Limitations of Topological Data Analysis for event-related fMRI

Cameron T. Ellis[1*], Michael Lesnick[2,3], Greg Henselman-Petrusek[2], Bryn Keller[4], & Jonathan D. Cohen[2]

[1]Department of Psychology, Yale University, New Haven, CT, USA

[2]Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA

[3]Department of Mathematics, University at Albany, Albany, NY, USA

[4]Intel Labs, Hillsboro, OR, USA

Article type: Methods

Abstract: 193 / 200 words

Main text: 3330 / 6000 words

Display elements: 2

References: 24

Keywords: topological data analysis, fMRI, simulation, image analysis, event-related design, representation

Corresponding Author: Cameron Ellis, cameron.ellis@yale.edu

Recent fMRI research has shown that perceptual and cognitive representations are instantiated in high-dimensional multi-voxel patterns in the brain. However, the methods for detecting these representations are limited. Topological Data Analysis (TDA) is a new approach, based on the mathematical field of topology, that can detect unique types of geometric structures in patterns of data. Several recent studies have successfully applied TDA to the study of various forms of neural data; however, to our knowledge, TDA has not been successfully applied to data from event-related fMRI designs in ways that reveal novel structure. Event-related fMRI is very common but limited in terms of the number of events that can be run within a practical time frame and the effect size that can be expected. Here, we investigate whether TDA can identify known signals given these constraints. We use fmrisim, a Python-based simulator of realistic fMRI data, to assess the plausibility of recovering a simple topological representation under a variety of signal and design conditions. Our results suggest that TDA has limited usefulness for event-related fMRI using current methods, as fMRI data is too noisy to allow representations to be reliably identified using TDA.

Introduction

A fundamental construct in cognitive psychology and neuroscience is that of a representational space, within which knowledge is stored. A representational space is a set of dimensions along which items of a particular type are described, in which each dimension reflects features of items of that type (e.g., perceptual, semantic, evaluative, etc.). Thus, a central goal of both cognitive psychology and neuroscience is to characterize the dimensions that define representational spaces for different types of information (e.g., objects in the world, goals, classes of actions, etc.). Progress towards this goal obviously relies on methods that can identify the structure of such spaces. Cognitive neuroscience seeks evidence for such structure in patterns of neural activity.

Topological Data Analysis (TDA) holds promise for this effort, as an unbiased (that is, "exploratory") method for identifying the latent structure in high-dimensional data. Consistent with this promise, TDA has recently been applied productively to data from direct neuronal recordings. For instance, realistic models (Dabaghian, Mémoli, Frank, & Carlsson, 2012) and real data (Giusti, Pastalkova, Curto, & Itskov, 2015) suggest that the co-firing of hippocampal place cells contains topological structure. Other studies have shown that TDA can find structure in fMRI data (Bassett & Sporns, 2017). Examples of TDA being applied to fMRI data show that when a signal is periodic, TDA can capture it (Knyazeva, Orlov, Ushakov, Makarenko, & Velichkovsky, 2016), and TDA can describe the topological structure identified in the state transitions between tasks using other methods (Saggar et al., 2018). However, to our knowledge, TDA has not yet been demonstrated to directly identify structure in event-related fMRI data.

Event-related fMRI, in which discrete events are presented one after the other, is an important method for identifying mental representations. An example of a representation is the one you have of a landmark near your home (inspired by an experiment from Nielson, Smith, Sreekumar, Dennis & Sederberg, 2015). This representation has many dimensions (e.g., its location relative to your home, the type of place it is, how much you like it, how often you go there, etc.). Your representation of different landmarks around your home will have different coordinates in this representational space. To identify these representations, an event-related fMRI experiment could be carried out in which you are shown images of these landmarks one after the other, separated by a pause. Each image presentation is an event and will evoke a neural response. If these neural responses capture the spatial dimensions of these landmarks, then the physical distances in these landmarks will relate to the differences in neural representations. With this logic, it should be possible to identify patterns of activity in the brain that reflect the distance between these landmarks; that is, their topological structure. This is a simple example but it applies to other types of representational spaces such as faces (Valentine, 1991), temporal relationships (Schapiro, Rogers, Cordova, Turk-Browne, & Botvinick, 2013), etc. Some fMRI studies have generated evidence for patterns of brain activity that reflect the metric properties inherent in a set of stimuli (Nau, Schröder, Bellmund, & Doeller, 2018; Schapiro et al., 2013), but these did not use TDA.

An appeal of TDA is that it has the potential to identify the topological structure of representational spaces (e.g., contiguity relationships) without requiring that affine properties (i.e., exact metric relationships) be preserved; for example, a rubber band preserves its topological identity even as it is stretched into various shapes. This may be

important for identifying representational spaces in the brain, if these have more topological structure than strict metric relationships — something that standard methods would find harder to detect.

In order to evaluate the potential usefulness of TDA for identifying the structure of representational spaces in the brain, we simulated an event-related fMRI dataset by embedding a simple topologically-structured signal (a loop/ring) within a realistically noisy model of fMRI data. We tested different design parameters (e.g., number of event types, and samples per event type), as well as different levels of signal strength, to evaluate the extent to which TDA is able to recover the topological structure of the signal imposed on the noisy data.

Methods

Materials and Procedure

We used fmrisim (Ellis, Baldassano, Schapiro, Cai & Cohen, forthcoming) to simulate realistic fMRI data using the public code published at https://github.com/CameronTEllis/event_related_fmri_tda. This package uses raw fMRI data to estimate noise parameters and simulate realistic data. The raw data of 20 participants, each with 5 runs, from Schapiro and colleagues (2013) were used to acquire these noise parameter estimates. This data was collected on a 3T scanner (Siemens Allegra) with a 16-channel head coil and a T2* gradient-echo planar imaging sequence (TR=2s, TE=30ms, flip angle=90°, matrix=64x64, slices=34, resolution=3x3x3mm, gap=1mm).

These noise parameters were used to generate a noise volume into which a known signal was embedded. We embedded this signal in a region of interest (ROI), identified by Schapiro et al., (2013) within the left superior temporal gyrus (spanning 442 voxels). The simulated signal was generated with an event-related design in which events of different types occurred one after the other. An event type could be thought of as a unique stimulus type, like a landmark or face. Datasets were simulated with 12, 15, or 18 different event types. Each event type in this simulation elicited activation across voxels of the ROI. For each run, event types were repeated either 5 or 10 times. This resulted in simulated run durations ranging from 45 to 135 minutes (more than 100 minutes would be a long fMRI experiment and might need to be completed across multiple sessions). Figure 1 shows a summary of the methods and analysis used here.

The key feature of the simulation is that the event types elicited a neural representation with topological structure. Specifically, each event type was a point in high-dimensional space, where an event type's position in each dimension was the activity elicited by signal voxels, and the pattern of activity across all signal voxels specifies the event type's location in representational space. For the purposes of this analysis, the coordinates of the event types were arranged to have a simple topological structure: a loop (Figure 1). These coordinates were specified first in two dimensions, and then an orthonormal transformation was used to embed these points in 442 dimensions (matching the number of ROI voxels) in a way that preserved Euclidean distance. This meant that each of these voxels' activity represented a dimension along which each event could vary and the pattern across the voxels determined the event's representation. Critically, the only part of this simulated brain that should have contained this kind of

representational structure was in this ROI; everywhere else in the brain served as a noise baseline.

These events each lasted for 1s, and occurred 7, 9 or 11s apart with equal probability and in a randomized order. This jittered delay ensured a clean measurement of the brain's delayed, hemodynamic response (Burock, Buckner, Woldorff, Rosen, & Dale,1998). Each voxel had a different response for each event (based on the voxel's embedding in the representation space) but the maximum response was set by a percent signal change value (0, 0.25, 0.5, 0.75, 1 or 5% change) validated in another study (Ellis et al., forthcoming). Percent signal change refers to the magnitude of evoked response relative baseline fluctuations. In typical experiments, we expect to find between 0.25% and 1% signal change, depending on the cortical area and stimulus (Gläscher, 2009). A signal change of 5% was included as a reference, representing an extremely strong signal. Each voxel's time course of evoked responses was convolved with a double gamma hemodynamic response function (Friston et al., 1998) to imitate the brain's response to event onsets. These evoked responses were then added to the noise and Z-scored across time to complete the simulation of a run. For each combination of the conditions (number of event types, repetitions per run, and percent signal change), a new simulation was made for all runs and all participants.

Analysis

To extract the evoked representation, a mass-univariate general linear model was performed on each run, following the standard approach with real fMRI data. To do this, all of the events for a given type were convolved with a double gamma hemodynamic

response function (Friston et al., 1998) to predict the neural response, creating a coefficient for each event type at each voxel. These coefficients were then averaged across runs to create a voxelwise response to each event type. These averaged responses to each event type were then used in a BrainIAK (http://brainiak.org) searchlight analysis. A searchlight analysis involves iterating the same computation over a subset ("searchlight") of voxels centered in different voxels in the brain. The searchlight is a tensor of 3 x 3 x 3 x N voxels centered on every voxel in the brain, where N was the number of distinct event types. For each searchlight, the voxels were reorganized into matrices of 27 x N, which reflects the pattern of activity in that searchlight for each event type. We then created an N x N Euclidean distance matrix by comparing the N distinct 27-voxel patterns across event types. Events evoking similar patterns of activity within a given searchlight had a low distance between them, and vice versa. We used these distance matrices in a TDA analysis for each searchlight (i.e., centered on each voxel in the brain).

We used persistent homology (Zomorodian & Carlsson, 2005), a standard TDA tool, to quantify the characteristics of each searchlight. A full description of persistent homology is outside of the scope of the present work; interested readers are pointed to standard introductions to the topic (Carlsson, 2009; Ghrist, 2008). In brief, persistent homology takes as input a distance matrix to identify topological features of the data. In the standard setup, these features are characterized as *N*-dimensional holes, where a 0-dimensional hole corresponds to a cluster, a 1-dimensional hole is a loop, a 2-dimensional hole is a hollow sphere, etc. The persistence of these features can be quantified with respect to a filtration procedure (e.g., systematic variation in the

resolution of measurement), by identifying the value of the filtration at which a feature appears (is "born") and disappears ("dies"). This is usually interpreted as a measure of 'importance' or robustness of the feature. Features that have short persistence are considered topological noise, whereas longer persistence is taken to reflect a meaningful topological feature (Cohen-Steiner, Edelsbrunner, & Harer, 2007). We can plot the persistent homology as a diagram showing the birth of each feature against its death. We often plot persistent homology for several different choices of the dimension (e.g., clusters, loops, etc.) in a single figure using different symbols to distinguish the dimensions. Figure 1 shows high persistence of one loop feature.

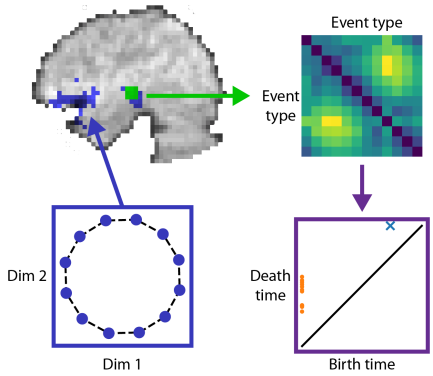

Figure 1. Application of persistent homology analysis to fMRI data. Lower left box shows a topological pattern of activity (a set of representations with distances in a high-dimensional space that constitute a loop — blue dots) inserted into an ROI in the brain, convolved with the hemodynamic response, and combined with noise (blue voxels in upper left image). A distance matrix is then calculated for each 3 x 3 x 3 searchlight of voxels by comparing the event-evoked activity for each event type (matrix at upper right). This distance matrix is used as the input to a persistent homology computation (purple box at lower right) to extract cluster features (orange) and loops (blue). The loop features are used to characterize the signal strength for the voxel at the center of the searchlight.

Our analysis focuses on a signal constituting a single loop feature imposed on the data. We consider two ways of testing for the presence of this signal: i) the number of loops identified in the data (which should be exactly 1); and ii) the persistence of the

largest loop (if any were identified). To do so, we created a voxelwise TDA metric by computing the persistent homology for every searchlight in the brain, using the BrainIAK-extra's (https://github.com/brainiak/brainiak-extras) Python wrapper of the PHAT algorithm (Bauer, Kerber, Reininghaus, & Wagner, 2017). For computational efficiency, we looked only for cluster and loop features. If TDA is an adequate tool for identifying the loop structure embedded in realistically noisy fMRI data, then ROIs in which the signal was embedded should produce persistent homology plots with exactly one persistent (long-lived) feature; this should not be so for control ROIs. For the latter, we used voxels on the symmetrically opposite side of the brain as the signal ROIs, and did not have signal added to the data (also 442 voxels). We conducted analyses both on simulated single brain data, as well as on the average of a full cohort of simulated brains' distance matrices — forming a single "super-brain" — which, by noise averaging, might strengthen the signal.

Note that the procedure described above represents a best case scenario: it assessed whether it was possible to detect known signal at a known location in the brain. If successful, then to demonstrate the usefulness of TDA as an exploratory method, it would be necessary to show that signal can be recovered under more realistic circumstances in which neither the signal nor its location are known and using whole-brain analyses that correct for multiple comparison issues.

Results

We first analyzed the size of the most persistent (or long-lived) loop feature. Figure 2A shows the brain-wise *t* statistic for the comparison of voxels in the ROI that

contained signal with control a ROI (i.e., voxels on the exact opposite side of the brain that did not contain signal). Figure 2B shows the results for the "super-brain" analyses. Each line represents a different set of experiment conditions (5 or 10 repetitions per run; 12, 15, or 18 event types), and each increment represents different signal magnitudes (0, 0.25, 0.5, 0.75, 1, or 5% signal change). To determine whether persistence of the loop feature was significantly greater in the signal vs. control ROIs, we used the critical *t* value for $p<.001$ (df=441). The plots show that for 5 repetitions per run a high (0.75%) signal change is needed to recover the loop structure. For 10 repetitions, only a moderate (0.5%) signal change is needed to recover loop signal. Figure 2B shows the same analysis for the "super-brain" with similar results. Note, the latter appears to be more noisy than the brain-wise analysis. This means it is less noisy to average participants in terms of TDA summary statistics (i.e., participants are more similar to each other) than when averaging the distance matrices used to create these summary statistics. So far, these findings suggest that even a single, simple, clear signal can be recovered by TDA only in designs with a strong evoked response and/or a large amount of data.

Figure 2C-D show the proportion of voxels in the signal ROI that contain only one loop feature (i.e. the associated 1-dimensional persistence diagram has exactly one point). The shaded bars on the right-hand side represent the total range of these values across all noise voxels (different event numbers have different noise properties so these are ranges are different). The results show that the values in the signal voxels fall within the range of noise values (based on the minimum and maximum proportion across all of the different simulations run with this number of nodes), indicating that TDA was no better than chance at detecting the specific topological structure of signal embedded in

the data. Additional analyses, not reported here, showed that when we ignored loops that were not sufficiently persistent (i.e., ignored topological noise), we got similar results. The only condition in which TDA was successful in identifying the signal was when the data is averaged across brains and the signal was the strongest one tested (and in excess of what can be expected from typical event-related fMRI data).

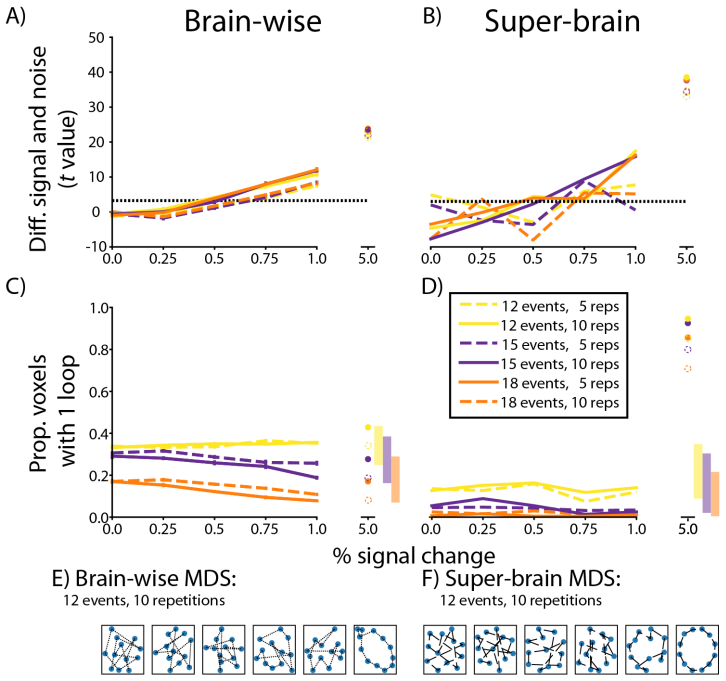

Figure 2. Results of persistent homology analysis of simulated event-related fMRI data. The top row shows the difference in maximum persistence of loop features between signal and matched noise voxels for brain-wise (A) and "super-brain" (B) analyses. Different colors depict numbers of event types (12, 15 or 18) and the different lines depict numbers of trial repetitions per run (5 or 10). The dotted line represents a cutoff for significance at $p<.001$. The middle row shows the proportion of voxels in the signal ROI exhibiting one loop feature (i.e. the 1-dimensional persistence diagram has exactly one point) for the brain-wise (C) and "super-brain" analyses (D). Shaded regions on the far edge represent the range of observed values for matched noise voxels. The bottom row shows multidimensional scaling plots of the brain-wise (E) and "super-brain" (F) data from the signal voxels for a typical condition (12 event types, 10 repetitions) that would take 90 minutes to run. Lines connect events that ought to be adjacent in representation space. The six panels in each group of multidimensional scaling plots correspond to the six levels of percent signal change. Error bars for (A) and (C) are standard error of the mean between participants.

The elements of Figure 2E shows examples of multi-dimensional scaling plots (Shepard, 1962) for all signal voxels in the in a single brain's data and Figure 2F shows the "super-brain" data, providing an estimate of the structure of the signal. These visualizations suggest that the only condition in which the loop structure of the signal was detectable was the 5% signal change condition (using either the brain-wise or "super-brain" analysis).

Discussion

Here we report simulation analyses evaluating the usefulness of TDA methods in recovering topological structure in event-related fMRI data. Specifically, we used fmrisim to simulate a realistic pattern of noise in fMRI data, and then inserted a specific topological feature (a single loop) into some voxels. We calculated the persistent homology of all voxels in the brain and evaluated the extent to which the loop was reliably identified by TDA. We showed that although it was possible for TDA to do so, the noise characteristics of event-related fMRI data made this likely only under conditions of either an extremely strong signal or large amounts of data — both in ranges that exceed what can be expected from current event-related fMRI designs.

These analyses suggest that TDA has limited usefulness for current event-related designs fMRI data. This interpretation is supported by our use of a very simple topological feature — a loop — described by a relatively small number of event types, and analysis methods that provided the best possible circumstances for finding signal. We had initially intended to examine more complicated topological features, such as "figures

of 8”, interlocking rings and branching structures, but initial tests, coupled with the results reported above, suggested that more complex topological structures (which inherently require more event types) are presently beyond the reach of TDA using current event-related fMRI data.

While these results are disappointing, they do not necessarily bear on the usefulness of TDA outside of the domain of event-related designs in fMRI. Empirical research has shown that TDA has value in other domains (Knyazeva et al., 2016; Lord et al., 2016; Petri et al., 2014; Saggar et al., 2018; Stolz, Harrington, & Porter, 2017) and it is possible that other experimental designs might profit from the use of TDA to identify mental representations. For example, designs using more naturalistic stimuli (e.g., movies or narratives; Hasson, Nir, Levy, Fuhrmann, & Malach, 2005; Yeshurun et al., 2017) generate large quantities of data (in which each time point can be treated as a node on a graph) with strong evoked responses. Such data may be more conducive to analysis using TDA.

In sum, we have shown that realistic patterns of noise limit the ability of TDA to recover topological signal from simulations of current forms of event-related fMRI data. TDA is a powerful tool that continues to hold promise for exploratory analysis of other forms of neural data. For now, however, we advise caution in its application to event-related fMRI data, and encourage the pursuit of advances that can overcome the limitations we have reported here.

Acknowledgements

Thank you to N. Malek for comments on an earlier draft.

Contributions

C.T.E., M. L., B. K. & J.D.C. designed the project. C.T.E., M. L., G.H.P., B. K. developed the analysis. C.T.E. performed the analyses. C.T.E. wrote the initial draft of the manuscript. All authors reviewed the manuscript.

Conflicting interests

B.K. is an employee of Intel Corporation which has a current grant with Princeton University. No other conflicts are declared.

Funding

This work was made possible by support from Intel Corporation and the John Templeton Foundation. The opinions expressed in this publication do not necessarily reflect the views of these agencies.

References

Bassett, D. S., & Sporns, O. (2017). Network neuroscience. *Nature Neuroscience, 20*(3), 353-364.

Bauer, U., Kerber, M., Reininghaus, J., & Wagner, H. (2017). Phat–persistent homology algorithms toolbox. *Journal of symbolic computation, 78*, 76-90.

Burock, M. A., Buckner, R. L., Woldorff, M. G., Rosen, B. R., & Dale, A. M. (1998). Randomized event-related experimental designs allow for extremely rapid presentation rates using functional MRI. *Neuroreport, 9*(16), 3735-3739.

Dabaghian, Y., Mémoli, F., Frank, L., & Carlsson, G. (2012). A topological paradigm for hippocampal spatial map formation using persistent homology. *PLoS computational biology, 8*(8), e1002581.

Carlsson, G. (2009). Topology and data. *Bulletin of the American Mathematical Society, 46*(2), 255-308.

Cohen-Steiner, D., Edelsbrunner, H., & Harer, J. (2007). Stability of persistence diagrams. *Discrete & Computational Geometry*, *37*(1), 103-120.

Ellis, C. T., Baldassano, C., Schapiro, A. C., Cai M. B., & Cohen J. D. (2018). *Facilitating open-science with realistic fMRI simulation: validation and application.* Manuscript in preparation.

Friston, K. J., Fletcher, P., Josephs, O., Holmes, A., Rugg, M., & Turner, R. (1998). Event-related fMRI: characterizing differential responses. *Neuroimage, 7*(1), 30-40.

Ghrist, R. (2008). Barcodes: the persistent topology of data. *Bulletin of the American Mathematical Society, 45*(1), 61-75.

Giusti, C., Pastalkova, E., Curto, C., & Itskov, V. (2015). Clique topology reveals intrinsic geometric structure in neural correlations. *Proceedings of the National Academy of Sciences, 112*(44), 13455-13460.

Gläscher, J. (2009). Visualization of group inference data in functional neuroimaging. *Neuroinformatics, 7*(1), 73-82.

Hasson, U., Nir, Y., Levy, I., Fuhrmann, G., & Malach, R. (2004). Intersubject synchronization of cortical activity during natural vision. *Science, 303*(5664), 1634-1640.

Knyazeva, I., Orlov, V., Ushakov, V., Makarenko, N., & Velichkovsky, B. (2016). On alternative instruments for the fMRI data analysis: General linear model versus

algebraic topology approach. In *Biologically Inspired Cognitive Architectures (BICA) for Young Scientists* (pp. 107-113). Springer, Cham.

Lord, L. D., Expert, P., Fernandes, H. M., Petri, G., Van Hartevelt, T. J., Vaccarino, F., ... & Kringelbach, M. L. (2016). Insights into brain architectures from the homological scaffolds of functional connectivity networks. *Frontiers in Systems Neuroscience*, *10*(85), 1-14.

Nau, M., Schröder, T. N., Bellmund, J. L., & Doeller, C. F. (2018). Hexadirectional coding of visual space in human entorhinal cortex. *Nature neuroscience, 21*(2), 188.

Nielson, D. M., Smith, T. A., Sreekumar, V., Dennis, S., & Sederberg, P. B. (2015). Human hippocampus represents space and time during retrieval of real-world memories. *Proceedings of the National Academy of Sciences, 112*(35), 11078-11083.

Petri, G., Expert, P., Turkheimer, F., Carhart-Harris, R., Nutt, D., Hellyer, P. J., & Vaccarino, F. (2014). Homological scaffolds of brain functional networks. *Journal of The Royal Society Interface*, *11*(101), 20140873.

Saggar, M., Sporns, O., Gonzalez-Castillo, J., Bandettini, P. A., Carlsson, G., Glover, G., & Reiss, A. L. (2018). Towards a new approach to reveal dynamical organization of the brain using topological data analysis. *Nature Communications, 9*(1), 1399.

Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B., & Botvinick, M. M. (2013). Neural representations of events arise from temporal community structure. *Nature Neuroscience, 16*(4), 486-492.

Shepard, R. N. (1962). The analysis of proximities: multidimensional scaling with an unknown distance function. I. *Psychometrika, 27*(2), 125-140.

Stolz, B. J., Harrington, H. A., & Porter, M. A. (2017). Persistent homology of time-dependent functional networks constructed from coupled time series. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, *27*(4), 047410.

Valentine, T. (1991). A unified account of the effects of distinctiveness, inversion, and race in face recognition. *The Quarterly Journal of Experimental Psychology Section A, 43*(2), 161-204.

Yeshurun, Y., Swanson, S., Simony, E., Chen, J., Lazaridi, C., Honey, C. J., & Hasson, U. (2017). Same story, different story: the neural representation of interpretive frameworks. *Psychological Science, 28*(3), 307-319.

Zomorodian, A., & Carlsson, G. (2005). Computing persistent homology. *Discrete & Computational Geometry, 33*(2), 249-274.

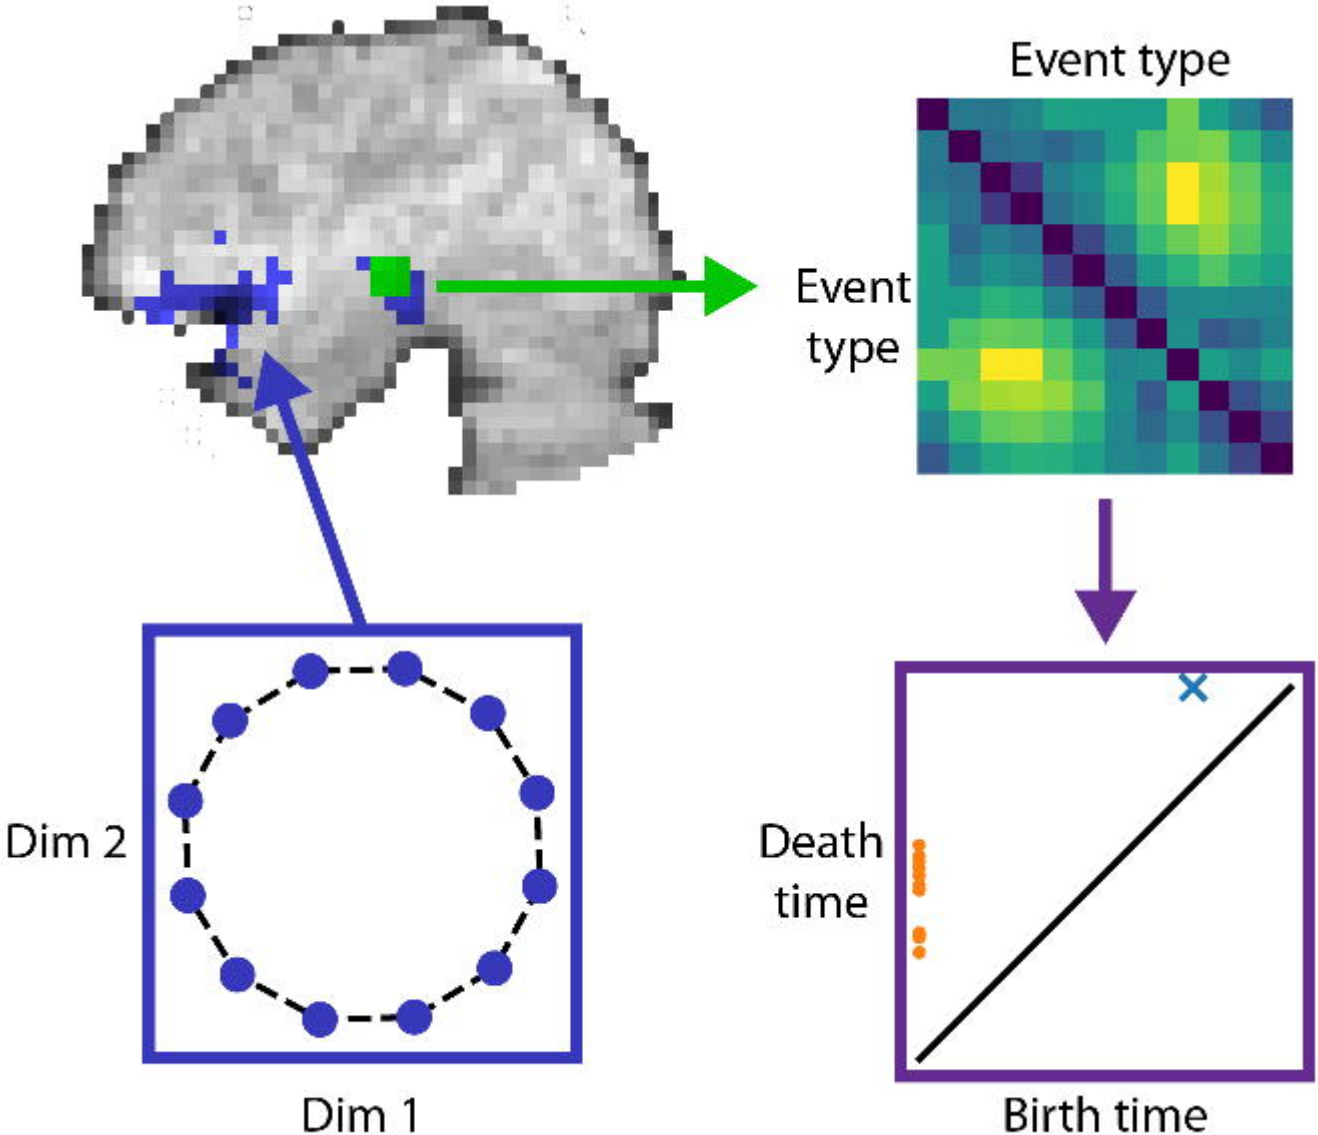
Event type
Event
type
Dim 2
Dim 1
Death
time
Birth time

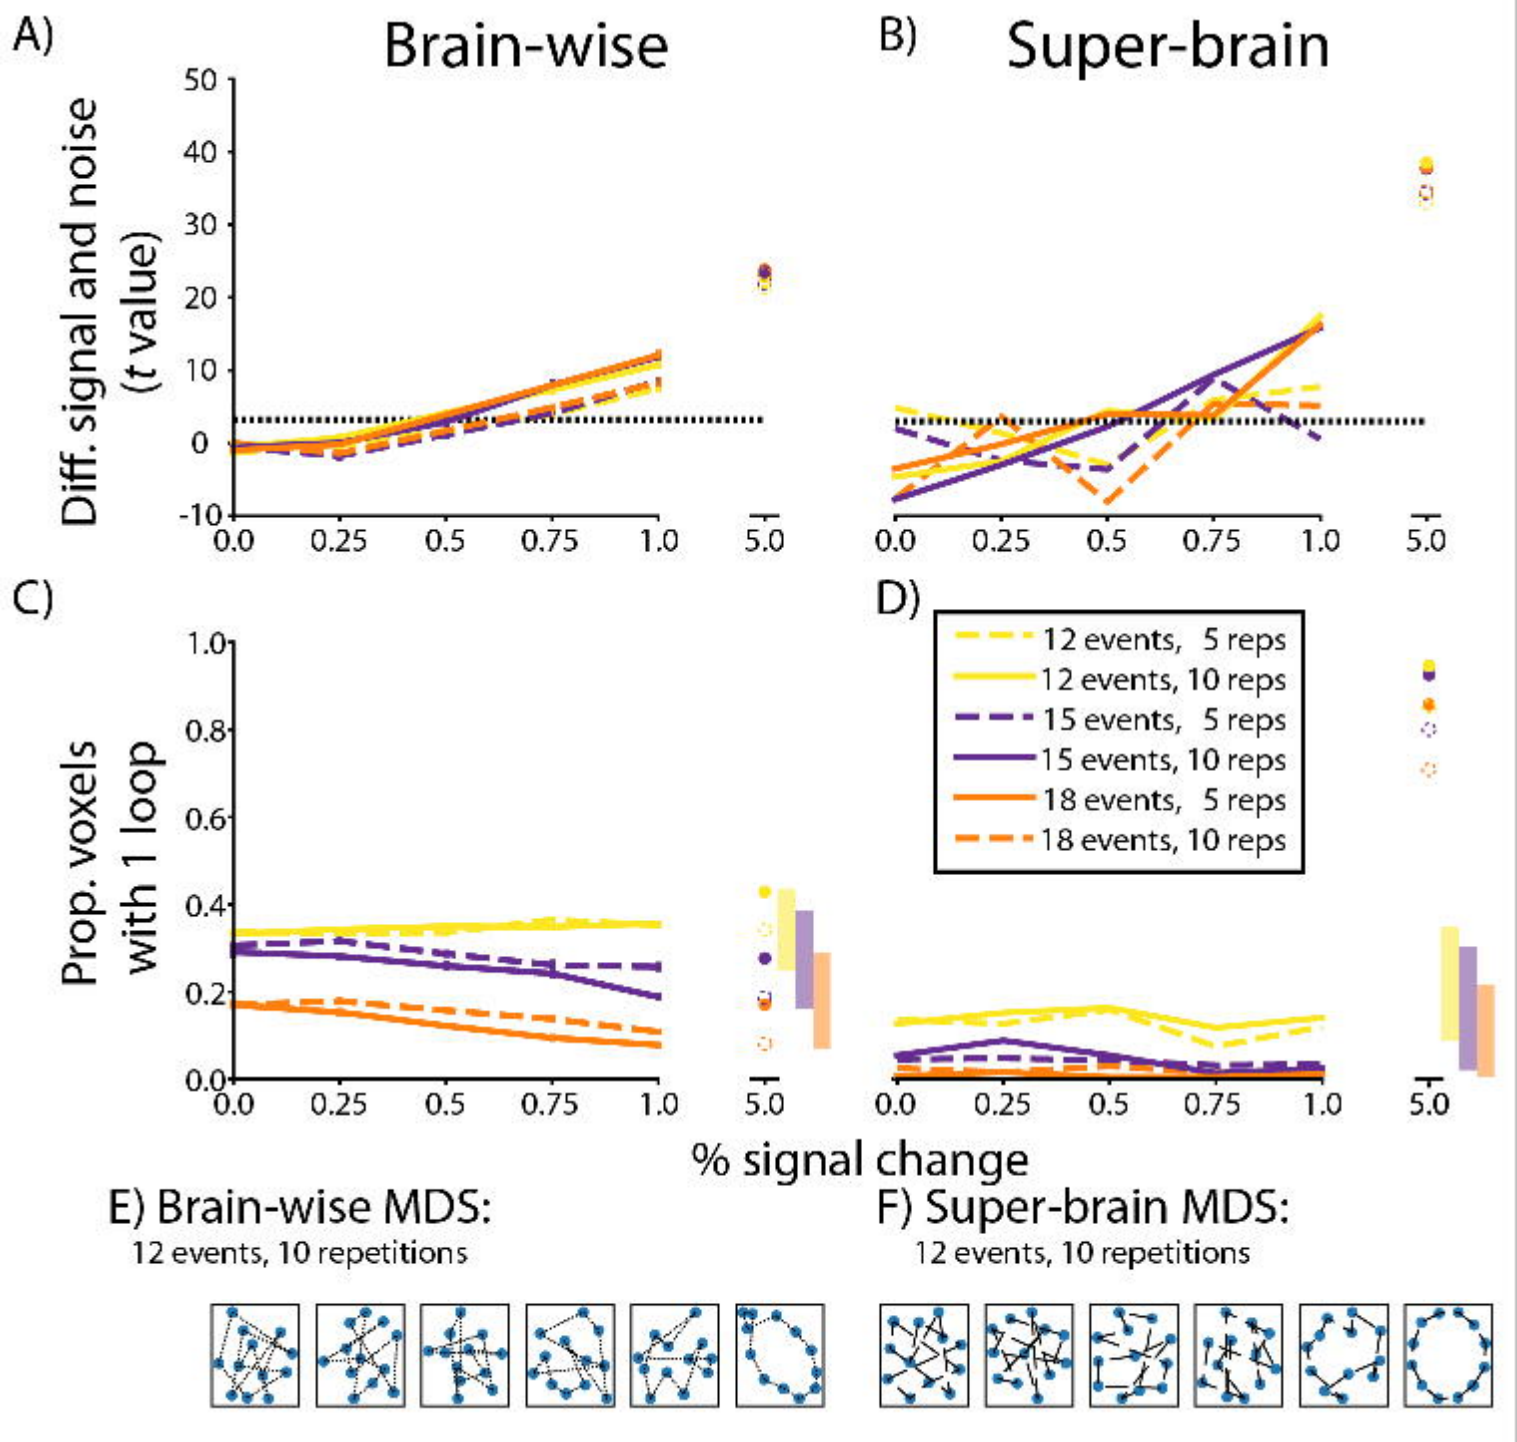
A)
Brain-wise
B)
Super-brain
Diff. signal and noise
(t value)
50
40
30
20
10
0
-10
0.0
0.25
0.5
0.75
1.0
5.0
C)
D)
Prop. voxels
with 1 loop
1.0
0.8
0.6
0.4
0.2
0.0
12 events, 5 reps
12 events, 10 reps
15 events, 5 reps
15 events, 10 reps
18 events, 5 reps
18 events, 10 reps
% signal change
E) Brain-wise MDS:
12 events, 10 repetitions
F) Super-brain MDS:
12 events, 10 repetitions