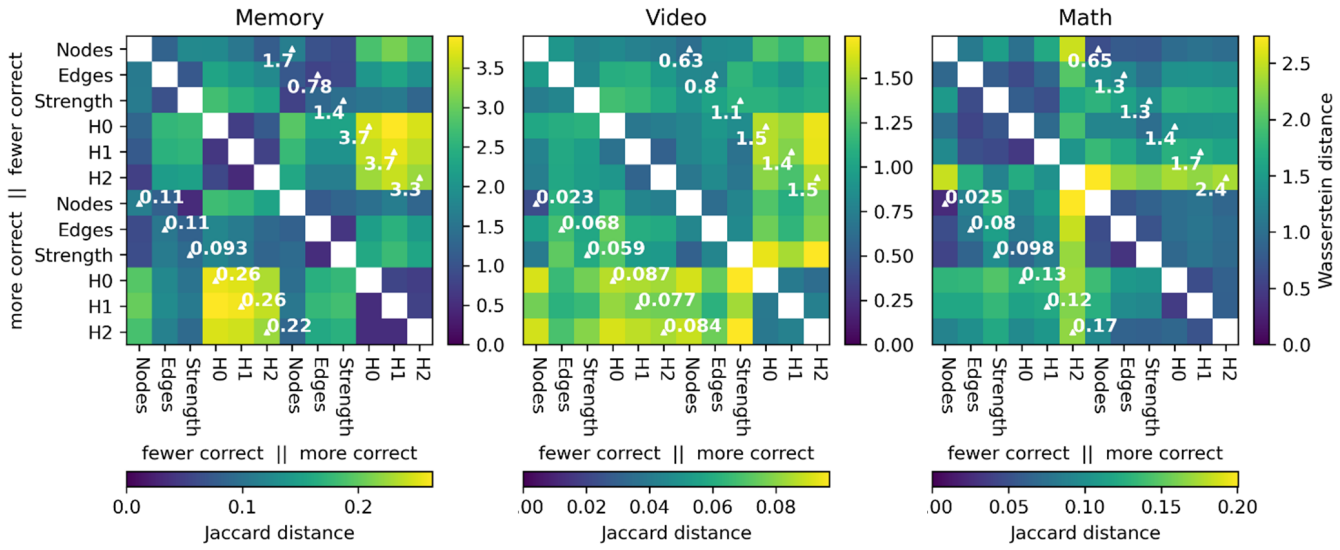

FIG. 5: **Distances between mean graphs corresponding to different task performance levels**. Mean performance graphs are calculated as in the Methods. The graph-graph distances between mean performance graphs is computed using both the Jaccard distance between 1-dimensional simplices, and for the Wasserstein distance between 1-dimensional persistence diagrams. For each of the six metric-spaces, the numeric value of distances between mean graphs is displayed.

IV. Discussion

Brain function is believed to emerge from extensive coordination among brain nuclei. But those features that typify state-specific brain organization remain a subject of intense and ongoing research [3, 35]. Describing interregional coordination across time via the framework of instantaneous coherence allows us to discern what features of the coherence graph best characterizes brain states. To do so, we compared the performance of two broad classes of TVFC metric spaces: one based upon the overlap of graph weights, and the other based upon the respective graph homological structure.

As a quantitative tool, the homological approach is tailor-made for defining topological similarities among metric spaces [8]. Indeed, fMRI studies have implemented persistent homology to discern group-level FC differences in task performance [30], and with respect to pharmacological treatments [46]. Similar findings are observed in MEG data [16]. Stateful segmentation was also achieved from homological features in $H0$ for 8-channel EEG TVFC as volunteers engaged in a visuo-motor task [67].

The results of the presents study provide ample evidence that the homology of coherence-based TVFC effectively disambiguates experimentally defined brain states in the population-general brain. Moreover, the results observe connectivity motifs in the homological cycles of mean graphs. These motifs provide novel insights into how brain function proceeds from spontaneous self-assembly over multiple scales of interaction.

By contrast, the performance of more traditional means for comparing TVFC graphs tends to be worse at distinguishing states (figure 4, "Totals"). This is especially true for the simplicial overlap distance across edges which tends to cluster volunteers separately (figure 3).

Interestingly, Finn et al. (2015) observed that the edge overlap distance from correlation-based FC matrices constitutes a biometric "fingerprint" of the individual [19]. Fingerprinting was found to be especially strong during the resting condition. Similarly, in the present study, edge-overlap distance is particular good at finding resting-state conditions. Moreover, the present findings of increased state segmentation during the video task is supported by the improved capacity to predict volunteer traits using naturalistic movie stimuli [18]. These findings suggest that the present observations of volunteer-specific edge overlap distances may be a long-term feature of the individual brain rather than simply an artifact of short-term autocorrelations.

Collectively, these results indicate that the edge overlap distance presents features that are similar within-volunteers, over both short and long time spans.

Distances measured with respect to properties of nodes—i.e., node activity, and node strength—are more group general. However, they perform worse at segmenting experiments Utilizing the same dataset as the present study [24], Saggar et al (2018) computed a node overlap distance to visualize two-dimensional mappings of within-volunteer temporal similarity [52]. In the majority of cases, the visualizations depict smooth transitions across the mapping. Smooth transitions are clearly depicted during the resting state, and during the transition to temporally adjacent task states. However, for some

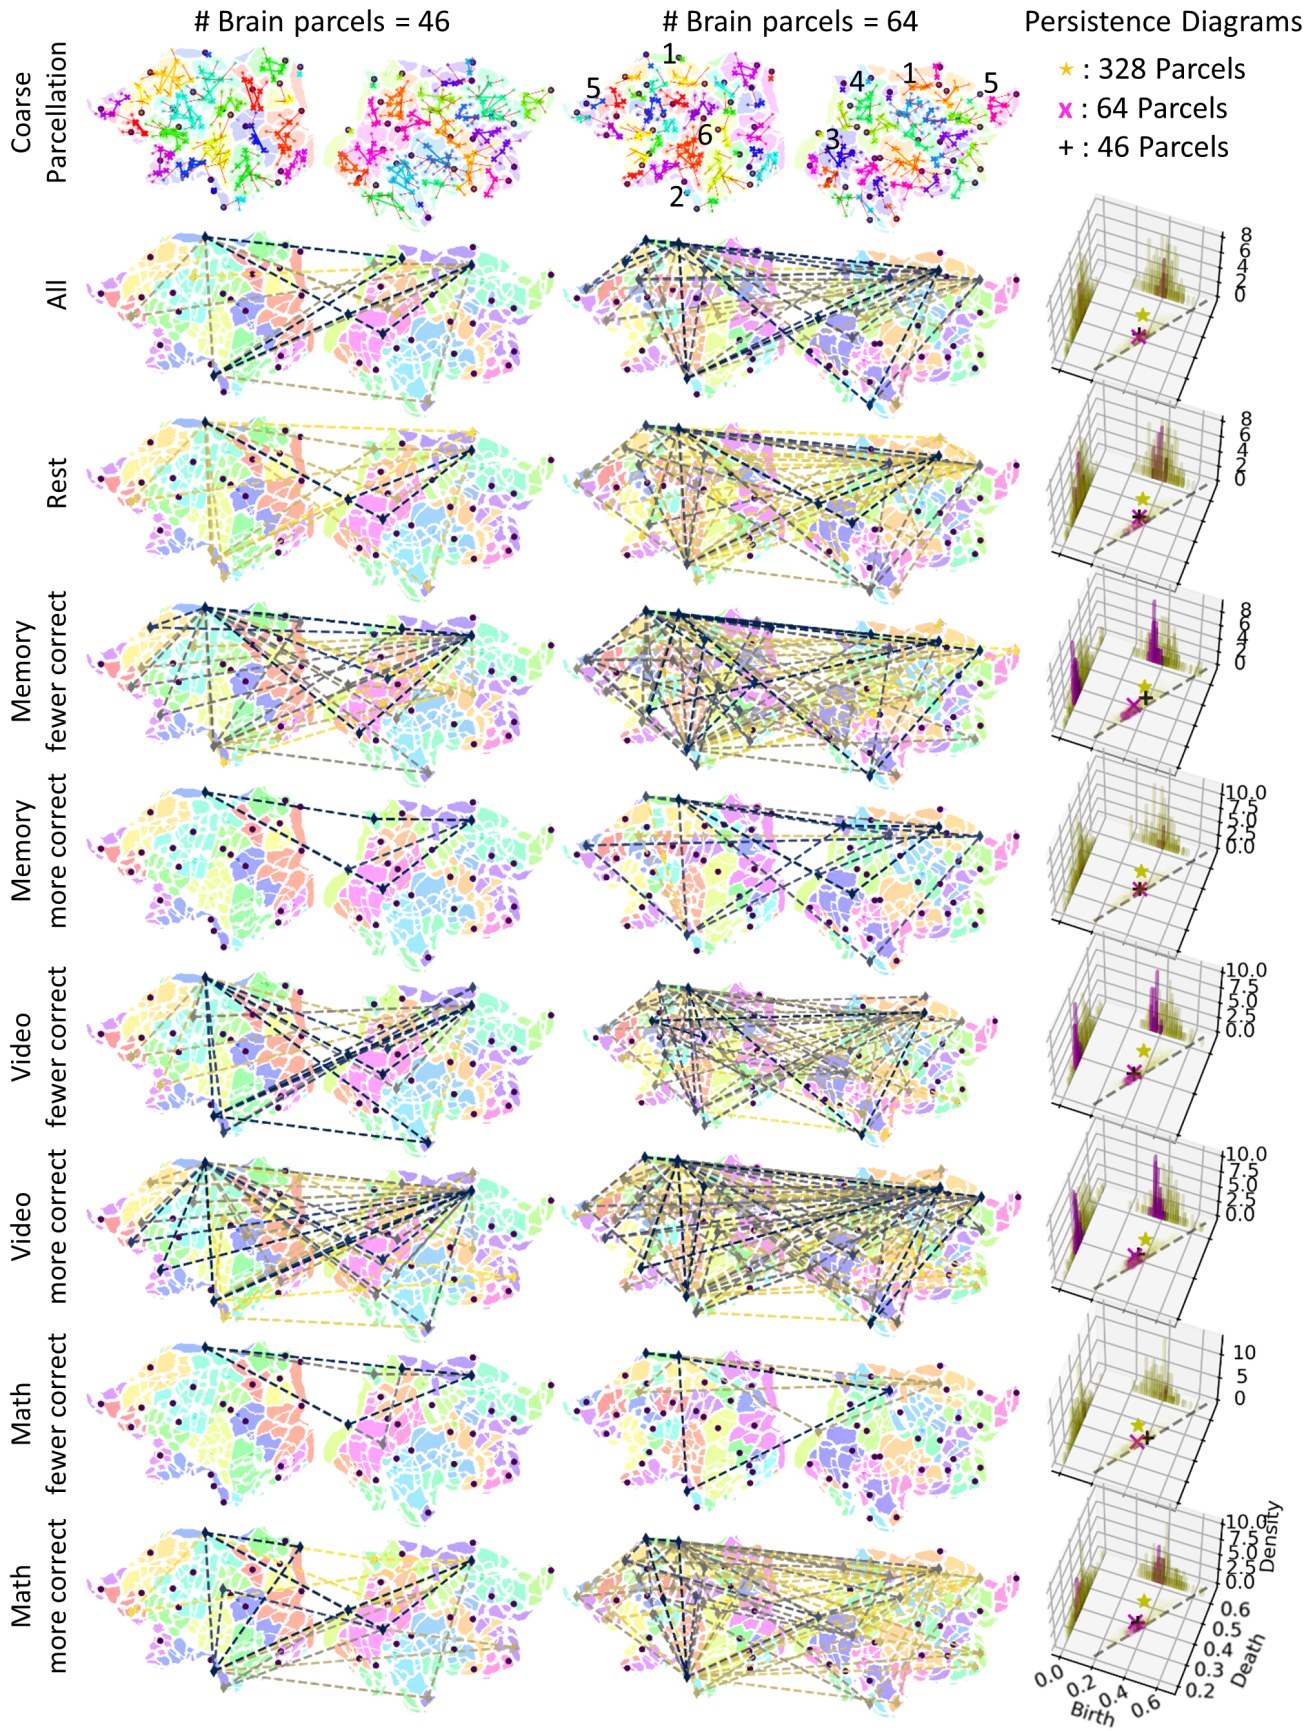

FIG. 6: (Continued on the following page.)

FIG. 6: **Illustrative examples of persistent homological cycles from the H_1 mean graphs referred to in figure 5**. The number of nodes in each graph was reduced to limit computational load and to improve visualization. Spatial coarsening was conducted as in the Methods. The original 328 brain regions are recolored to display their coarsened clusters. Black dots point one of the original 328 brain parcels that has the highest weighted degree within the coarsened parcel. Numbers refer to brain regions indicated in the body of the manuscript. 1: cingulate cortex, 2: insula, 3: visual cortex, 4: posterior default mode network, 5: anterior default mode network, 6: dorsal attention network. Darker lines represent homological cycles having longer lifetimes. Diamonds having the same color as lines highlight the clusters included in each cycle. Additionally, the column labeled "Persistence Diagrams" displays the persistence diagrams from each mean graph over each level of granularity. Projecting persistence diagrams onto the z axis highlights the density of homological cycles across lifespans. The larger markers indicate the longest lived homology.

volunteers, the mapping depicts modularized transitions within the context of a single experiment.

Using a complimentary dataset [23], Billings et al. (2017) also computed maps of node overlap distances [6]. Distances were mapped across the time points from a population of volunteers. These authors also observed a general pattern of variable activity across the map, punctuated by moments of highly co-localized brain states.

It is interesting to consider that the "strength" metric is very similar to the metric over 0-dimensional nodes. The complexity of the problem is the same, as the dimension of each metric's distance vectors equals the number of nodes. But while improved results are gathered from the strength metric during task epochs, node activity overlap better characterizes populations during rest. Given the findings from Billings (2017) and Saggar (2018) that rest is mapped to a different space than tasks; and given that the current results show the $H0$ metric space to out-perform the $H1$ metric space during rest; it is interesting to think that rest and task differ in terms of their fundamental topologies.

The empirical and theoretical support for emphasizing topological descriptors beyond apparent network connectivity has prompted several authors to reinterpret neuronal networks as functioning through their explicitly topological organization [12, 20, 32, 49, 50, 55, 57]. Indeed, there is ample experimental evidence that n-dimensional cliques and homological cavities in the mesoscopic space of structural brain images reflect known brain networks [54]. Further findings evidence that n-dimensional cliques and homologies encode microscopic interactions among neuronal circuits within the hippocampal place field [2, 13, 21] and in the somatomotor representation of the head [10]. Our results provide further support for the utility of the topological approach to discern the evolution of brain states through time, and to better comprehend the brain's multiscale complexities.

Utilizing a spatially coarsened brain parcellation to visualize the homological cycles that form the boundaries of homologies is one way to better grasp the topological underpinnings of TVFC. There is an evident role for the bilateral cingulate and the left insula in establishing homological cycles. Cycles often include regions for visual sensation and default mode activity (see figure 6) Interestingly, the cingulate, insula, and thalamus are key components of the cingulo-opercular task control network, understood to be active during tonic attention [51]. In addition, the mathematics task recruits the dorsal attention network, known to be active during acute tasks [51]. Thus it follows that coarse homology is describing the integration of sensory and attention networks [66].

Pursuing the hypothesis that functional BOLD homology highlights preferential connectivity in the attention networks may offer promising insights into the preferential route of information flow across distributed functional networks. Homology may be especially useful when topological features are expected to be important, i.e., when one expects patterned connectivity among clusters in $H0$ AND/or cycles in $H1$. Certainly there are ongoing questions regarding the pathways through which otherwise structurally distributed brain networks form TVFC networks [14].

V. Limitations and future directions

Future research should strive to make a more detailed catalogue of the homologies that commonly appear among brain regions. While the present study resorted to a very coarse brain parcellation to visualize homology (see figure 6), the effort to catalogue minimal cycles over more granular parcellations is an attainable goal. Moreover, given that sliced Wasserstein distances between persistence diagrams only preserve information about the lifetime of homologies across a filtration, and ignore information about the nodes included in the homology's equivalence class; future work that includes information about the particular subset of brain regions forming a homological boundary would certainly improve the capacity to distinguish states of the brain. This catalogue ought to continue to leverage controlled changes in the spatial resolution to observe multivariate persistent homology.

Another limitation of the present study is the reliance on clustering in the low-dimensional space. Even while low-dimensional embeddings provide an efficient means for visualizing data, there is always some loss of information. For instance, the UMAP method for embedding point cloud data transduces an explicit nearest-neighbor approximation of the high-dimensional simplicial complex into the low-dimensional space. This approxima-

tion may be causally related to the observation that metric spaces based upon, especially, 1-dimensional simplicial overlap organize into temporally-adjacent clusters. While edge overlap may be a volunteer-specific trait. And while the trait may be partially alleviated by deconvolution of the volunteer-specific hemodynamic response function. Future work that biases the low-dimensional embedding in a more appropriate way—perhaps by learning a transductive vector embedding as in [1]—may offer some additional improvements. In any case, approaches that circumvent dimensionality reduction entirely by operating in the native high-dimensional space may offer the most general solution to the loss of information during low-dimensional embedding.

VI. Conclusion

To understand the dynamic self-organization of complex systems like the brain, it helps to view system dynamics through lenses that highlight the presence and the structure of complexes. Given the kinds of weighted graphs typical of TVFC analysis, persistent homology is well-suited for interpreting complexes of brain regions. The view from homology outperforms more traditional graph metrics —i.e., the activity measures of 0-dimensional nodes, or the weights of 1-dimensional edges—at the task of segmenting experimentally defined brain states into patterns that generalize well across multiple volunteers. By examining coarsened approximations of these homological features, we gain some intriguing insights into how brains operate via the recruitment of a distributed set of brain regions into distinct functional topologies.

[1] Yunsheng Bai, Hao Ding, Yang Qiao, Agustin Marinovic, Ken Gu, Ting Chen, Yizhou Sun, and Wei Wang. Unsupervised Inductive Graph-Level Representation Learning via Graph-Graph Proximity. apr 2019.

[2] Edward Basso, Mamiko Arai, and Yuri Dabaghian. Gamma Synchronization Influences Map Formation Time in a Topological Model of Spatial Learning. *PLOS Computational Biology*, 12(9):e1005114, sep 2016.

[3] Demian Battaglia and Andrea Brovelli. Functional connectivity and neuronal dynamics: insights from computational methods. Technical report, 2020.

[4] Federico Battiston, Giulia Cencetti, Iacopo Iacopini, Vito Latora, Maxime Lucas, Alice Patania, Jean-Gabriel Young, and Giovanni Petri. Networks beyond pairwise interactions: structure and dynamics. *Physics Reports*, 2020.

[5] Jacob C. W. Billings, Garth J. Thompson, Wen-Ju Pan, Matthew E. Magnuson, Alessio Medda, and Shella Keilholz. Disentangling Multispectral Functional Connectivity With Wavelets. *Frontiers in Neuroscience*, 12:812, nov 2018.

[6] Jacob C.W. Billings, Alessio Medda, Sadia Shakil, Xiaohong Shen, Amrit Kashyap, Shiyang Chen, Anzar Abbas, Xiaodi Zhang, Maysam Nezafati, Wen Ju Pan, Gordon J. Berman, and Shella D. Keilholz. Instantaneous brain dynamics mapped to a continuous state space. *NeuroImage*, 162(Supplement C):344–352, nov 2017.

[7] Bharat Biswal, F Zerrin Yetkin, Victor M Haughton, and James S Hyde. Functional connectivity in the motor cortex of resting human brain using echo-planar mri. *Magnetic Resonance in Medicine*, 34(4):537–541, oct 1995.

[8] Gunnar Carlsson. Topology and Data. Technical Report 2, 2009.

[9] Mathieu Carrì, Marco Cuturi, and Steve Oudot. Sliced Wasserstein Kernel for Persistence Diagrams. Technical report.

[10] Rishidev Chaudhuri, Berk Gerçek, Biraj Pandey, Adrien Peyrache, and Ila Fiete. The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nature Neuroscience*, 22(9):1512–1520, 2019.

[11] Moo K. Chung, Hyekyoung Lee, Alex Di Christofano, Hernando Ombao, and Victor Solo. Exact topological inference of the resting-state brain networks in twins. *Network Neuroscience*, 3(3):674–694, jan 2019.

[12] Carina Curto. What can Topology Tell us About the Neural Code? Technical report.

[13] Yuri Dabaghian, Vicky L. Brandt, and Loren M. Frank. Reconceiving the hippocampal map as a topological template. *eLife*, 3:e03476, 2014.

[14] Jessica S Damoiseaux and Michael D Greicius. Greater than the sum of its parts: a review of studies combining structural connectivity and resting-state functional connectivity. *Brain Structure and Function*, 213(6):525–533, 2009.

[15] C Du, N D Volkow, A P Koretsky, and Y Pan. Low-frequency calcium oscillations accompany deoxyhemoglobin oscillations in rat somatosensory cortex. *Proc Natl Acad Sci U S A*, 111(43):E4677–86, 2014.

[16] Ali Nabi Duman, Ahmet Emin Tatar, Harun Pirim, Ali Nabi Duman, Ahmet Emin Tatar, and Harun Pirim. Uncovering Dynamic Brain Reconfiguration in MEG Working Memory n-Back Task Using Topological Data Analysis. *Brain Sciences*, 9(6):144, jun 2019.

[17] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. *Discrete and Computational Geometry*, 28(4):511–533, nov 2002.

[18] Emily S Finn and Peter A Bandettini. Movie-watching outperforms rest for functional connectivity-based prediction of behavior. *bioRxiv*, page 2020.08.23.263723, aug 2020.

[19] Emily S Finn, Xilin Shen, Dustin Scheinost, Monica D Rosenberg, Jessica Huang, Marvin M Chun, Xenophon Papademetris, and R Todd Constable. Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. *Nature neuroscience*, 18(11):1664–1671, 2015.

[20] Chad Giusti, Robert Ghrist, and Danielle S Bassett. Two's company, three (or more) is a simplex: Algebraic-

topological tools for understanding higher-order structure in neural data. Technical report, 2016.

[21] Chad Giusti, Eva Pastalkova, Carina Curto, and Vladimir Itskov. Clique topology reveals intrinsic geometric structure in neural correlations. 2015.

[22] Matthew F. Glasser, Timothy S. Coalson, Emma C. Robinson, Carl D. Hacker, John Harwell, Essa Yacoub, Kamil Ugurbil, Jesper Andersson, Christian F. Beckmann, Mark Jenkinson, Stephen M. Smith, and David C. Van Essen. A multi-modal parcellation of human cerebral cortex. *Nature*, 536(7615):171–178, aug 2016.

[23] Matthew F. Glasser, Stamatios N. Sotiropoulos, J. Anthony Wilson, Timothy S. Coalson, Bruce Fischl, Jesper L. Andersson, Junqian Xu, Saad Jbabdi, Matthew Webster, Jonathan R. Polimeni, David C. Van Essen, and Mark Jenkinson. The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage*, 80:105–124, oct 2013.

[24] Javier Gonzalez-Castillo, Colin W Hoy, Daniel A Handwerker, Meghan E Robinson, Laura C Buchanan, Ziad S Saad, and Peter A Bandettini. Tracking ongoing cognition in individuals using brief, whole-brain functional connectivity patterns. *Proceedings of the National Academy of Sciences of the United States of America*, 112(28):8762–7, jul 2015.

[25] Evan M Gordon, Timothy O Laumann, Babatunde Adeyemo, Jeremy F Huckins, William M Kelley, and Steven E Petersen. Generation and Evaluation of a Cortical Area Parcellation from Resting-State Correlations. *Cerebral cortex (New York, N.Y. : 1991)*, 26(1):288–303, jan 2016.

[26] Marco Guerra, Alessandro De Gregorio, Ulderico Fugacci, Giovanni Petri, and Francesco Vaccarino. Homological scaffold via minimal homology bases. *arXiv preprint arXiv:2004.11606*, 2020.

[27] Enrique CA Hansen, Demian Battaglia, Andreas Spiegler, Gustavo Deco, and Viktor K Jirsa. Functional connectivity dynamics: modeling the switching behavior of the resting state. *Neuroimage*, 105:525–535, 2015.

[28] W Hasenkamp, C D Wilson-Mendenhall, E Duncan, and L W Barsalou. Mind wandering and attention during focused meditation: a fine-grained temporal analysis of fluctuating cognitive states. *Neuroimage*, 59(1):750–760, 2012.

[29] R Matthew Hutchison, Thilo Womelsdorf, Elena A Allen, Peter A Bandettini, Vince D Calhoun, Maurizio Corbetta, Stefania Della Penna, Jeff H Duyn, Gary H Glover, Javier Gonzalez-Castillo, et al. Dynamic functional connectivity: promise, issues, and interpretations. *Neuroimage*, 80:360–378, 2013.

[30] Esther Ibáñez-Marcelo, Lisa Campioni, Angkoon Phinyomark, Giovanni Petri, and Enrica L. Santarcangelo. Topology highlights mesoscopic functional equivalence between imagery and perception: The case of hypnotizability. *NeuroImage*, 200:437–449, oct 2019.

[31] Adam Kohn, Amin Zandvakili, and Matthew A Smith. Correlations and brain states: from electrophysiology to functional imaging. *Current Opinion in Neurobiology*, 19(4):434–438, 2009.

[32] Guilia Lerda. *University of Torino, Torino, Italy.* PhD thesis, University of Turino, 2016.

[33] Jingfeng M Li, William J Bentley, Abraham Z Snyder, Marcus E Raichle, and Lawrence H Snyder. Functional connectivity arises from a slow rhythmic mechanism. *Proceedings of the National Academy of Sciences*, 112(19):E2527–E2535, 2015.

[34] Louis David Lord, Paul Expert, Henrique M. Fernandes, Giovanni Petri, Tim J. Van Hartevelt, Francesco Vaccarino, Gustavo Deco, Federico Turkheimer, and Morten L. Kringelbach. Insights into brain architectures from the homological scaffolds of functional connectivity networks. *Frontiers in Systems Neuroscience*, 10(NOV), nov 2016.

[35] Daniel J. Lurie, Daniel Kessler, Danielle S. Bassett, Richard F. Betzel, Michael Breakspear, Shella Kheilholz, Aaron Kucyi, Raphaël Liégeois, Martin A. Lindquist, Anthony Randal McIntosh, Russell A. Poldrack, James M. Shine, William Hedley Thompson, Natalia Z. Bielczyk, Linda Douw, Dominik Kraft, Robyn L. Miller, Muthuraman Muthuraman, Lorenzo Pasquini, Adeel Razi, Diego Vidaurre, Hua Xie, and Vince D. Calhoun. Questions and controversies in the study of time-varying functional connectivity in resting fMRI. *Network Neuroscience*, 4(1):30–69, jan 2020.

[36] M E Magnuson, G J Thompson, W J Pan, and S D Keilholz. Effects of Severing the Corpus Callosum on Electrical and BOLD Functional Connectivity and Spontaneous Dynamic Activity in the Rat Brain. *Brain Connect*, 4(1):15–29, 2014.

[37] J A Maldjian, E M Davenport, and C T Whitlow. Graph theoretical analysis of resting-state MEG data: Identifying interhemispheric connectivity and the default mode. *Neuroimage*, 96:88–94, 2014.

[38] Daniel S. Margulies, Joachim Böttger, Aimi Watanabe, and Krzysztof J. Gorgolewski. Visualizing the human connectome. *NeuroImage*, 80:445–461, oct 2013.

[39] N Marwan, M Thiel, and N R Nowaczyk. Application of the cross wavelet transform and wavelet coherence togeophysical time series. *Nonlinear Processes in Geophysics*, 9(3):325–331, 2002.

[40] Leland McInnes, John Healy, and James Melville. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. feb 2018.

[41] Yuriy Mileyko, Sayan Mukherjee, and John Harer. Probability measures on the space of persistence diagrams. *Inverse Problems*, 27(12):124007, dec 2011.

[42] F Palhano-Fontes, K C Andrade, L F Tofoli, A C Santos, J A Crippa, J E Hallak, S Ribeiro, and D B de Araujo. The Psychedelic State Induced by Ayahuasca Modulates the Activity and Connectivity of the Default Mode Network. *PLoS ONE*, 10(2):e0118143, 2015.

[43] W J Pan, G J Thompson, M E Magnuson, D Jaeger, and S Keilholz. Infraslow LFP correlates to resting-state fMRI BOLD signals. *Neuroimage*, 74:288–297, 2013.

[44] Alice Patania, Francesco Vaccarino, and Giovanni Petri. Topological analysis of data. *EPJ Data Science*, 6:1–6, 2017.

[45] G. 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(101):20140873–20140873, oct 2014.

[46] Giovanni Petri, Paul Expert, Federico Turkheimer, Robin Carhart-Harris, David Nutt, Peter J Hellyer, and Francesco Vaccarino. Homological scaffolds of brain functional networks. *Journal of The Royal Society Interface*, 11(101):20140873, 2014.

[47] Giovanni Petri, Martina Scolamiero, Irene Donato, and Francesco Vaccarino. Topological strata of weighted complex networks. *PloS one*, 8(6), 2013.

[48] Angkoon Phinyomark, Esther Ibanez-Marcelo, and Giovanni Petri. Resting-state fmri functional connectivity: Big data preprocessing pipelines and topological data analysis. *IEEE Transactions on Big Data*, 3(4):415–428, 2017.

[49] Mario Rasetti. The 'Life Machine': A Quantum Metaphor for Living Matter. *International Journal of Theoretical Physics*, 56(1):145–167, 2017.

[50] Michael W. Reimann, Max Nolte, Martina Scolamiero, Katharine Turner, Rodrigo Perin, Giuseppe Chindemi, Paweł Dłotko, Ran Levi, Kathryn Hess, and Henry Markram. Cliques of Neurons Bound into Cavities Provide a Missing Link between Structure and Function. *Frontiers in Computational Neuroscience*, 11:48, 2017.

[51] Sepideh Sadaghiani and Mark D'Esposito. Functional characterization of the cingulo-opercular network in the maintenance of tonic alertness. *Cerebral Cortex*, 25(9):2763–2773, sep 2015.

[52] Manish Saggar, Olaf Sporns, Javier Gonzalez-Castillo, Peter A. Bandettini, Gunnar Carlsson, Gary Glover, and Allan L. Reiss. Towards a new approach to reveal dynamical organization of the brain using topological data analysis. *Nature Communications*, 9(1), 2018.

[53] Ann Sizemore, Chad Giusti, and Danielle S Bassett. Classification of weighted networks through mesoscale homological features. *Journal of Complex Networks*, 5(2):245–273, 2017.

[54] Ann E. Sizemore, Chad Giusti, Ari Kahn, Jean M. Vettel, Richard F. Betzel, and Danielle S. Bassett. Cliques and cavities in the human connectome. *Journal of Computational Neuroscience*, 44(1):115–145, 2018.

[55] Ann E Sizemore, Jennifer E Phillips-Cremins, Robert Ghrist, and Danielle S Bassett. The importance of the whole: Topological data analysis for the network neuroscientist. 2018.

[56] Stephen M Smith, Peter T Fox, Karla L Miller, David C Glahn, P Mickle Fox, Clare E Mackay, Nicola Filippini, Kate E Watkins, Roberto Toro, Angela R Laird, and Christian F Beckmann. Correspondence of the brain's functional architecture during activation and rest. *Proceedings of the National Academy of Sciences*, 2009.

[57] Bernadette Stolz. *Computational Topology in Neuroscience*. PhD thesis, 2014.

[58] Andrew Tausz, Mikael Vejdemo-Johansson, and Henry Adams. JavaPlex: A research software package for persistent (co)homology. In Han Hong and Chee Yap, editors, *Proceedings of ICMS 2014*, Lecture Notes in Computer Science 8592, pages 129–136, 2014. Software available at `http://appliedtopology.github.io/javaplex/`.

[59] W H Thompson and P Fransson. The frequency dimension of fMRI dynamic connectivity: Network connectivity, functional hubs and integration in the resting brain. *Neuroimage*, 121:227–242, 2015.

[60] Christopher Torrence, Gilbert P. Compo, Christopher Torrence, and Gilbert P. Compo. A Practical Guide to Wavelet Analysis. *Bulletin of the American Meteorological Society*, 79(1):61–78, jan 1998.

[61] Christopher Torrence, Peter J. Webster, Christopher Torrence, and Peter J. Webster. Interdecadal Changes in the ENSO–Monsoon System. *Journal of Climate*, 12(8):2679–2690, aug 1999.

[62] Christopher Tralie, Nathaniel Saul, and Rann Bar-On. Ripser.py: A lean persistent homology library for python. *The Journal of Open Source Software*, 3(29):925, Sep 2018.

[63] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(2579-2605):85, 2008.

[64] D C Van Essen, K Ugurbil, E Auerbach, D Barch, T E Behrens, R Bucholz, A Chang, L Chen, M Corbetta, S W Curtiss, S Della Penna, D Feinberg, M F Glasser, N Harel, A C Heath, L Larson-Prior, D Marcus, G Michalareas, S Moeller, R Oostenveld, S E Petersen, F Prior, B L Schlaggar, S M Smith, A Z Snyder, J Xu, and E Yacoub. The Human Connectome Project: a data acquisition perspective. *Neuroimage*, 62(4), 2012.

[65] B A Vogt, L Vogt, and S Laureys. Cytology and functionally correlated circuits of human posterior cingulate areas. *Neuroimage*, 29(2):452–466, 2006.

[66] Simone Vossel, Joy J. Geng, and Gereon R. Fink. Dorsal and ventral attention systems: Distinct neural circuits but collaborative roles. *Neuroscientist*, 20(2):150–159, 2014.

[67] Jaejun Yoo, Eun Young Kim, Yong Min Ahn, and Jong Chul Ye. Topological persistence vineyard for dynamic functional brain connectivity during resting and gaming stages. *Journal of Neuroscience Methods*, 267:1–13, jul 2016.

A. Supplemental Materials

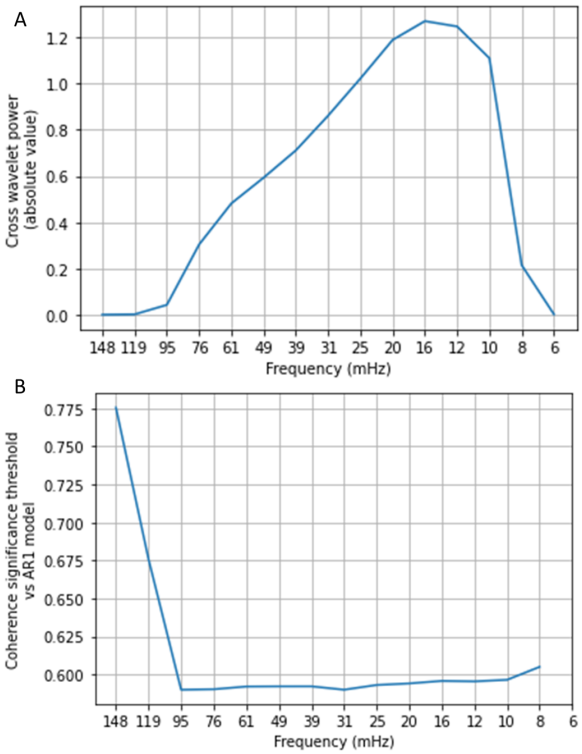

FIG. A.1: **Describes the mean values of the input data across frequency bands.** Part A of the figure displays the absolute value of the mean cross wavelet power. Frequencies between 0.02 Hz and 0.01 Hz are most likely to contribute to the mean wavelet coherence used to quantify TVFC edges. Part B of the figure displays a significance threshold level indicating significantly high coherence. The threshold was calculated as an average against the background power spectrum. A distribution over the background power spectrum was calculated from 300 lag-1 approximations of each time series. Following [60], the 95% confidence interval is the product of the background power spectrum and the 95th percentile value of a chi-squared distribution with two degrees of freedom.

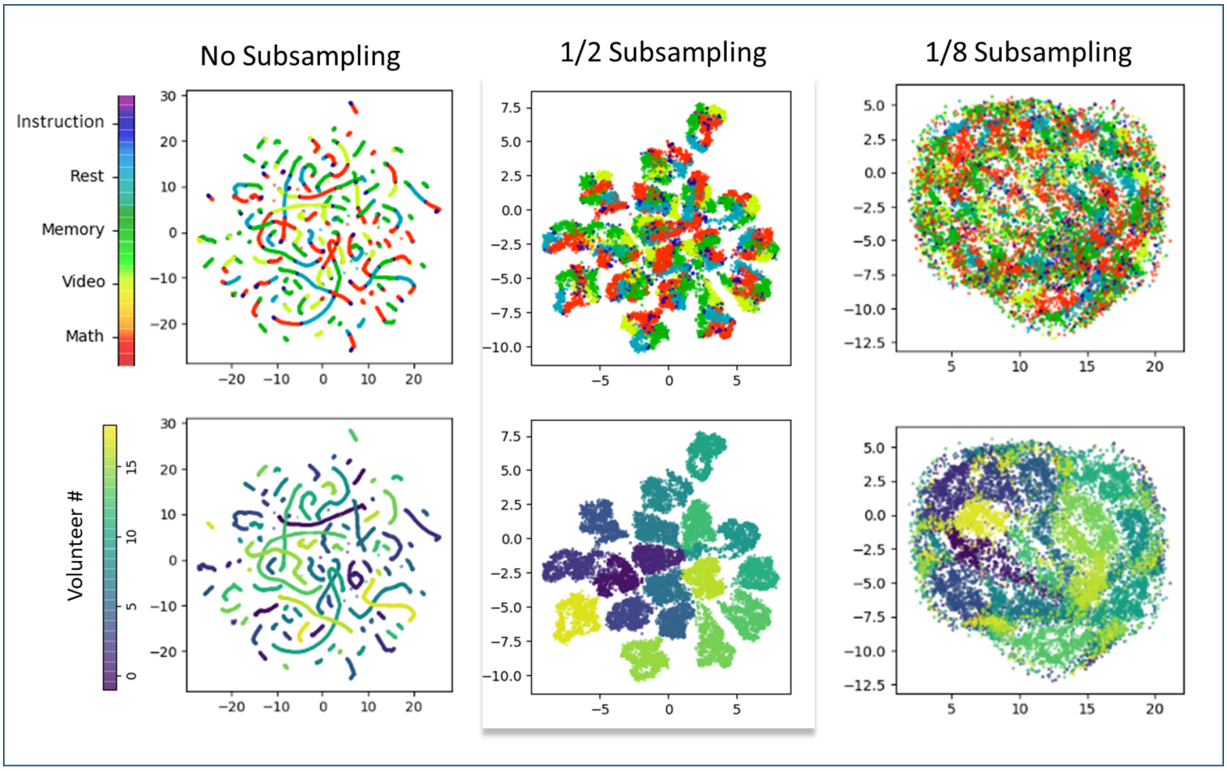

FIG. A.2: **Observes an embedding of the weighted Jaccard distance between edges at several levels of subsampling.** With no subsampling, successive time points from within the same scan retain strong attractive forces in the UMAP nearest-neighbor embedding. This produces string-like masses. Subsampling the data by half retains the similarities present within each volunteer's scan, thereby grouping volunteers into their own cluster. A noticeable degree of volunteerwise clustering—e.g., temporal self-similarity—is still present when seeding the embedding with 1/8 the total number of data points. None-the-less, some mutual attraction does emerge among edge-centric brain states measured across multiple volunteers.