TOPOLOGICAL CORRECTION OF SUBJECT-LEVEL INTRINSIC CONNECTIVITY NETWORKS

Noah Lewis*†*, Armin Iraji*†*, Robyn Miller*†*, Vince D. Calhoun*†,‡

* School of Computational Science and Engineering, Georgia Institute of Technology
† Tri-institutional Center for Translational Research in Neuroimaging
Data Science (TReNDS), Georgia Institute of Technology, Georgia State University,
Emory University
‡School of Electrical and Computer Engineering, Georgia Institute of Technology

ABSTRACT

Over the last several decades, researchers have sought to capture the underlying functional activity of the human brain from functional magnetic resonance imaging (fMRI). One well-studied and promising avenue of research is independent component analysis (ICA) to capture the maximally independent set of elements known as functional networks. These functional networks are represented as spatial maps or 3D images in which each voxel has a value associated with the given network. The current state-of-the-art methods use group-level spatial maps as a reference to provide estimates of subject-level maps, which are vulnerable to the low contrast-to-noise (CNR) ratio of fMRI signals. However, such approaches do not account for all group-level spatial information. As such, their subject-level estimate is still quite noisy. This work presents a novel method that leverages the topological properties of the group maps to improve subject-level estimations. We show that adding topological similarity constraints also improves subject-specific information.

***Index Terms*—** Brain Networks, fMRI, Topology, ICA, Topologically-informed Network Estimation

1. INTRODUCTION

Group independent component analysis (ICA) [1] is a widely used method for estimating brain networks, referred to as intrinsic connectivity networks (ICNs). These ICNs, estimated at the group level, contain valuable information about brain functionality universal to all humans. To capture individual networks, most methods leverage post-hoc estimations using group-level maps and the subject's whole brain fMRI image [2]. One approach for estimating single subject maps, dual regression [3], is a linear transform that estimates both the subject-level spatial map and component time series. However, linear subject-level estimates do not consider spatial information, as each voxel acts as a single regressor. This loss of spatial information impacts both the visualization and analyses of these subject maps. Thus, a gold standard would be a method that adapts the current methodology to regard spatial information without interfering with the subject information. This is where we suggest that an effective method would correct the linear regression while constraining the maps to maintain group-level information that would reduce noise and make the subject-level maps more quantifiably similar to the group-level maps.

One aspect of the human brain and neuroimaging is that these images often contain informative topological properties. Previous work has studied the topological information of networks and graphs of fMRI images [4] and the topological correction of images [5]. This topological information is generally estimated by the persistent homology of the image or graph. Persistence estimates the homological classes (objects, holes, hollows, etc.) after defining the image or graph as a set of simplicial complexes or a set of points, lines, triangles, and the n-dimensional counterparts. To date, no work has leveraged persistence to capture or improve subject-level spatial maps of ICNs.

This work presents a novel method that adapts and amalgamates previous persistent homology methodologies to constrain the subject-level spatial maps to be topologically similar to the group-level images. The assumption is that this will reduce the overall CNR by considering the group-level spatial information. Firstly, this work adapts previous work [6] that regularizes machine learning algorithms (linear regression, specifically) to constrain the regression coefficients to have specific topological properties. However, as we do not want to enforce prior knowledge-based topological information, we adapt this topology layer to use a continuous loss function between the group and subject maps. Based on previous work [7, 8], we chose the Wasserstein distance between the persistence information of the group and subject maps. Finally, as a proof-of-concept, we present the preliminary results from our method obtained from 2D coronal slices from the MRI images. To show the effectiveness of this method, we compare the difference between the topologically corrected subject spatial map and the group maps with the difference between the dual-regression-only maps and the group maps. Although it is beneficial for the subject-level maps to be similar to the group-level maps, we must also ensure that we do not lose between-subject variability. This kind of variability is vital to researchers, and we show that our method enhances subject-level variability.

2. METHODS

2.1. Data and Preprocessing

For this work, we chose a dataset of 50 fMRI subjects from the Function Biomedical Informatics Research Network (FBIRN). The images were preprocessed with the Statistical Parametric and Mapping (SPM) (https://www.fil.ion.ucl.ac.uk/spm/) and Analysis of Functional NeuroImages (AFNI) (https://afni.nimh.nih.gov/) toolboxes. Motion correction was performed with the INRIAlign toolbox. A slice-timing correction was performed using the exact middle slice as the reference. Despiking, warping to the Montreal Neurological Institute (MNI) template, resampling to 3mm^3 voxels, spatial smoothing with a 6mm full-width/half-max gaussian kernel, and z-scoring were all performed.

Once processed, we computed the ICNs using the Group ICA of fMRI Toolbox (GIFT) [9]. This pipeline starts with the data's subject-level principal components analysis (PCA) to select the principal components (PCs) with a variance greater than 99%. Then, group-level PCA was applied to the subject-level PCs concatenated across time to select the 20 group-level PCs with the highest variance. The infomax ICA algorithm [10] was applied using the ICASSO framework for 100 runs to obtain the 20 most stable components. We select the components with the largest overlap with the gray matter and the components where low-frequency fluctuations dominate the timecourses. Finally, we select two components for our experiments: one from the default mode network (DMN) and one from the dorsal attention network (ATN).

2.2. ICNs and Dual Regression

Group spatial maps of ICNs have been vital for neuroimaging research. However, group maps are not useful for studying subject variation. This is why subject estimations have become popular. These methods include directly estimating the subject maps [11] and estimating the subject maps from the group components [1]. These include PCA-based and regression-based approaches, called back-reconstruction. One approach, dual regression, uses linear regression to estimate the subject ICN timeseries and a secondary regression that estimates the subject spatial maps from the computed timeseries. However, this does not maintain the group's spatial information.

2.3. Persistent Homology

Persistent homology, a common method to estimate the topological information from images, point clouds, or graphs, has found great utility in neuroimaging [12]. Persistent homology computes the homological features at different resolutions, a process known as filtration. These features are elements within a given homological class. In this work, we use two classes, H0 and H1. H0 is the 0-dimensional connected components, H1 is the set of holes in the image. The features are estimated by breaking the image down into simplicial complexes. Each simplicial complex is a set of simplices: points, lines, and triangles. The filtration, specifically level-set filtration, is the set of linearly-spaced resolutions in which the simplices are computed based on a threshold of the voxel intensities. During the persistence computation, we keep track of each H0 and H1 element as they are created or destroyed during the filtration (i.e., a feature exists if it is within the threshold and is destroyed if it is outside the threshold). The creation and destruction of these features are called "birth" and "death". These birth and death values define the overall topological information.

2.4. Topology Layer for Machine Learning

The topology layer for machine learning, from [6], defines a loss function on a given machine learning algorithm learned via gradient descent. One of their primary examples uses this loss function to regularize linear regression. Their method begins by estimating the regression coefficients with ordinary least squares, then iteratively learning the topological loss to produce a final set of coefficients constrained to contain pre-defined topological information.

2.5. Topological Loss

To properly constrain the spatial maps to include the topological information from the group maps, a loss function must be defined to compare the subject and group maps. As shown in previous work, the Wasserstein distance [7] can be adapted to compute the error between two distributions of homological persistence (i.e., the birth and deaths of all homological features)[4].

$$\mathcal{L}(\Theta) = \min_{\Theta} \sum_{i}^{V} [b_i(\Theta) - b_i(group)]^2 + [d_i(\Theta) - d_i(group)]^2$$

The minimization problem for the Wasserstein-based loss function. b_i and d_i are the births and deaths of the i^{th} object, respectively. Θ is the estimated subject spatial maps from the $group$ spatial maps. Note that the objects are sorted by total persistence value ($d - b$), allowing us to match objects between the two maps. And for every object that exists in one spatial map but not the other (i.e. $|\Theta| \neq |group|$), the object is mapped to zero, thus penalizing erroneous objects in the estimated subject maps. This mapping of erroneous objects to zero is part of this work's novelty and a vital part of the estimation.

2.6. Our Methodology

Although the topology layer fits the problem of dual regression nicely, we do not know the "ground truth" or pre-defined topological properties of a group spatial map. Thus, we adapt the Wasserstein distance to replace the loss function defined by the machine learning topology layer. This adaptation can be seen in the previous section. As this loss function is differentiable, it is trivial to see how it can easily be adapted to the topological layer.

Previous work adapts this loss function to neuroimaging data, but most of it has focused on networks [4] or brain segmentation [8]. To our knowledge, this is the first time someone has adapted a dimension-specific Wasserstein loss function to model parameters. This is the first time someone has used topological spatial constraints for subject spatial maps.

To estimate each coronal slice of the subject spatial maps, we initialize the spatial maps with dual regression. However, this is only an initial position. From here, using gradient descent, we optimize the topological loss function as well as the MSE between the estimated spatial maps and the fMRI images (in order to preserve the linear regression loss). Both the topological loss and the MSE have their own learning rates which are user-specified hyperparameters.

2.7. Experiments

In order to show the validity of our method, we suggest that there is one over-arching goal: to show that the topologically corrected maps are more similar to the group maps while increasing between-subject variability. Subject maps that are more similar to the group maps will intrinsically contain more group spatial information and hopefully increase the CNR. However, we also show that the subject maps now have higher subject variability after the topological correction. We argue that with these two properties, the maps have a higher CNR while also being more visually appealing.

3. RESULTS

3.1. Similarity to Group Maps

We argue that one aspect of cleaner-looking subject maps with higher CNR is that they should be more similar to the group map. To show this, for both the ATN and DMN network, we plotted the Pearson correlation between each subject map and the group map in figure 1. We see that the average correlation over all subjects is for the topological maps than the OLS maps for both the ATN and DMN.

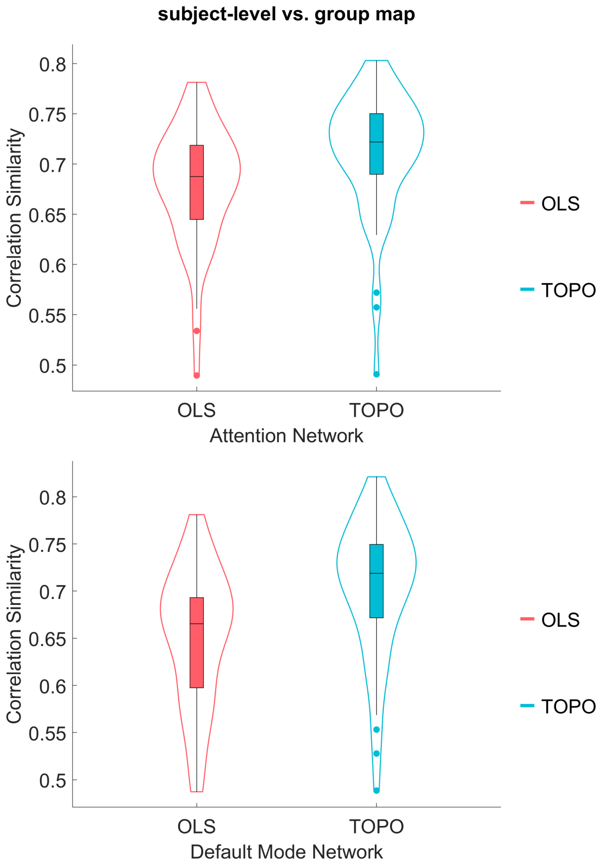

Fig. 1. A violin plot of the correlation between each topologically corrected subject map and the group map (blue) as well as the correlations for the dual-regression only maps (red) for the ATN (top) and DMN (bottom).

3.2. Between-Subject Variability

The second half of our preliminary result shows that topologically correcting the subject maps improves between-subject variability, as seen in figure 2. From this figure, we see a large increase in subject variability. While subject variability does increase with the topological correction, it also increases in the voxels that most contribute to the group network.

4. DISCUSSION

Estimations of subject-specific spatial maps can often be noisy or dismissive of group-level information. So, we present this method that adapts current subject-level estimations to maintain spatial information from the group maps. By constraining the subject maps to be topologically similar to the group maps, we keep relevant spatial information without destroying subject-specific variability. From figure 2, we see that the subject variability increases overall, but more

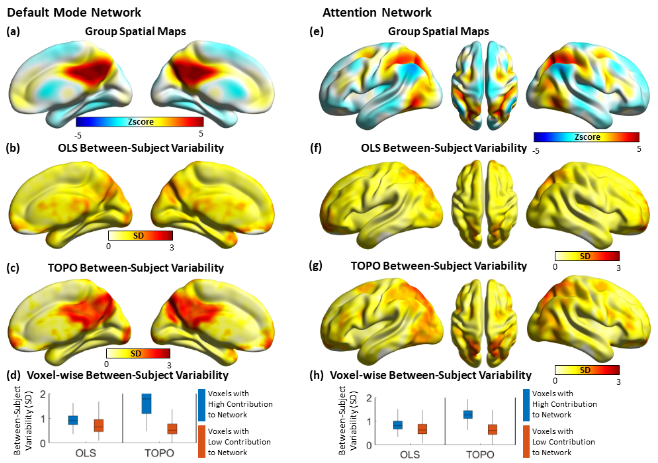

Fig. 2. The average group spatial map (a), the voxel-wise, between-subject standard deviation of the OLS-only maps (b), the voxel-wise, between-subject standard deviation for the topologically corrected maps (c), and a plot of the between-subject variability of both methods (d) for the DMN. As well as the same configuration for the ATN (e-h).

importantly, it increases in voxels that contribute the most to the group network. In 2.b (DMN) and 2.f (ATN), we see that the OLS-only maps show subject variability across the entire brain, whereas the topologically corrected maps in 2.c and 2.g show much higher variability, mostly in voxels that contribute the most to the group network. This is backed by 2.d and 2.h, which plot the subject variability for the network-relevant voxels (blue) and the less relevant voxels (red). This increase in variability in voxels that contribute to the network is a promising sign, as it suggests that the topological correction enhances network-specific signals within the subjects, possibly increasing the CNR. Future work will extend this method from 2D slices to 3D volumes. It should be noted that this extension is not trivial, and requires much more computational overhead. However, working in the native 3D setting of the data will preserve the natural properties. Whereas the 2D setting induces biases associated with the orientation of the given slice.

5. REFERENCES

[1] Vince D Calhoun, Tulay Adali, Godfrey D Pearlson, and James J Pekar, "A method for making group inferences from functional mri data using independent component analysis," *Human brain mapping*, vol. 14, no. 3, pp. 140–151, 2001.

[2] Erik Barry Erhardt, Srinivas Rachakonda, Edward J Bedrick, Elena A Allen, Tülay Adali, and Vince D Calhoun, "Comparison of multi-subject ica methods for analysis of fmri data," *Human brain mapping*, vol. 32, no. 12, pp. 2075–2095, 2011.

[3] Christian Beckmann, Clare E Mackay, Nicola Filippini, Stephen M Smith, et al., "Group comparison of resting-state fmri data using multi-subject ica and dual regression," *Neuroimage*, vol. 47, no. Suppl 1, pp. S148, 2009.

[4] Tananun Songdechakraiwut and Moo K Chung, "Topological learning for brain networks," *arXiv preprint arXiv:2012.00675*, 2020.

[5] Bruce Fischl, Arthur Liu, and Anders M Dale, "Automated manifold surgery: constructing geometrically accurate and topologically correct models of the human cerebral cortex," *IEEE transactions on medical imaging*, vol. 20, no. 1, pp. 70–80, 2001.

[6] Rickard Brüel Gabrielsson, Bradley J Nelson, Anjan Dwaraknath, and Primoz Skraba, "A topology layer for machine learning," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 1553–1563.

[7] David Cohen-Steiner, Herbert Edelsbrunner, John Harer, and Yuriy Mileyko, "Lipschitz functions have l p-stable persistence," *Foundations of computational mathematics*, vol. 10, no. 2, pp. 127–139, 2010.

[8] Xiaoling Hu, Fuxin Li, Dimitris Samaras, and Chao Chen, "Topology-preserving deep image segmentation," *Advances in neural information processing systems*, vol. 32, 2019.

[9] Armin Iraji, Ashkan Faghiri, Noah Lewis, Zening Fu, Srinivas Rachakonda, and Vince D Calhoun, "Tools of the trade: estimating time-varying connectivity patterns from fmri data," *Social Cognitive and Affective Neuroscience*, vol. 16, no. 8, pp. 849–874, 2021.

[10] Anthony J Bell and Terrence J Sejnowski, "An information-maximization approach to blind separation and blind deconvolution," *Neural computation*, vol. 7, no. 6, pp. 1129–1159, 1995.

[11] Vesa Kiviniemi, Tapani Vire, Jukka Remes, Ahmed Abou Elseoud, Tuomo Starck, Osmo Tervonen, and Juha Nikkinen, "A sliding time-window ica reveals spatial variability of the default mode network in time," *Brain connectivity*, vol. 1, pp. 339–347, 2011.

[12] Ben Cassidy, Caroline Rae, and Victor Solo, "Brain activity: Conditional dissimilarity and persistent homology," in *2015 IEEE 12th International Symposium on Biomed. Imaging (ISBI)*. IEEE, 2015, pp. 1356–1359.