Spectral Permutation Test on Persistence Diagrams

Yuan Wang[1], *Moo K. Chung*[2], *Julius Fridriksson*[3]
[1]Department of Epidemiology and Biostatistics, University of South Carolina.
[2]Department of Biostatistics and Medical Informatics, University of Wisconsin - Madison
[3]Department of Communication Sciences and Disorders, University of South Carolina.
Correspondence: wang578@mailbox.sc.edu

Paper ID 5167

Brain networks constructed from diffusion and functional magnetic resonance imaging (dMRI and fMRI) are typically investigated through graph theoretic models. It has recently been noted that the complexity of brain connectivity may not be sufficiently captured by single-scale models and multi-scale models are needed. Persistent homology (PH) is an algorithm that extracts multi-scale features in brain networks that cannot be easily decoded by standard graph theoretic models. It summarizes topological structures in a network through multi-scale descriptors such as persistence diagram (PD). Various statistical inference procedures have been developed for PDs. In this study, we propose a spectral permutation test on PDs through a new scale-space representation, where the upper-triangular domain of PDs is represented using a finite number of Fourier coefficients with respect to the Laplace-Beltrami (LB) eigenfunction expansion of the domain. The scale-space representation provides a powerful vectorized algebraic representation for comparisons of PDs at the same coordinates, foregoing the need for matching across PDs due to their arbitrary point locations. We evaluate the empirical performance of the proposed spectral permutation test in detecting an innate shape with a hole in a two-dimensional image. The test is found to be sensitive in detecting the topological structure under noisy perturbations. It is also applied to compare diffusion and rest-state functional brain networks at baseline and first treatment visits within two types of post-stroke aphasia. We find that the structural connectivity in the diffusion networks alters between visits, whereas the resting-state functional connectivity does not.

Methods

Suppose we have a network represented by the weighted graph $G=(V,w)$ with the node set $V=\{1,\dots,p\}$ and unique positive undirected edge weights $w=(w_{ij})$ constructed from a similarity measure such as Pearson's correlation. We define the binary network $G_\varepsilon=(V,w_\varepsilon)$ as a subgraph of G consisting of the node set V and the binary edge weights w_ε defined by

$$w_{\varepsilon,ij}=\begin{cases}1 & \text{if } w_{ij}<\varepsilon;\\ 0 & \text{otherwise.}\end{cases} \tag{1}$$

As we increase ε, which we call the *filtration value*, more edges are included in the binary network G_ε and so the size of the edge set increases. Since edges connected in the network do not get disconnected again, we observe a sequence of nested subgraphs

$$G_{\varepsilon_0}\subset G_{\varepsilon_1}\subset G_{\varepsilon_2}\subset\cdots, \tag{2}$$

for any $\varepsilon_0\le\varepsilon_1\le\varepsilon_2\le\cdots$. This sequence of nested subgraphs make up a *Rips filtration* where two nodes with a weight w_{ij} smaller than ε are connected, and the birth and death of clusters and holes are tracked through the filtration. We pair the birth and death times of clusters and holes as coordinates of scatter points on a planar graph $\{(a_i,b_i)\}_{i=1}^L$, i.e., *persistence diagram* (PD). PDs do not possess a natural statistical framework and requires additional manipulation such as kernel smoothing.

Heat-kernel estimation of persistence diagram. We estimate a PD based on a spectral representation. Let $\mathscr{T}$ be the upper triangular region above $y=x$ where the scatter points $\{(a_i,b_i)\}_{i=1}^L$ are located. We constrain $\mathscr{T}$ at some fixed y-coordinates so that $\mathscr{T}$ is bounded. The heat kernel (HK) in $\mathscr{T}$ is given by

$$K_\sigma(p,q)=\sum_{k=0}^{\infty}e^{-\lambda_k\sigma}\psi_k(p)\psi_k(q) \tag{3}$$

with respect to the eigenfunctions ψ_k of Laplace-Beltrami (LB) operator Δ satisfying $\Delta\psi_k(p)=\lambda_k\psi_k(p)$ for $p\in\mathscr{T}$. The first eigenvalue $\lambda_0=0$ corresponds to eigenfunction $\psi_0=\frac{1}{\sqrt{\mu(\mathscr{T})}}$, where $\mu(\mathscr{T})$ is the area of triangle $\mathscr{T}$ and σ is the bandwidth of the HK. Consider heat diffusion

$$\frac{\partial\mathbf{h}(\sigma,p)}{\partial\sigma}=\Delta\mathbf{h}(\sigma,p) \tag{4}$$

with the initial condition $\mathbf{h}(\sigma=0,p)=\sum_{i=1}^L\delta_{(a_i,b_i)}(p)$, where $\delta_{(a_i,b_i)}$ is the Dirac-delta function at (a_i,b_i). The scatter points in the PD serve as the heat sources. A unique solution to (4) is thus given by the HK expansion

$$\begin{aligned}\mathbf{h}(\sigma,p)&=\int_{\mathscr{T}}K_\sigma(p,q)\mathbf{h}(\sigma=0,q)\,d\mu(q)\\&=\sum_{k=0}^{\infty}e^{-\lambda_k\sigma}f_k\psi_k(p),\end{aligned} \tag{5}$$

where

$$f_k=\int_{\mathscr{T}}\mathbf{h}(\sigma=0,q)\psi_k(q)\,d\mu(q)=\sum_{i=1}^L\psi_k(a_i,b_i) \tag{6}$$

are the Fourier coefficients with respect to the the LB eigenfunctions. In practice, we include a finite number of terms for PD estimation:

$$\mathbf{h}_K(\sigma,p)=\sum_{k=0}^{K}e^{-\lambda_k\sigma}f_k\psi_k(p), \tag{7}$$

with sufficiently large degree $K=10000$ for convergence. As $\sigma\to0$, we can completely recover the initial scatter points. As $\sigma\to\infty$, we are doing kernel density estimation with uniform kernel on $\mathscr{T}$.

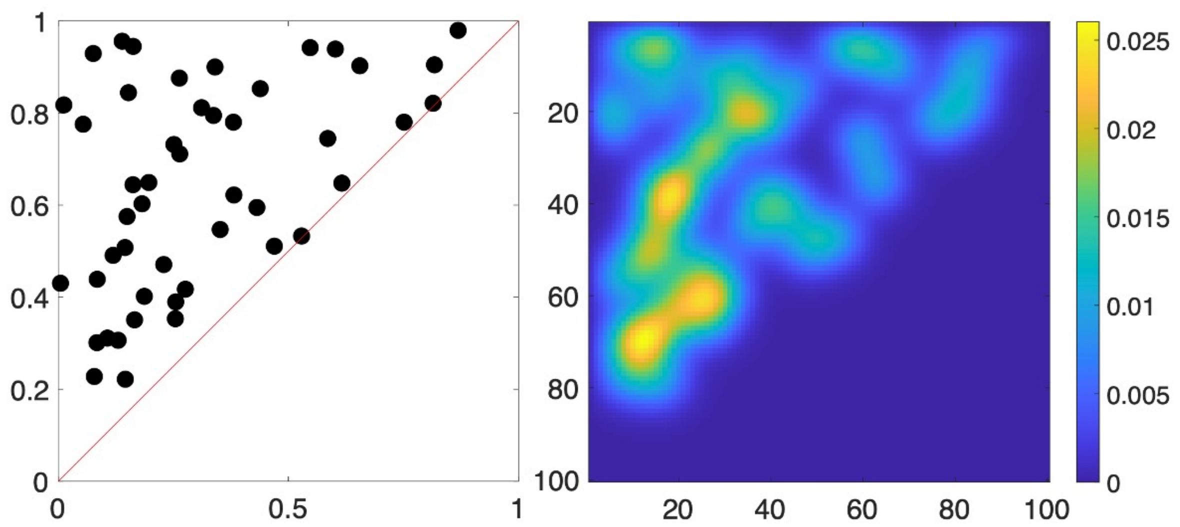

Figure 1. Heat-kernel (HK) estimation of a persistence diagram (PD) with $\sigma=10$: PD (left) and its HK estimate (right).

Spectral permutation test on persistence diagrams. We use permutation test to compare across PDs. We propose a *spectral transposition test* that performs the permutation test on the spectrum of PDs. Suppose we obtain PDs through filtrations on two groups of correlation brain networks with rsizes m and n. The degree-K HK-estimates of the PDs $\{\mathbf{f}^i\}$ and $\{\mathbf{g}^j\}$ are

$$\mathbf{f}^i(p)=\sum_{k=0}^{K}e^{-\lambda_k\sigma}f_k^i\psi_k(p), i=1,\dots,m, \tag{8}$$

$$\mathbf{g}^j(p)=\sum_{k=0}^{K}e^{-\lambda_k\sigma}g_k^j\psi_k(p), j=1,\dots,n, \tag{9}$$

where f_k^i and g_k^j, $k=0,\dots,K$, are the Fourier coefficients with respect to the k-th LB eigenfunction ψ_k. Their functional means are

$$\bar{\mathbf{f}}(p)=\sum_{k=0}^{K}e^{-\lambda_k\sigma}\bar{f}_k\psi_k(p), \tag{10}$$

$$\bar{\mathbf{g}}(p)=\sum_{k=0}^{K}e^{-\lambda_k\sigma}\bar{g}_k\psi_k(p), \tag{11}$$

where $\bar{f}_k=\frac{1}{m}\sum_{i=1}^m f_k^i$ and $\bar{g}_k=\frac{1}{n}\sum_{j=1}^n g_k^j$ are the mean Fourier coefficients. We will use the L_2-norm difference between the functional means $\|\bar{\mathbf{f}}-\bar{\mathbf{g}}\|_2^2$ as a test statistic for measuring the group differences. We can algebraically show

$$\|\bar{\mathbf{f}}-\bar{\mathbf{g}}\|_2^2=\sum_{k=0}^{K}e^{-2\lambda_k\sigma}(\bar{f}_k-\bar{g}_k)^2. \tag{12}$$

In a standard permutation test, the subject labels of the two groups are randomly exchanged. Here, we consider the permutation π_{ij} that only exchanges the i-th and j-th subject labels between $\{\mathbf{f}^i, i=1,\dots,m\}$ and $\{\mathbf{g}^j, j=1,\dots,n\}$ and keeps all the other PDs fixed, i.e.

$$\pi_{ij}(\mathbf{f}^1,\dots,\mathbf{f}^m)=(\mathbf{f}^1,\dots,\mathbf{g}^j,\dots,\mathbf{f}^m), \tag{13}$$

$$\pi_{ij}(\mathbf{g}^1,\dots,\mathbf{g}^n)=(\mathbf{g}^1,\dots,\mathbf{f}^i,\dots,\mathbf{g}^n), \tag{14}$$

which we call a *spectral transposition*. Any permutation of the two groups of m and n subjects is reachable by a sequence of transpositions, which has been shown to be computationally much more efficient than the standard permutation testing procedure of exchanging all labels at once. We generate the empirical distribution for the permutation test through the spetral transpositions. Over one spectral transposition π_{ij}, we obtain the L_2 distance between the functional means of the degree-K HK-estimates based on transposed PDs:

$$\|\bar{\mathbf{f}}'-\bar{\mathbf{g}}'\|_2^2=\sum_{k=0}^{K}e^{-2\lambda_k\sigma}(\bar{f}_k'-\bar{g}_k')^2, \tag{15}$$

where $\bar{f}_k'=\bar{f}_k+\frac{1}{m}(g_k^j-f_k^i)$ and $\bar{g}_k'=\bar{g}_k+\frac{1}{n}(f_k^i-g_k^j)$ are the means of transposed Fourier coefficients. Since we know $\bar{f}_k$ and $\bar{g}_k$ already, we simply update the terms $\frac{1}{m}(g_k^j-f_k^i)$ and $\frac{1}{n}(f_k^i-g_k^j)$ in an *online* fashion. The p-value of the spectral permutation test is then calculated as the proportion of L_2 distances in the empirical distribution exceeding the L_2 distance between the observed PDs. To ensure convergence, we perform 100,000 permutations in the subsequent analysis.

Performance evaluation. We evaluate the proposed test's power in detecting the shape of a key, or part of the key, with a distinct hole. In each simulation, two groups of five 100-point point clouds are generated: the 100 points in each point cloud of the first group are generated randomly from the rectangular image or part of it above the threshold, whereas the 100 points in each point cloud of the second group are generated randomly with a varied percentage (90%, 95%, 100%) of points from the shape of the key. Rips filtration is constructed on each point cloud. The proposed spectral permutation test is then applied to compare the PDs of the Rips filtrations in the two groups. When there are respectively 90%, 95%, and 100% points sampled from the shape of the key in the second group, the spectral permutation test rejects (p-value <0.05) the null hypothesis of no group difference in 91, 100, and 100% (whole key) and 76, 88, and 93% (partial key) of 100 simulations (corresponding means $\pm$ standard deviations of p-values: 0.0124 ± 0.0327, 0.0041 ± 0.0125, 0.0008 ± 0.0057 (whole key), and 0.0417 ± 0.0794, 0.0200 ± 0.0545, 0.0082 ± 0.0217 (partial key), showing that the test stays sensitive in detecting the group shape difference when points in the second group are not entirely sampled from the shape of the key.

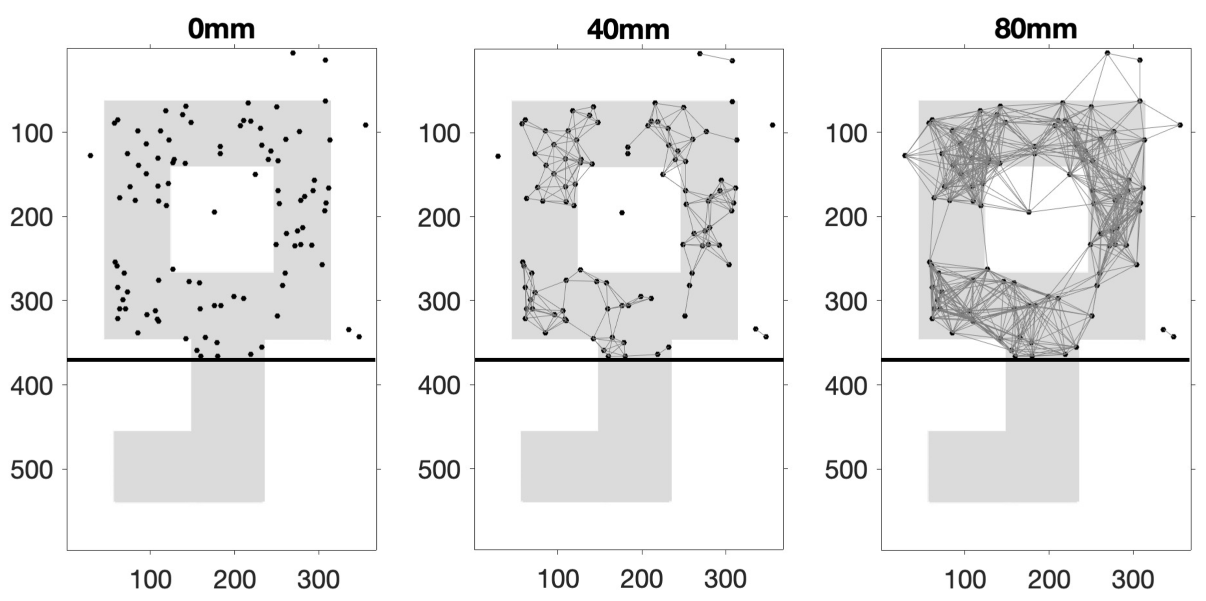

Figure 2. The 1-skeletons of Rips complexes built on 100 points randomly sampled from the image with an innate shape of a key.

Application

Aphasia is an acquired speech-language disorder that commonly develops after a left-hemisphere stroke. It affects an estimated one million people in the US. Quantification of brain functional patterns in fMRI allows for an objective assessment of aphasia impairment.

Data. Participants were recruited locally in Columbia, South Carolina, as part of the Predicting Outcome of Language Recovery (POLAR) in Aphasia study of post-stroke aphasia by the Center for the Study of Aphasia Recovery at the University of South Carolina. The study was approved by the Institutional Review Board and adhered to the ethics guidelines. Only participants with a single ischemic or a hemorrhagic stroke in the left hemisphere were included. Aphasia types were classified based on the Western Aphasia Battery-Revised (WAB-R). Among the participants included in the study, 14 were diagnosed with anomia or anomic aphasia (a mild, fluent type of aphasia where individuals have word retrieval failures and cannot express the words they want to say, particularly nouns and verbs), and 28 were diagnosed with Broca's aphasia (type of aphasia characterized by partial loss of the ability to produce spoken or written language, although comprehension generally remains intact). Every participant underwent resting-state fMRI (rs-fMRI) and diffusion MRI (dMRI) scans at a baseline and first treatment visit with a Siemens Prisma 3T scanner with a 20-channel head coil. The preprocessing procedures of the fMRI data include motion correction, brain extraction and time correction. This modality is processed using a novel method developed for stroke patients. The automated anatomical label (AAL) atlas was used for brain parcellation. A single correlation matrix representing functional connectivity between 90 AAL ROIs (excluding cerebellum and vermis) was computed for each individual. Average fractional anisotropy (FA) values were computed for AAL ROIs for each participant. A structural correlation matrix was computed on the average FA values by leaving one participant out in each group.

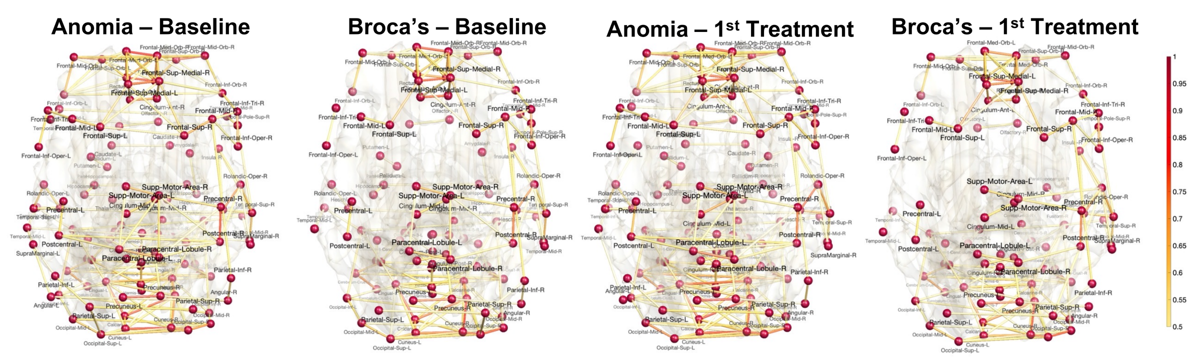

Figure 3. Average resting-state functional correlation matrices of the anomia and Broca's groups in two visits.

Topological network analysis. We construct Rips filtrations and PDs over the individual structural and resting-state functional correlation networks within the anomic and Broca's groups. The HK-estimated PDs are then respectively compared between the two visits using the spectral transposition test. The test on resting-state functional networks does not detect strong difference in holes between the two visits in both anomia (p-value = 0.7048) and Broca's aphasia (p-value = 0.3641), which could indicate connectivity in resting-state functional network is not yet altered by the first treatment. On the other hand, there is significant difference between the structural networks in both anomia (p-value = 0.0151) and Broca's aphasia (p-value = 0.0221), indicating changes in structural connectivity between the two visits.

Acknowledgment

NIH R01-EB028753 and NSF DMS-2010778 (PI: Chung), NIH R21-DC014170 and NIH P50-DC014664 (PI: Fridriksson)