



Published in final edited form as:

Neuroimage. 2021 February 01; 226: 117538. doi:10.1016/j.neuroimage.2020.117538.

rest2vec: Vectorizing the resting-state functional connectome using graph embedding

Zachery D. Morrissey^{a,b}, Liang Zhan^{e,f}, Olusola Ajilore^b, Alex D. Leow^{b,c,d,*}

^aGraduate Program in Neuroscience, University of Illinois at Chicago, Chicago, IL 60612, USA

^bDepartment of Psychiatry, University of Illinois at Chicago, Chicago, IL 60612, USA

^cDepartment of Bioengineering, University of Illinois at Chicago, Chicago, IL 60607, USA

^dDepartment of Computer Science, University of Illinois at Chicago, Chicago, IL 60607, USA

^eDepartment of Electrical and Computer Engineering, University of Pittsburgh, Pittsburgh, PA 15261, USA

^fDepartment of Bioengineering, University of Pittsburgh, Pittsburgh, PA 15261, USA

Abstract

Resting-state functional magnetic resonance imaging (rs-fMRI) is widely used in connectomics for studying the functional relationships between regions of the human brain. rs-fMRI connectomics, however, has inherent analytical challenges, such as how to properly model negative correlations between BOLD time series. In addition, functional relationships between brain regions do not necessarily correspond to their anatomical distance, making the functional topology of the brain less well understood. Recent machine learning techniques, such as word2vec, have used embedding methods to map high-dimensional data into vector spaces, where words with more similar meanings are mapped closer to one another. Inspired by this approach, we have developed the graph embedding pipeline *rest2vec* for studying the vector space of functional connectomes. We demonstrate how *rest2vec* uses the phase angle spatial embedding (PhASE) method with dimensionality reduction to embed the connectome into lower dimensions, where the functional definition of a brain region is represented continuously in an intrinsic “functional space.” Furthermore, we show how the “functional distance” between brain regions in this space can be applied to discover biologically-relevant connectivity gradients. Interestingly, *rest2vec* can be

This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>)

*Corresponding author at: Department of Psychiatry, University of Illinois at Chicago, Chicago, IL 60612, USA.

alexfeuille@gmail.com (A.D. Leow).

Credit authorship contribution statement

Zachery D. Morrissey: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing original draft, Writing review & editing, Visualization, Funding acquisition. **Liang Zhan:** Conceptualization, Methodology, Software, Validation, Formal analysis, Resources, Data curation, Writing review & editing, Supervision, Funding acquisition. **Olusola Ajilore:** Conceptualization, Methodology, Formal analysis, Writing review & editing, Supervision, Project administration, Funding acquisition. **Alex D. Leow:** Conceptualization, Methodology, Formal analysis, Resources, Writing review & editing, Supervision, Project administration, Funding acquisition.

Declaration of Competing Interest

The authors declare no conflicts of interest.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at [10.1016/j.neuroimage.2020.117538](https://doi.org/10.1016/j.neuroimage.2020.117538).

conceptualized in the context of the recently proposed maximum mean discrepancy (MMD) metric, followed by a double-centering approach seen in kernel PCA. In sum, `rest2vec` creates a low-dimensional representation of the rs-fMRI connectome where brain regions are mapped according to their functional relationships, giving a more informed understanding of the functional organization of the brain.

Keywords

Functional neuroimaging; Resting-state fMRI; Connectomics; Gradients

1. Introduction

Neuroimaging data acquired from magnetic resonance imaging (MRI) tend to be vast and high-dimensional. In particular, resting-state functional MRI (rs-fMRI) produces temporal snapshots of the brain's default activity in the absence of tasks, offering a window into the functional macroscale organization of the brain. As computational tools have become more widely available over the past two decades, researchers have applied graph theory-based models to neuroimaging data to study the network properties of the brain, which has grown into the field of connectomics (Sporns et al., 2005). In connectomics analyses, the brain can be represented as an $N \times N$ matrix, where the rows and columns are N brain regions of interest (ROI), and the elements of the matrix represent some measure of connection between them (e.g., number of white matter fibers, Pearson correlation of blood oxygenation level-dependent (BOLD) time series). Given this volume of high-dimensional data, however, one quickly runs into the “curse of dimensionality.” Originally coined by Bellman (1961), the term refers to the challenge of visualizing and analyzing high-dimensional data. Because the number of points in a Cartesian space grows exponentially with increasing dimensions, high-dimensional spaces become extremely sparse, an effect known as the “empty space phenomenon.” Consequently, this makes understanding the properties of these data more difficult, as metric comparisons become less effective with increasing dimensionality (Lee and Verleysen, 2007).

There are a variety of dimensionality reduction techniques that address this problem. These methods work by embedding a high-dimensional manifold, represented by the discrete points of the data, into a lower dimension (e.g., two or three dimensions), which can then be visualized. This process becomes complicated, however, if the manifold of the underlying data is nonlinear, as is thought to be the case with neuroimaging data (Gerber et al., 2010; McClurkin et al., 1991; Wolz et al., 2012; Ye et al., 2015). One of the most well-known example cases of a nonlinear manifold is the 3D Swiss roll. Nonlinear dimensionality reduction techniques, such as isometric mapping (isomap), address the characteristic Swiss roll problem by preserving the intrinsic geometry of nonlinear manifolds (i.e., unrolling the Swiss roll) in lower-dimensional spaces (Tenenbaum et al., 2000; Ye et al., 2015).

Negative correlations also remain a challenging factor in rs-fMRI connectomics, as they are more difficult to interpret using network models. Simpler models generally either threshold or apply other transformations, such as taking the absolute value, to remove negative

Author Manuscript

correlations; this process, however, likely removes substantive dynamics of brain connectivity (Rubinov and Sporns, 2011). Although some analyses account for negative correlations, these often introduce additional parameters that must be arbitrarily set to determine their relative contribution (Rubinov and Sporns, 2011). Previously, we introduced probability associated community estimation (PACE) (Zhan et al., 2017) and phase angle spatial embedding (PhASE) (Morrissey et al., 2018) to address these challenges. These methods take inspiration from the Ising model from statistical mechanics, where magnetic ions are designated with either in-phase or out-of-phase spin state configurations (Pekalski, 2001). We adapted this model to describe the phase relationship between regions of the brain, where each brain region is an N -dimensional vector whose elements are defined by its phase coupling with every other region in the brain.

Author Manuscript

Having generated this embedding space, how might the organization of the functional connectome best be visualized and understood? In the domain of natural language processing, Mikolov et al. (2013) created the word2vec method to map words with similar meanings near one another in a vector space (e.g., “king” is close to “man,” “queen” is close to “woman”). Inspired by this approach, we propose a novel graph embedding pipeline, rest2vec, that uses this phase angle representation with the nonlinear dimensionality reduction method isomap to embed the functional connectome in a lower-dimensional embedding based on its functional relationships. In this space, the functional connectome is arranged by its *intrinsic geometry*, where the regions of the brain are mapped according to their functional connectivity independent of anatomical constraints. Here, the Euclidean distance between regions in this space can be thought of as an intrinsic “functional distance,” with similar regions having a short functional distance between one another and dissimilar regions having a large functional distance between one another.

Author Manuscript

We show how this vectorized approach has implications for detecting connectivity gradients by linking rest2vec to the maximum mean discrepancy (MMD) metric. The MMD was originally developed as a metric describing the distance between probability distributions (Gretton et al., 2012). Here, we treated the MMD as a modularity index, similar to Q -based maximization methods (Blondel et al., 2008), such that, when maximized, it detects the sets of brain regions with the most dissimilar functional connectivity. By reformulating this connectome modularity problem in a probabilistic sense, we are able to generate continuous community assignment values for each region, as opposed to a binary classification. Finally, we also show how brain regions in the rest2vec embedding space can be mapped to behavioral metrics using the Neurosynth meta-analysis database (Yarkoni et al., 2011). Together, rest2vec uses nonlinear dimensionality reduction and manifold learning techniques to create a low-dimensional representation of the rs-fMRI connectome where brain regions are mapped according to their functional relationships, giving a more informed understanding of the functional organization of the brain.

2. Materials and methods

2.1. Dataset description

Two independent and publicly available rs-fMRI connectome datasets composed of healthy subjects were used: one from the Functional 1000 (F1000) Connectomes Project (Biswal et

al., 2010) with 177 regions of interest (ROI) available through the USC Multimodal Connectivity Database (1000_Functional_Connectomes study, <http://umcd.humanconnectomeproject.org/umcd/default/index>), and one by Diez et al. (2015) with 2514 ROIs available through the NeuroImaging Tools & Resources Collaboratory (NITRC) (https://www.nitrc.org/projects/biocr_hcatlas/). Each of these studies were performed in accordance with their institution’s respective ethics committees. These are referred to as the “F1000” and “Diez” datasets hereafter.

The F1000 dataset has data from 986 subjects collected across multiple sites using a common scanning and preprocessing protocol. Data were motion corrected and spatially smoothed with a 6 mm FWHM Gaussian kernel. A band-pass filter was applied between 0.005–0.01 Hz. Nuisance parameters, CSF, white matter, and the average global signal were regressed out prior to spatial normalization to the MNI152 template. Spatially constrained spectral clustering was used to determine the ROIs (Craddock et al., 2012). The average difference in age between male ($N = 426$, $M \pm SD = 28.7 \pm 12.7$) and female ($N = 560$, $M \pm SD = 27.9 \pm 12.7$) subjects in the F1000 dataset was 0.83 years and was not statistically significant ($t(984) = 1.025$, $p = 0.306$).

The Diez dataset has 12 subjects (6 male) with a mean age of 33.5 ± 8.7 years; no individual subject ages were reported. An interleaved gradient-echo EPI sequence was used to acquire BOLD T2* images (scanning time = 7.28 min, 200 volumes total). Motion correction, slice-timing correction, and smoothing (6 mm FWHM Gaussian kernel) were applied. A band-pass filter was applied between 0.001–0.08 Hz; linear and quadratic trends were also removed. Motion, CSF, white matter, and average global signal were regressed out prior to spatial normalization to the MNI152 template. Spatially constrained clustering (Craddock et al., 2012) was used to determine the 2514 ROIs. Finally, the Pearson correlation coefficient of the BOLD time series between ROIs was calculated as the measure of functional connectivity between ROIs (Diez et al., 2015).

2.2. rest2vec

The pipeline for rest2vec is shown in Fig. 1. Rest2vec aims to create a graph embedding of rs-fMRI connectomes by transforming positive and negative edges into N -dimensional phase angle vectors that can then be represented in a low-dimensional embedding using nonlinear dimensionality reduction. Briefly, we first computed the probability of observing a negative edge between all pairs of regions across all subjects to form the probability matrix \mathbf{P}^- . This probability is then used to determine the phase angle $\Theta_{i,j}$ between regions to create the phase angle spatial embedding (PhASE) matrix Θ . This process embeds the phase relationship between all regions in an N -Euclidean space and transforms the values between 0 (fully in-phase) and $\pi/2$ (fully out-of-phase). The intrinsic geometry of the connectome was then visualized in two dimensions using the nonlinear dimensionality reduction method isomap (Tenenbaum et al., 2000). Finally, we use kernel functions to link rest2vec to the maximum mean discrepancy MMD metric (Gretton et al., 2012) to demonstrate how rest2vec can be used to study functional connectivity gradients. The representative matrices for each step are displayed in Figure A.1.

2.2.1. Phase angle spatial embedding (PhASE)—A functional connectome derived from rs-fMRI is defined as an undirected graph $G(V, E)$, composed of a set of vertices V representing the brain regions of interest (ROI), and signed, weighted edges E describing the measure of connectivity between them based on their BOLD time series. Typically, some measure of correlation, e.g., Pearson correlation, between BOLD time series is used to describe the functional connectivity between ROIs.

Previously, we introduced probability associated community estimation (PACE) (Zhan et al., 2017), and phase angle spatial embedding (PhASE) (Morrissey et al., 2018) for encoding resting-state fMRI connectomes based on the phase relationship between brain regions (Morrissey et al., 2018) to account for negative correlations in functional connectomes. We begin by briefly summarizing these procedures in the context of rest2vec.

Let \mathbf{W} be an $N \times N \times S$ array (i.e., a tensor) composed of $N \times N$ weighted, signed functional connectomes for N regions and S observations. In our case, we consider the observations from a group of S subjects. Given some weight of functional coupling between regions i and j (e.g., Pearson correlation), we define the probability of negative correlation matrix \mathbf{P}^- where each element $P_{i,j}^-$ is the probability of observing a negative edge between i and j defined as $P_{i,j}^- = \text{Prob}([W_{i,j} < 0])$, where, in the case of S subjects, we estimate this probability by the following

$$P_{i,j}^- = \frac{1}{S} \sum_s [W_{i,j,s} < 0], \quad (1)$$

where $W_{i,j,s}$ is the edge between regions i and j for the s th subject, and the Iverson bracket expression $[W_{i,j,s} < 0]$ equals 1 if $W_{i,j,s} < 0$, and 0 otherwise. Because $P_{i,j}^- \in [0, 1]$, it also follows naturally that $P_{i,j}^- + P_{i,j}^+ = 1$.

One advantage of this procedure is that the probability measure defined in Eq. (1) can be defined by the user for their specific context. By taking advantage of this relationship, we then define the phase angle spatial embedding (PhASE) matrix Θ , where the phase angle $\Theta_{i,j}$ between regions i and j is defined as

$$\Theta_{i,j} = \arctan \sqrt{\frac{P_{i,j}^-}{P_{i,j}^+}}. \quad (2)$$

Thus $\Theta_{i,j} \in [0, \pi/2]$, where 0 represents a fully in-phase (co-activating) relationship and $\pi/2$ represents a fully out-of-phase (anti-activating) relationship. Each column of Θ is a vector embedding each region in an N -dimensional Euclidean space such that

$$\Theta_{:,i} = [\Theta_{i,1} \ \Theta_{i,2} \ \dots \ \Theta_{i,N}]^T \in [0, \pi/2]^N.$$

2.2.2. Relation of PhASE to the maximum mean discrepancy—Here we describe how PhASE can be linked to the maximum mean discrepancy (MMD) developed by Gretton et al. (2012) to address the connectome modularity problem. Following the formulation

defined in Gretton et al. (2012), consider the random variables x and y defined on a metric space \mathcal{X} equipped with the metric d , with the corresponding Borel probabilities p and q (i.e., $x \sim p$ and $y \sim q$). Given observations $X := \{x_1, \dots, x_m\}$ and $Y := \{y_1, \dots, y_n\}$ drawn from the probability distributions p and q , $p = q$ if and only if $\mathbf{E}_x[f(x)] = \mathbf{E}_y[f(y)] \forall f \in C(\mathcal{X})$ where $C(\mathcal{X})$ is the space of bounded continuous functions on \mathcal{X} . Next, given a class of functions \mathcal{F} such that $f: \mathcal{X} \rightarrow \mathbb{R}$, the maximum mean discrepancy (MMD) between p and q with respect to \mathcal{F} is defined as

$$\text{MMD}[\mathcal{F}, p, q] := \sup_{f \in \mathcal{F}} (\mathbf{E}_x[f(x)] - \mathbf{E}_y[f(y)]). \quad (3)$$

This can be empirically estimated given X and Y as

$$\text{MMD}[\mathcal{F}, X, Y] := \sup_{f \in \mathcal{F}} \left(\frac{1}{m} \sum_{i=1}^m f(x_i) - \frac{1}{n} \sum_{i=1}^n f(y_i) \right), \quad (4)$$

where m is equal to the number of observations in X and n is equal to the number of observations in Y .

To apply these definitions in the context of connectomics, we use the same definitions of x , y , p , q , X and Y defined above to assign each region to one of the two distributions p or q . Under the working assumption that the distributions of functional modules in the connectome are far apart (i.e., their within-module connections are greater than their between-module connections (Fortunato, 2010)), we thus seek to discover the arrangement of regions such that the MMD between them is maximized.

Using a reproducible kernel Hilbert space (RKHS), the squared form of Eq. (4) can be evaluated using kernel functions as

$$\begin{aligned} \text{MMD}^2[\mathcal{F}, X, Y] &:= \frac{1}{m(m-1)} \sum_{i=1}^m \sum_{j \neq 1}^m k(x_i, x_j) + \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j \neq 1}^n k(y_i, y_j) \\ &\quad - \frac{2}{mn} \sum_{i=1}^m \sum_{j=1}^n k(x_i, y_j). \end{aligned} \quad (5)$$

From Eq. (5), kernel functions can be used, in our case, to compute the kernel matrix \mathbf{K} where the similarity $K_{i,j}$ between regions i and j , in the case of the radial basis function (RBF) kernel k_{RBF} , is given by

$$K_{i,j} = k_{\text{RBF}}(\Theta_{i,:}, \Theta_{j,:}) = \exp\left(-\sigma \sum_{\ell=1}^N |\Theta_{i,\ell} - \Theta_{j,\ell}|^2\right), \quad (6)$$

for phase angle Θ between regions i and j in reference to all other regions indexed by ℓ , for N regions, using the scaling factor σ .

Similarly, we let the cosine kernel k_{cos} evaluating the similarity between regions i and j be defined as

$$K_{i,j} = k_{\cos}(\Theta_{i,:}, \Theta_{j,:}) = \frac{1}{N} \sum_{\ell=1}^N \cos(\Theta_{i,\ell} - \Theta_{j,\ell}), \quad (7)$$

using the same variable definitions as RBF kernel. Because the RBF kernel has an additional parameter, and the cosine kernel has a geometric relation to angles, the cosine kernel is used here; the Taylor expansion of both these kernels can be shown to have similar leading terms.

2.2.3. From the maximum mean discrepancy to rest2vec connectivity

gradients—Following the kernel definitions above and the equation as described by Gretton et al. (2012) (with a modified notation for our purposes), assuming a two-community partition, let the maximum mean discrepancy (MMD) between two modules V_A and V_B be defined as

$$\text{MMD}(V_A, V_B)^2 = \frac{1}{m^2} \sum_{i,j \in V_A} K_{i,j} - \frac{2}{mn} \sum_{i \in V_A, j \in V_B} K_{i,j} + \frac{1}{n^2} \sum_{i,j \in V_B} K_{i,j}, \quad (8)$$

Where $|V_A| = m$, $|V_B| = n$, $|V| = m + n = N$, $V_A \cup V_B = V$, $V_A \cap V_B = \emptyset$, and i is allowed to equal j .

We seek to find a partition between V_A and V_B such that Eq. (8) is maximized. First, we can rewrite $\text{MMD}(V_A, V_B)^2 = \mathbf{y}^T \mathbf{K} \mathbf{y}$ for $\mathbf{y} \in \mathbb{R}^{N \times 1}$ and $\mathbf{K} \in \mathbb{R}^{N \times N}$, where

$$y_i = \begin{cases} \frac{1}{m} & \text{if } i \in V_A \\ -\frac{1}{n} & \text{if } i \in V_B \end{cases}. \quad (9)$$

Thus we define the optimal partition Modularity(V) into modules V_A and V_B as

$$\text{Modularity}(V) = \underset{\substack{V_A, V_B \\ V_A \cup V_B = V \\ V_A \cap V_B = \emptyset}}{\text{argmax}} \text{MMD}(V_A, V_B)^2. \quad (10)$$

This maximization problem can be approximated in a simplified way by relaxing Eq. (10) to a Rayleigh quotient maximization problem. Letting \mathbf{y} be defined as above, where

$$\|\mathbf{y}\| = \left(\frac{N}{mn}\right)^{1/2}, \quad (11)$$

we perform change of variables to the unit length vector \mathbf{v} , where

$$\mathbf{v} = \left(\frac{mn}{N}\right)^{1/2} \mathbf{y}, \quad (12)$$

and $\|\mathbf{v}\|^2 = 1$, $\mathbf{v}^T \mathbf{1} = 0$, where $\mathbf{1} = [1 \dots 1]^T$, $\mathbf{1} \in \mathbb{R}^{N \times 1}$. Then we can rewrite Eq. (10) in terms of \mathbf{v} to define the partition that maximizes $\text{MMD}(V_A, V_B)^2$ as

$$\text{Modularity}(V) = \underset{\|\mathbf{v}\| = 1, \mathbf{v}^T \mathbf{1} = 0}{\text{argmax}} \frac{N}{mn} \mathbf{v}^T \mathbf{K} \mathbf{v}. \quad (13)$$

To compute $\text{MMD}(V_A, V_B)^2$ in Eq. (13), in the context of connectomics, requires ground truth knowledge of m and n in advance. Thus, to account for this, we assume that N is large and that the two communities V_A and V_B are approximately the same size such that $|m-n| \in \mathcal{O}(N)$. The normalization factor in Eq. (13) can then be simplified to

$$\text{Modularity}(V) \approx \frac{4}{N} \underset{\|\mathbf{v}\| = 1, \mathbf{v}^T \mathbf{1} = 0}{\text{argmax}} \mathbf{v}^T \mathbf{K} \mathbf{v}. \quad (14)$$

Finally we relax the constraints of \mathbf{v} from $v_i \in \left\{ \sqrt{\frac{|V_B|}{N|V_A|}}, -\sqrt{\frac{|V_A|}{N|V_B|}} \right\}$ taking only two values to taking any real values such that $\mathbf{v}^* \in \mathbb{R}^N$. These relaxed constraints allow us to conveniently reframe Eq. (13) as a Rayleigh quotient maximization problem. We account for arbitrary origin for the Rayleigh quotient maximization by centering the kernel similarity matrix \mathbf{K} to $\tilde{\mathbf{K}} = \mathbf{C}_N \mathbf{K} \mathbf{C}_N$, where the centering matrix $\mathbf{C}_N = \mathbf{I}_N - \frac{1}{N} \mathbf{J}_N$, $\mathbf{C}_N \in \mathbb{R}^{N \times N}$, $\mathbf{I}_N \in \mathbb{R}^{N \times N}$ is the identity matrix, and $\mathbf{J}_N \in \mathbb{R}^{N \times N}$ is the ones matrix (i.e., $\mathbf{1}\mathbf{1}^T$).

Rather than finding $\text{MMD}(V_A, V_B)^2$ as a function of the partition, we approximate the optimal partition $\text{Modularity}(V)$ by finding the vector \mathbf{v}^* that maximizes the Rayleigh quotient such that

$$\text{Modularity}(V) \approx \frac{4}{N} \underset{\mathbf{v}^* \neq 0, \mathbf{v}^* \in \mathbb{R}^N}{\text{argmax}} \frac{\mathbf{v}^{*T} \tilde{\mathbf{K}} \mathbf{v}^*}{\mathbf{v}^{*T} \mathbf{v}^*}. \quad (15)$$

We can then compute the mapping vector \mathbf{v}^* that maximizes the Rayleigh quotient by computing the eigenvector \mathbf{q} of $\tilde{\mathbf{K}}$ corresponding to the largest eigenvalue λ_{\max} of $\tilde{\mathbf{K}}$. Similar to the Fiedler vector in spectral clustering methods (Weisstein, 1999), the elements of \mathbf{q} assign both community affiliation based on its sign (+ or -) as well as magnitude. Further, \mathbf{q} can be binarized to determine discrete community labels for each region as

$$\mathbf{v}^* = \begin{cases} i \in V_A & \text{if } q_i \geq 0 \\ i \in V_B & \text{if } q_i < 0 \end{cases}, \quad \forall i = 1, 2, 3, \dots, N. \quad (16)$$

In sum, these derivations suggest that connectivity gradients informed by `rest2vec` can be extracted by considering the top eigenvectors of the centered kernel matrix $\tilde{\mathbf{K}}$.

2.3. Visualizing `rest2vec` in 2D using nonlinear dimensionality reduction

Isomap (Tenenbaum et al., 2000) was used to reduce the PhASE matrix Θ as a manifold $\mathcal{M} \in \mathbb{R}^{N \times N}$ to a d -dimensional embedding $\mathbf{Y} \in \mathbb{R}^{N \times d}$, where $d < N$. Isomap is advantageous for this procedure as it is a nonlinear technique, using methods such as

Dijkstra's algorithm (Dijkstra, 1959) to compute the geodesic distances between vertices in high-dimensional space. Nonlinear methods importantly address the Swiss roll problem faced by traditional linear methods such as PCA and MDS (Tenenbaum et al., 2000).

In our case, we used $k = 12$ and $k = 50$ nearest neighbors, for the F1000 and Diez datasets, respectively, to reduce to two dimensions using the Isomap implementation in the Scikit-learn version 0.21.3 library (Pedregosa et al., 2011). To account for any bias due to the choice of nearest neighbors, we also re-calculated the isomap embedding using $k = \{20, 40, 60, 80, 100\}$ with the Diez dataset as an example. The residual variance (Tenenbaum et al., 2000) $1 - R^2(D_M, D_Y)$ was used to quantify how similar the pairwise distances were between the original and low-dimensional spaces, and was shown to be consistent across a reasonable range of values for k (Figure A.3).

Because the isomap procedure centers data about the origin, and by Eq. (2) the phase angle between perfectly in-phase regions is zero, we analyzed each region's Euclidean distance to the origin in this space to observe how the phase relationship between regions is preserved with respect to its low-dimensional embedding. After generating the isomap embedding, the distance D_i to the origin of the isomap space $[0 \dots 0] \in \mathbb{R}^{1 \times d}$ for the i th region was calculated using the Euclidean distance in the 2D isomap embedding.

Because of its natural representation for distance to the origin, the data was transformed to polar coordinates of radius r and angle θ using the polar transformation

$$r = \sqrt{x^2 + y^2} \quad (17)$$

$$\theta = \text{atan2}(y, x) \quad (18)$$

to visualize the functional embedding space.

2.4. Statistical analysis of rest2vec

2.4.1. Functional community mapping analysis—We used the Louvain community algorithm (Blondel et al., 2008) as implemented in the Brain Connectivity Toolbox (Rubinov and Sporns, 2010) to independently evaluate the mapping of functional communities in the low-dimensional embedding, as well as the contribution of negative edges on functional community clustering. The Louvain algorithm was applied on the group average Pearson correlation connectome (\mathbf{W}), the thresholded connectome with only positive weights (\mathbf{W}^+) and the absolute value of the thresholded connectome ($|\mathbf{W}|$). The parameter γ was set to 1 for each connectome. For \mathbf{W} , the negative weights were treated symmetrically.

The silhouette score (Pedregosa et al., 2011; Rousseeuw, 1987), a metric commonly used to evaluate clustering algorithms such as k -means, was used to assess the clustering performance of the functional communities in the embedding space. We also computed the silhouette score from a random permutation of the Louvain labels from \mathbf{W} as a negative control for comparison. Because the Louvain algorithm returns different results with each

run, we performed 100 iterations of the algorithm to compute the mean and standard deviation of the overall silhouette coefficient for each connectome.

2.4.2. Using k-means clustering to define gradient clusters in the rest2vec embedding—*k*-means clustering was used to formally classify clusters for regions (such as the precuneus) that had heterogeneous mappings in the isomap embedding. The *k*-means clustering algorithm was performed using the Scikit-learn implementation (Pedregosa et al., 2011) for $k = 2$ clusters in the isomap embedding. The same seed value was used to ensure reproducible results.

To determine how affiliated other (non-precuneus) regions were to either of the two clusters, regions were first assigned to the precuneus cluster they were closest to in the isomap embedding. A diverging cluster affiliation scale was computed based on the Euclidean distance of each region to its precuneus cluster’s centroid in the isomap embedding, which we termed “intrinsic functional distance,” such that regions with more positive or negative values were closer to the centroid of their respective precuneus cluster. The cluster affiliation a_i was defined as

$$a_i = \begin{cases} \max(d_{C_0}) - d_{C_0,i} & \text{if } i \in C_0 \\ d_{C_1,i} - \max(d_{C_1}) & \text{if } i \in C_1 \end{cases} \quad (19)$$

where d is the intrinsic functional distance from region i to the centroid of cluster C .

2.4.3. rest2vec consistency with Neurosynth meta-analysis association terms—The Neurosynth meta-analysis database (Yarkoni et al., 2011) was used to search for behavioral and psychiatric terms to obtain a NIfTI image containing the association measured between the key term and fMRI activity across studies in the database. For this study, we used the “association test” images provided by Neurosynth, which provides the z -score describing the association between the key term and each voxel.

In order to map this data to the rest2vec embedding, we used the MNI coordinates from the Diez dataset as a lookup value to obtain the z -score of the corresponding voxel in the Neurosynth association test NIfTI file. The glass brain and embedding data was then plotted such that the regions with z -scores > 0 were mapped to a color gradient; all other regions were set to a uniform gray value in the background.

2.5. Code and data availability

All code used to produce the results and figures is available online via GitHub (<https://github.com/zmorrissey/rest2vec>) and our laboratory website (<http://brain.uic.edu/>). We also provide the MNI (x, y, z) coordinates and the rest2vec embedding (x, y)-coordinates for each dataset.

2.5.1. Statistics—The StatsModels library version 0.10.1 for Python (Seabold and Perktold, 2010) was used for statistical analyses. Student’s independent t-test was used to test if there were any differences in age between male and female subjects for the F1000 dataset. The ordinary least squares (OLS) method was used to fit the parameters for the

linear regression between isomap distance to origin and phase angle. For silhouette score analysis, a one-way ANOVA followed by Tukey's *post hoc* test was used.

2.5.2. Visualization—Graphics were drawn using the Matplotlib version 3.1.1 (Hunter, 2007) and Seaborn version 0.9.0 (Seaborn, 2019) libraries using Python version 3.7.3 from the Anaconda distribution (Anaconda, 2018). Glass brain figures were visualized using the `plot_connectome` function from the Nilearn version 0.6.2 library (Abraham et al., 2014). Inkscape version 0.92 was used for final arrangement of some figures (Albert et al., 2019).

Brain surface plots were created by representing the $N \times 4$ array, consisting of the MNI (x, y, z)-coordinates for all N -regions, and the $N \times 1$ vector containing the data value associated with each region, as a 3D volume. For brain distance maps, the intrinsic functional distance vector was made by computing the Euclidean distance between the mean (x, y)-coordinates of the anatomical region in the isomap embedding and all other regions. For regions that had heterogeneous mapping (i.e., multiple clusters) in the isomap space, k -means clustering was performed to calculate cluster affiliations for each region as described above in Eq. (19).

The 3D volume containing the original data was then interpolated using a linear grid interpolation and registered to the MNI template volume with 12 degrees of freedom using the FLIRT tool in the FSL (Jenkinson et al., 2012) interface from the Nipype version 1.3.0-rc1 library (Gorgolewski et al., 2011). The interpolated 3D volume was mapped to the Freesurfer pial surface template using the `vol_to_surf` function from the Nilearn library. The surface data was then visualized using the `plot_surf_stat_map` function from the Nilearn library.

3. Results

3.1. Functional connectivity gradients with rest2vec

3.1.1. The maximum eigenvector from rest2vec defines a principal functional connectivity gradient—After applying the rest2vec pipeline to the F1000 and Diez datasets, we used the maximum eigenvector of the double-centered kernel similarity matrix $\tilde{\mathbf{K}}$ to extract the principal connectivity gradient of the functional connectome (Fig. 2). We visualized the anatomical and functional embedding by anatomical lobe to visualize how each lobe participates in the principal gradient. We observed a symmetrical partition when viewed in the functional embedding space (Fig. 2, middle), primarily demarcating between the task-negative network (TNN) and the default mode network (DMN). When the brain is faceted by lobe affiliation, several notable patterns emerge. For example, the frontal lobe is demarcated into the prefrontal cortex (PFC) and pre-motor areas and dorsolateral PFC anatomically, which are respectively situated in opposite quadrants of the functional embedding. In addition, the parietal lobe is split largely into default mode network (DMN) regions – including regions of the inferior parietal lobule and precuneus – and primary and secondary unimodal areas, including somatosensory cortices and areas involved in visual processing. Consistent with this observation, the occipital lobe has the largest proportion of regions belonging to the putative TPN at 83.4%. Similar demarcations can be observed in the F1000 dataset as well (Fig. A.7).

In addition, the second and third eigenvalues of the kernel similarity matrix $\widetilde{\mathbf{K}}$ account for a notable portion of the variance of the data (Fig. A.5), so we also examined the gradients given by these two eigenvectors (Figure A.8, Figure A.9). The second eigenvector gradient revealed a rostral-caudal gradient, most notable for the sub-lobar, limbic, and temporal lobes, and a medial-lateral gradient for the parietal and frontal lobes. The third eigenvector gradient revealed a largely anterior-posterior gradient across all lobes except for the parietal and occipital lobes, which were dominated by a single community. Interestingly, there is a notable hemispheric asymmetry in the frontal lobe in areas of the left hemisphere, suggestive of language and speech regions of the brain. Taken together, these data suggest that rest2vec is able to recover biologically-relevant connectome gradient properties, while also accounting for the presence of negative edges, thereby removing heuristic steps that may bias downstream analyses as a result.

3.1.2. Region-specific gradients defined by intrinsic functional distance in the rest2vec space

While the kernel similarity matrix eigenvectors from rest2vec can reveal unbiased functional connectivity gradients, we then asked whether one could investigate specific brain regions of interest (ROIs) to query region-specific gradient connectivity information. We hypothesized that the Euclidean distance within the rest2vec isomap embedding space could be used as a metric of “intrinsic functional distance” between brain regions, where if two regions are close together in the embedding space they are more functionally coupled. From Eq. (2), lower values of $\Theta_{i,j}$ indicate a more in-phase relationship between regions. Thus we hypothesized that more in-phase regions would be embedded closer to the origin of the isomap space, whereas more out-of-phase regions would be embedded further from the origin. The 2-norm of each N -dimensional vector of the PhASE matrix $\|\Theta_{i,:}\|$ was used as a summary measure of each region’s overall phase. For each dataset, there was a statistically significant positive correlation between each region’s $\|\Theta_{i,:}\|$ and its distance from the origin of the 2D isomap embedding (F1000 dataset: $F(1, 175) = 87.55$, $R^2 = 0.33$, $r = 0.58$, $p < 0.0001$; Diez dataset: $F(1, 2512) = 340.2$, $R^2 = 0.119$, $r = 0.34$, $p < 0.0001$) (Figure A.2, left). This pattern can be seen when the rows and columns of the PhASE matrix are sorted by ascending $\|\Theta_{i,:}\|$ values, in particular for the coarser parcellation from the F1000 dataset (Figure A.2, right). Together this suggests that regions mapped closer to the origin were more in-phase with other regions, whereas more out-of-phase regions were mapped further from the origin.

To examine this relationship further, we faceted the anatomical and functional embeddings of the Diez dataset by anatomical lobe affiliation ranked by ascending distance to the origin (Fig. 3). Notably, the brainstem displayed the most centrally embedded regions (median distance = 6.9), followed by (in ascending order): sub-lobar, limbic lobe, temporal lobe, frontal lobe, cerebellum, parietal lobe, and occipital lobe regions. At the other extreme, the occipital lobe displayed the most distant and densely clustered representation in the embedding space (median distance = 24). Examination of the phase angle vectors for occipital lobe regions revealed highly in-phase relationships within the occipital lobe, while regions outside the occipital lobe were mostly out-of-phase (Fig. A.10). Since the occipital lobe and large portions of the parietal lobe (e.g., motor cortices), and cerebellum are mapped further in the periphery, this suggests that regions involved in primary sensory processing are

mapped further in the periphery, while regions such as the brainstem, thalamus, and heteromodal areas have more in-phase relationships and are mapped closer to the origin. Interestingly, some lobes, such as the frontal lobe, appeared to have non-normal distance distributions. Further inspection revealed this bimodal distribution to be composed of primarily ROIs from the middle frontal gyrus below the median, and superior and medial frontal and precentral gyri, above the median, among other regions (Figure A.11). Additionally, those regions above the median were more strongly affiliated with different functional communities, whereas those below the median were less affiliated with either functional community, discussed in more detail below.

The F1000 dataset also showed the occipital lobe being furthest from the origin and densely clustered in the embedding space (Figure A.4). Although there were fewer brainstem ROIs in the F1000 dataset, the brainstem, sub-lobar, and temporal regions likewise were more centrally-embedded in the F1000 dataset. Interestingly, the coarser parcellation of the F1000 dataset showed similar median distances across the parietal, frontal, and limbic lobes. This suggests that while the overall distance distributions are similar across parcellations, finer parcellations may intrinsically be able to detect more subtle differences in how regions are mapped in the embedding space. Furthermore, the Diez dataset also include ROIs from the cerebellum, which are likely to influence the topology of the embedding space compared to the F1000, which does not include the cerebellum. Together, these results suggest that the overall distance distributions reveal consistent trends in different parcellations, and that more granular parcellations, as well as contribution of brainstem and cerebellum may contribute additional information to the embedding topology.

Given that the distance to the origin of the isomap embedding preserved phase coupling characteristics across anatomical regions, we next asked if the intrinsic functional distance between regions in this space could reveal biologically-relevant connectivity patterns. As a proof of concept, we computed the intrinsic functional distance for regions that occupy distinct locations in the embedding space, namely the occipital lobe, which was densely clustered in the periphery, and the hippocampus, which was clustered centrally in the embedding. When the intrinsic functional distance to the occipital lobe was mapped as a color gradient on the brain surface, the dorsal and ventral visual streams (Goodale, 2011; Goodale and Milner, 1992; Ungerleider, 1982) became apparent (Fig. 4, left), consistent with the hypothesis that distance in this embedding space preserves functionally relevant information. In contrast, the hippocampus also had a relatively homogeneous cluster in the isomap embedding, but showed a much more distributed surface map gradient to regions of the default mode network (DMN), such as the precuneus, prefrontal cortex, thalamus, and inferior parietal lobule (Fig. 4, right).

While certain anatomical regions showed a relatively homogeneous clustering in the isomap embedding, such as the occipital lobe, others showed heterogeneous clustering patterns. Thus we hypothesized that *rest2vec* could be used to identify functional subnetworks within individual regions based on their clustering within the isomap embedding. As a test case, we examined the isomap embedding pattern for the precuneus, which is known to participate in different networks across its dorsal-anterior/ventral-posterior axes (Cavanna and Trimble, 2006; Zhang and Li, 2012). The bivariate kernel density estimate plot of the precuneus ROIs

in the Diez dataset suggested the presence of two primary clusters (Fig. 5, top left). To investigate this further, we used *k*means clustering to formally assign precuneus ROIs to one of two clusters (Fig. 5, top middle). A meta-cluster was made that included all other regions in the Diez dataset by assigning regions to the precuneus cluster they were closer to. We then measured the intrinsic functional distance between each region to its precuneus cluster centroid to assign an affiliation value to each region (Fig. 5, top right). The brain surface map projection of these data demarcated these two cluster centroids into the dorsal-anterior precuneus and the ventral-posterior precuneus (Fig. 5, bottom). The dorsal-anterior cluster of the precuneus was most strongly affiliated with the occipital and superior parietal regions, as well as the paracentral lobule, middle and superior temporal cortices, and thalamus (Fig. 5, middle). The ventral-posterior cluster of the precuneus was most strongly affiliated with the hippocampus, cuneus, cerebellum, parahippocampal cortex, posterior cingulate cortex, calcarine cortex, amygdala, and superior occipital cortices. Thus we showed that performing clustering analyses on a particular region in this space (e.g., precuneus) revealed multiple, distinct connectivity gradients, which have been previously reported in the literature. Together, this demonstrates how *rest2vec*'s intrinsic functional distance can be used to discover functional connectivity gradients.

3.2. *rest2vec* provides a framework for functional connectome parcellation

Since *rest2vec* could identify functionally relevant connectivity gradients within anatomical lobes, we next asked if *rest2vec* could be used to study functional community parcellation. We hypothesized that if the *rest2vec* embedding is preserving the functional organization of the connectome, then these functional communities would be clustered together in the embedding. Additionally, the contribution of negative edges should also yield the best clustering in the embedding space.

To test this, we computed the community assignment labels for the group average Pearson correlation connectome (\mathbf{W}) using the Louvain community algorithm (Blondel et al., 2008; Rubinov and Sporns, 2010). We did the same procedure using the connectome with negative edges removed (\mathbf{W}^+) and the absolute value of the connectome ($|\mathbf{W}|$). We calculated the silhouette coefficient (Rousseeuw, 1987) as a metric of clustering performance for each connectome, as well as a random permutation of the labels from \mathbf{W} as a negative control, across 100 iterations of the Louvain algorithm. One-way ANOVA followed by Tukey's *post hoc* test revealed that the silhouette coefficient was statistically significantly higher in both datasets for the Louvain community partition using the original connectome \mathbf{W} , followed by \mathbf{W}^+ , $|\mathbf{W}|$, and randomly assigned labels having the lowest silhouette score ($p < 0.001$ for all pairwise comparisons) (Fig. 6, bottom).

These results indicate that (i) the *rest2vec* embedding preserves functional modularity, as demonstrated independently using the Louvain community algorithm, and (ii) the contribution of negative edges yields the highest silhouette score, indicative of better clustering of functional communities. Together, this suggests that *rest2vec* is able to map functional community relationships into a lower-dimensional embedding space, and that the contribution of negative edges best preserves this mapping.

In addition to the Louvain community algorithm, which provides discrete community labels for each ROI, we also used the primary eigenvector \mathbf{q} of the centered kernel similarity matrix $\tilde{\mathbf{K}}$ as a way to partition the connectome into the first two communities, similar to the Fiedler vector in traditional spectral clustering methods. To set the partition, the index of regions corresponding to $q_i \geq 0$ were assigned to community V_A , and the index of regions corresponding to $q_i < 0$ were assigned to community V_B . This approach was validated by iteratively evaluating the maximum mean discrepancy (MMD) across 50 threshold values of \mathbf{q} (Figure A.6). The results suggest that the MMD is maximized when the partition yields communities of approximately equal size, which occurs for both datasets when the partition threshold for $q_i \approx 0$.

When the magnitude and sign of q_i are mapped to a diverging colormap in the isomap space, it was observed that regions closer to the vertical axis appeared more neutral, whereas regions further from the vertical axis were polarized into either community, suggesting these regions are more strongly mapped into that community. As observed earlier in Fig. 2, we found that the partition demarcated into the putative task-positive network (TPN) and task-negative network (TNN, also called default mode network (DMN)). This relationship can be seen when the primary eigenvector gradient is used to sort the Θ matrix for each dataset, where the resulting grid communities show out-of-phase relationships with the other community (Fig. 7, right). This can also be observed anatomically when the eigenvector gradient is mapped to the brain surface (Fig. 7, bottom).

Taken together, these data suggest that rest2vec is able to recover functional community information, while also accounting for the presence of negative edges, thereby removing heuristic steps that may bias downstream analyses as a result. Finally, we show that rest2vec can model these communities using a continuous gradient mapping, going beyond traditional discrete community labeling methods.

3.3. rest2vec is consistent with Neurosynth meta-analysis association terms

Finally, we sought to connect rest2vec to behavioral metrics. We used meta-analysis terms from the Neurosynth database (Yarkoni et al., 2011) to obtain brain MR images containing the FDR-corrected z -score association of each voxel to the key term. Then, we were able to query each MNI coordinate of the Diez dataset with the Neurosynth images to map each region's association z -score in the embedding space (Fig. 8).

We examined the embeddings for a selection of behavioral, psychiatric, and physiological terms to show a range of different traits: “anxiety,” “auditory,” “decision making,” “default mode,” “motor control,” “reward,” “speech,” and “visual.” Interestingly, “reward,” “anxiety,” and “decision making” were centrally located, suggesting that they are mapped with more in-phase regions and cluster together within the embedding space. As a comparison, more physiological terms mapped closely to their respective anatomical clusters, such as “visual,” “auditory,” and “speech.” “Auditory” and “speech” share notably similar mappings, suggestive of their shared function in language processing. Finally, the “default mode” and “motor control” mappings are consistent with both the MMD gradient and Louvain modularity results. Together, this demonstrates that rest2vec can be used to map

Neurosynth meta-analysis terms in the embedding space, offering additional insight into how the organization of the functional connectome can be linked to behavior.

4. Discussion

There have been many diverse and interesting approaches using dimensionality reduction and gradient mapping for studying brain organization. Margulies et al. demonstrated the diffusion map embedding method (Margulies et al., 2016), revealing a principal unimodal/transmodal gradient. In addition, there have been studies linking cytoarchitecture to macroscale neuroimaging data for informing cortical gradient mapping (Huntenburg et al., 2018; Paquola et al., 2019), as well as seeing how conditions such as ischemic stroke affect the gradients within the embedding space (Bayrak et al., 2019). Zhang et al. constructed intrinsic connectivity maps using network motifs reported in the literature with MDS to create a low-dimensional embedding, identifying three gradients across external-internal, modulation-representation, and anatomical centrality features (Zhang et al., 2019). Recently, Bethlehem et al. (2020) investigated how the functional communities in their diffusion map embedding space change across age using a novel dispersion metric. In a different approach, Shafiei et al. (2020) have proposed using tools from the time series analysis field to perform feature extraction on the temporal dynamics of functional connectivity to study the intrinsic dynamics of cortical gradients. Interestingly, the principal gradient identified across many of these techniques bear a striking similarity, suggesting this gradient accounts for much of the variance of intrinsic functional connectivity.

In this study, we presented a novel graph embedding approach for rs-fMRI connectivity using rest2vec. Rest2vec improves upon current methods by using the full range of correlative information and representing the functional relationships of the brain in a low-dimensional embedding. Previous studies have suggested that negative correlations may have important – but still not fully understood – biological roles (Rubinov and Sporns, 2010). There is evidence that weak functional connections significantly contribute to the explanation of cognitive functions in both health and disease (Bassett et al., 2012; Santarnecchi et al., 2014), and have been shown to be important in other types of complex networks (Csermely, 2004; Granovetter, 1983; Ma and Gao, 2012; Onnela et al., 2007; Santarnecchi et al., 2014).

Previous work from our group has demonstrated that using a probability-based divisive approach with permutation testing could recover the hierarchical community structure of rs-fMRI connectomes while preserving negative edges, which we called probability-associated community estimation (PACE) (Zhan et al., 2017). In addition, our previous study (Ye et al., 2015) demonstrated how nonlinear dimensionality reduction and manifold learning techniques could be used to investigate the intrinsic geometry of structural connectomes derived from diffusion imaging. Inspired by these approaches, we sought to develop a method by which rs-fMRI functional connectomes could be represented in their intrinsic geometry while also preserving negative edge relationships. We chose to use the isomap method because it uses a geodesic distance metric for generating the lower-dimensional embedding (Tenenbaum et al., 2000). By doing so, distance in the lower-dimensional

embedding conveys meaningful information, and also yields consistent results across each run, in comparison to stochastic methods, such as *t*-SNE (van der Maaten and Hinton, 2008).

In the context of functional connectivity, converting the embedding coordinate system to a polar representation was an intuitive visualization decision, as it centers the data around the origin, where regions with lower Θ_{ij} values are mapped closer to the origin and regions with higher Θ_{ij} values are mapped in the periphery. Interestingly, regions with a greater number of high Θ_{ij} values (i.e., more out-of-phase relationships) tended to be unimodal and also have low within-cluster Θ_{ij} values, as seen most clearly in the occipital lobe (Fig. A.10). In contrast, more centrally-embedded regions tended to be located in brainstem and sub-lobar regions, and also situated in the “transition zone” of the MMD gradient. This suggests that these regions may be important for between-module communication.

By using lower-dimensional embedding distance metrics, we were able to recover functionally relevant relationships. In the case of the occipital lobe, mapping the intrinsic functional distance to its cluster centroid in the isomap embedding generated a gradient map in the anatomical space of the dorsal and ventral visual streams (Goodale, 2011; Goodale and Milner, 1992). On the dorsal surface, the gradient proximal to the occipital lobe can be seen going to the posterior parietal regions, whereas on the ventral surface the proximal gradient extends from the occipital lobe to the inferior temporal lobe (Fig. 4, left). In another example, the precuneus had two primary clusters in the isomap embedding. When projected onto the brain surface, these two clusters demarcated the dorsal-anterior and ventral-posterior portions of the precuneus (Fig. 5). The dorsal-anterior gradient appeared to primarily consist of the superior parietal, somatomotor, and occipital cortices. The ventral-posterior gradient appeared to be composed of the posterior cingulate, parahippocampal, and superior occipital cortices and the hippocampus.

There is evidence for the dorsal-anterior and ventral-posterior portions of the precuneus being involved in different functions. A rs-fMRI study by Zhang and Li (2012) identified the dorsal and anterior portions of the precuneus having stronger connectivity with the occipital, somatomotor, and posterior parietal cortices and the superior temporal gyri. In addition, they identified the ventral precuneus as being more strongly associated with the middle frontal gyrus, posterior cingulate cortex, cuneus, and calcarine sulcus. This demarcation is thought to be due to the diverse roles of the precuneus. In particular, the dorsal-anterior portion of the precuneus, which has strong connectivity with the occipital and superior parietal cortices, is involved in processing polymodal imagery and visuospatial information, whereas the ventral-posterior precuneus is thought to be more involved in episodic memory retrieval (Cavanna and Trimble, 2006). While the study by Zhang and Li (2012) further subdivided the precuneus into eight clusters in their study, our results were largely consistent with their observations, suggesting that rest2vec can detect heterogeneous connectivity patterns within individual regions.

In addition to representing the intrinsic geometry of functional connectomes, we proposed using the maximum mean discrepancy (MMD) method by Gretton et al. (2012) to partition the connectome into maximally functionally distinct modules. The MMD was originally implemented to detect how different two probability distributions were to test if they were

from the same population (Gretton et al., 2012). For our use case, we maximized the MMD as an objective function to find two populations of brain regions such that their distributions are as distant as possible to identify functional communities. One advantage is this is a vectorized approach and does not rely on iterative methods. In addition, this method offers flexibility in the choice of probability and kernel similarity measures used as input, and so are not limited to only Pearson correlation measures.

When the functional connectome is represented in its intrinsic embedding using nonlinear dimensionality reduction, the MMD gradient elicited a strikingly symmetric representation. Upon closer observation, these two communities were split approximately between the canonical task-positive network (TPN) and the default mode network (DMN), consisting of the precuneus, inferior parietal lobule (IPL), posterior cingulate cortex, hippocampus, and areas of the prefrontal cortex (PFC), among others (Buckner et al., 2008). This initial bifurcation gradient is consistent with previous studies (Margulies et al., 2016; Zhan et al., 2017; Zhang et al., 2019), and is a validation that this embedding procedure is capturing functionally-relevant characteristics. In addition, we showed lobe-specific affiliations for the two communities. These results were consistent with the putative DMN/TPN split. Notably, the IPL and precuneus are shown in contrast to the postcentral regions within the parietal lobe; similarly, the PFC and pre-motor areas show clear boundaries. In addition, the contribution of the second and third eigenvectors offer additional insight into the functional demarcations of the connectome. Together, these results demonstrated that using this MMD maximization approach yielded reproducible and biologically-meaningful connectome partition gradients, and that the properties of these communities can be represented using dimensionality reduction.

In this study, we focused on analyzing the resting-state connectome's intrinsic geometry with *rest2vec*, as it is thought to represent the brain's inherent functional connectivity. In principle, *rest2vec* could be applied similarly to task fMRI connectomes to study how the architecture of the embedding space is altered due to task activity. Cole et al. have previously suggested that there is a "task-general" architecture of the functional connectome that is shared across a variety of tasks (Cole et al., 2014). Future work investigating how task embedding spaces are distinct from rest could lead to interesting insights into the gradient organization of the brain as it engages between rest and task states.

An open question is how the functional embedding space may relate to structural connectivity. For example, structural connectomes have an intuitive relationship to network navigation (Seguin et al., 2018). Seguin et al. (2018) discovered that the navigation routing score for brain regions is negatively correlated with their functional connectivity strength. In our present study, we demonstrated the link between Euclidean distance in the functional embedding space and gradients in the anatomical space. An interesting approach for future work would be to determine the relationship between structural network navigation and intrinsic functional connectome geometry.

Allard and Serrano (2020) recently published a study using hyperbolic space to study structural connectome navigation across multiple species. Of particular interest, the hyperbolic representation of the structural connectome with the Yeo 17 functional network

labels (Yeo et al., 2011) revealed a noticeable segregation of the functional communities. Together, this suggests that representing both structural and functional networks in their intrinsic geometries is preserving organizational properties that are otherwise obscured by the anatomical spatial embedding of the brain; future work that could integrate or compare the structural and functional connectome embeddings may help to shed more light on the brain's communication properties as inferred by these two modalities.

4.1. Limitations and future directions

In this paper, we used rs-fMRI connectomes from a group of subjects in order to compute the probability of there being a negative correlation between each pairwise edge between regions. While this approach led to consistent results across two independent datasets, we did not assess how robust this procedure was to inter-subject variability or the size of groups. In addition, while average rs-fMRI connectomes yield a wealth of functional connectivity information, they are a static representation of a dynamic process. Furthermore, there has been increasing emphasis on individual connectome analysis with aims towards personalized medicine (Finn et al., 2015; Miranda-Dominguez et al., 2014). To that end, future improvements on these methods will need to incorporate dynamic as well as subject-specific analyses of functional connectivity. Recent works by Betzel and Bassett (2017) and Gosak et al. (2019) have suggested the concept of hierarchical or multi-scale networks, which could lead to natural extensions of this work via subject embedding spaces which are in turn composed of network embedding spaces.

While we show how rest2vec can be linked to the maximum mean discrepancy and show the partition gradients for the first three eigenvectors of the kernel similarity matrix, further work will need to be done to develop a hierarchical way to detect N communities with this approach. In addition, more robust methods could be used for maximizing the MMD objective function to avoid the possibility of local maxima to achieve better accuracy. While we highlight the theoretical relationship between rest2vec and the MMD, as rest2vec is a low-dimensional embedding technique users are not limited to using the MMD to study functional communities in the embedding space, as we demonstrated with the widely-used Louvain community algorithm (Fig. 6).

4.2. Conclusions

Rest2vec incorporates both positive and negative edge connectivity using a model inspired by statistical mechanics to transform functional connectome data into phase angle relationships. This representation of the connectome can be combined with nonlinear dimensionality reduction techniques to represent the intrinsic geometry of the functional connectome in a lower-dimensional embedding. Together, these methods allow for a vectorized approach to investigate the functional relationships and connectivity gradients of rs-fMRI data. In addition, we connected rest2vec to the maximum mean discrepancy metric to demonstrate how rest2vec can be used to address the modularity problem as a kernel two-sample test. Finally, we show how rest2vec can be connected to behavioral data using the Neurosynth meta-analysis database to project behavioral scores in the embedding space. In summary, we presented a rs-fMRI connectome graph embedding technique that uses

nonlinear dimensionality reduction and statistical learning methods to create a low-dimensional representation of the intrinsic geometry of the functional connectome.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

This study was supported by the National Institutes of Health (T32AG057468, R21AG056782, R01AG071243) and the National Science Foundation (IIS-1837956).

References

- Albert M, Andler JA, Bah T, Barbry-Blot P, Barraud J-F, Barton C, Baxter B, Beard J, Bintz J, Biro A, Bishop N, Blocher JL, Böck H, Boczkowski T, Boldewyn HB, Borgmann D, Bouclet B, Breuer H, Broberg G, Brown C, Brubaker M, Bruno L, Brynn B, Buculei N, Byak B, Caclin P, Caldwell I, Carmichael G, Catmur E, Celorio C, Cenoz JA, Ceuppens J, Chyla Z, Clausen A, Cliff J, Cook K, Cromwell B, Crosbie R, Cruz J, De-Cooman A, Derezynski M, Díaz D, Dilly B, Doolittle L, Dufour N, Dwyer T, Dziuromanenko MV, Eberl M, Engelen J, Erdelyi M, Erikson U, Falzon N, Faubel S, Felfe F, Fitzsimon A, Flick E, Floryan M, Fowler B, Fred, Gemy C, Giannini S, Gondouin O, Gould T, de Greef T, Grosberg M, Gussem KD, Harrington B, Harvey D, Heckert AA, Hetherington C, Hirth J, 2019. Inkscape <https://inkscape.org/>.
- Abraham A, Pedregosa F, Eickenberg M, Gervais P, Mueller A, Kossaifi J, Gramfort A, Thirion B, Varoquaux G, 2014. Machine learning for neuroimaging with scikit-learn. *Front. Neuroinform* 8. doi:10.3389/fninf.2014.00014.
- Allard A, Serrano MA, 2020. Navigable maps of structural brain networks across species. *PLoS Comput. Biol* 16 (2), e1007584. doi:10.1371/journal.pcbi.1007584. [PubMed: 32012151]
- Anaconda, 2018. Software distribution, anaconda. 12. <https://anaconda.com/>
- Bassett DS, Nelson BG, Mueller BA, Camchong J, Lim KO, 2012. Altered resting state complexity in schizophrenia. *NeuroImage* 59 (3), 2196–2207. doi:10.1016/j.neuroimage.2011.10.002. [PubMed: 22008374]
- Bayrak c.S., Khalil AA, Villringer K, Fiebach JB, Villringer A, Margulies DS, Ovadia-Caro S, 2019. The impact of ischemic stroke on connectivity gradients. *NeuroImage* 24, 101947. doi:10.1016/j.neuroimage.2019.101947. [PubMed: 31376644]
- Bellman RE, 1961. *Adaptive Control Processes: A Guided Tour* Princeton university press.
- Bethlehem RA, Paquola C, Seidlitz J, Ronan L, Bernhardt B, Consortium C-C, Tsvetanov KA, 2020. Dispersion of functional gradients across the adult lifespan. *NeuroImage* 117299. doi:10.1016/j.neuroimage.2020.117299. [PubMed: 32828920]
- Betzl RF, Bassett DS, 2017. Multi-scale brain networks. *NeuroImage* 160, 73–83. doi:10.1016/j.neuroimage.2016.11.006. [PubMed: 27845257]
- Biswal BB, Mennes M, Zuo X-N, Gohel S, Kelly C, Smith SM, Beckmann CF, Adelstein JS, Buckner RL, Colcombe S, Dogonowski A-M, Ernst M, Fair D, Hampson M, Hoptman MJ, Hyde JS, Kiviniemi VJ, Kötter R, Li S-J, Lin CP, Lowe MJ, Mackay C, Madden DJ, Madsen KH, Margulies DS, Mayberg HS, McMahon K, Monk CS, Mostofsky SH, Nagel BJ, Pekar JJ, Peltier SJ, Petersen SE, Riedl V, Rombouts SAR, Rypma B, Schlaggar BL, Schmidt S, Seidler RD, Siegle GJ, Sorg C, Teng G-J, Veijola J, Villringer A, Walter M, Wang L, Weng X-C, Whitfield-Gabrieli S, Williamson P, Windischberger C, Zang Y-F, Zhang H-Y, Castellanos FX, Milham MP, 2010. Toward discovery science of human brain function. *Proc. Natl. Acad. Sci* 107 (10), 4734–4739. doi:10.1073/pnas.0911855107. [PubMed: 20176931]
- Blondel VD, Guillaume J-L, Lambiotte R, Lefebvre E, 2008. Fast unfolding of communities in large networks. *J. Stat. Mech* 2008 (10), P10008. doi:10.1088/1742-5468/2008/10/P10008.

- Buckner RL, Andrews-Hanna JR, Schacter DL, 2008. The brain's default network: anatomy, function, and relevance to disease. *Ann. N. Y. Acad. Sci* 1124 (1), 1–38. doi:10.1196/annals.1440.011. [PubMed: 18400922]
- Cavanna AE, Trimble MR, 2006. The precuneus: a review of its functional anatomy and behavioural correlates. *Brain* 129 (3), 564–583. doi:10.1093/brain/awl004. [PubMed: 16399806]
- Cole MW, Bassett DS, Power JD, Braver TS, Petersen SE, 2014. Intrinsic and task-evoked network architectures of the human brain. *Neuron* 83 (1), 238–251. doi:10.1016/j.neuron.2014.05.014. [PubMed: 24991964]
- Craddock RC, James G, Holtzheimer PE, Hu XP, Mayberg HS, 2012. A whole brain fMRI atlas generated via spatially constrained spectral clustering. *Hum. Brain Map* 33 (8), 1914–1928. doi:10.1002/hbm.21333.
- Csermely P, 2004. Strong links are important, but weak links stabilize them. *Trends Biochem. Sci* 29 (7), 331–334. doi:10.1016/j.tibs.2004.05.004. [PubMed: 15236738]
- Dijkstra EW, 1959. A note on two problems in connexion with graphs 3
- Diez I, Bonifazi P, Escudero I, Mateos B, Muñoz MA, Stramaglia S, Cortes JM, 2015. A novel brain partition highlights the modular skeleton shared by structure and function. *Sci. Rep* 5 (1). doi:10.1038/srep10532.
- Finn ES, Shen X, Scheinost D, Rosenberg MD, Huang J, Chun MM, Papademetris X, Constable RT, 2015. Functional connectome fingerprinting: Identifying individuals using patterns of brain connectivity. *Nat. Neurosci* 18 (11), 1664–1671. doi:10.1038/nn.4135. [PubMed: 26457551]
- Fortunato S, 2010. Community detection in graphs. *Phys. Rep* 486 (3–5), 75–174. doi:10.1016/j.physrep.2009.11.002.
- Gerber S, Tasdizen T, Fletcher PT, Joshi S, Whitaker R, 2010. Manifold modeling for brain population analysis. *Med. Image Anal* 14 (5), 643–653. doi:10.1016/j.media.2010.05.008. [PubMed: 20579930]
- Goodale MA, 2011. Transforming vision into action. *Vis. Res* 51 (13), 1567–1587. doi:10.1016/j.visres.2010.07.027. [PubMed: 20691202]
- Goodale MA, Milner AD, et al., 1992. Separate visual pathways for perception and action
- Gorgolewski K, Burns CD, Madison C, Clark D, Halchenko YO, Waskom ML, Ghosh SS, 2011. Nipype: a flexible, lightweight and extensible neuroimaging data processing framework in python. *Front. Neuroinform* 5. doi:10.3389/fninf.2011.00013.
- Gosak M, Markovi R, Dolenšek J, Rupnik MS, Marhl M, Stožer A, Perc M, 2019. Network science of biological systems at different scales: a review. *Phys. Life Rev* doi:10.1016/j.plrev.2017.11.003.
- Granovetter M, 1983. The strength of weak ties: a network theory revisited. *Sociol. Theory* 1, 201. doi:10.2307/202051.
- Gretton A, Borgwardt KM, Rasch MJ, Schölkopf B, Smola A, 2012. A kernel two-sample test. *J. Mach. Learn. Res* 13 (3), 723–773.
- Huntenburg JM, Bazin P-L, Margulies DS, 2018. Large-scale gradients in human cortical organization. *Trends Cognit. Sci* 22 (1), 21–31. doi:10.1016/j.tics.2017.11.002. [PubMed: 29203085]
- Hunter JD, 2007. Matplotlib: a 2d graphics environment. *Comput. Sci. Eng* 9 (3), 90–95. doi:10.1109/MCSE.2007.55.
- Jenkinson M, Beckmann CF, Behrens TE, Woolrich MW, Smith SM, 2012. FSL. *NeuroImage* 62 (2), 782–790. doi:10.1016/j.neuroimage.2011.09.015. [PubMed: 21979382]
- Lee JA, Verleysen M, 2007. *Nonlinear Dimensionality Reduction* Springer Science & Business Media.
- Ma X, Gao L, 2012. Discovering protein complexes in protein interaction networks via exploring the weak ties effect. *BMC Syst. Biol* 6 (Suppl 1), S6. doi:10.1186/1752-0509-6-S1-S6.
- van der Maaten L, Hinton G, 2008. Visualizing data using t-SNE. *J. Mach. Learn. Res* 9 (11), 2579–2605.
- Margulies DS, Ghosh SS, Goulas A, Falkiewicz M, Huntenburg JM, Langs G, Bezgin G, Eickhoff SB, Castellanos FX, Petrides M, Jefferies E, Smallwood J, 2016. Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proc. Natl. Acad. Sci* 113 (44), 12574–12579. doi:10.1073/pnas.1608282113. [PubMed: 27791099]

- McClurkin J, Optican L, Richmond B, Gawne T, 1991. Concurrent processing and complexity of temporally encoded neuronal messages in visual perception. *Science* 253 (5020), 675–677. doi:10.1126/science.1908118. [PubMed: 1908118]
- Mikolov T, Sutskever I, Chen K, Corrado GS, Dean J, 2013. Distributed representations of words and phrases and their compositionality. In: *Advances in Neural Information Processing Systems*, pp. 3111–3119.
- Miranda-Dominguez O, Mills BD, Carpenter SD, Grant KA, Kroenke CD, Nigg JT, Fair DA, 2014. Connectotyping: model based fingerprinting of the functional connectome. *PLoS One* 9 (11), e111048. doi:10.1371/journal.pone.0111048. [PubMed: 25386919]
- Morrissey Z, Zhan L, Lee H, Keiriz J, Forbes A, Ajilore O, Leow A, Chung M, 2018. Phase angle spatial embedding (phASE): a kernel method for studying the topology of the human functional connectome. In: Frangi AF, Schnabel JA, Davatzikos C, Alberola-López C, Fichtinger G (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*, Vol. 11072. Springer International Publishing, Cham, pp. 367–374. doi:10.1007/978-3-030-00931-1_42.
- Onnela J-P, Saramäki J, Hyvönen J, Szabó G, Lazer D, Kaski K, Kertész J, Barabási AL, 2007. Structure and tie strengths in mobile communication networks. *Proc. Natl. Acad. Sci* 104 (18), 7332–7336. doi:10.1073/pnas.0610245104. [PubMed: 17456605]
- Paquola C, De Wael RV, Wagstyl K, Bethlehem RAI, Hong S-J, Seidlitz J, Bullmore ET, Evans AC, Misisic B, Margulies DS, Smallwood J, Bernhardt BC, 2019. Microstructural and functional gradients are increasingly dissociated in transmodal cortices. *PLoS Biol* 17 (5), e3000284. doi:10.1371/journal.pbio.3000284. [PubMed: 31107870]
- Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M, Duchesnay E, 2011. Scikit-learn: machine learning in python. *J. Mach. Learn. Res* 12, 2825–2830.
- Pełkalski A, 2001. Ising model on a small world network. *Phys. Rev. E* 64 (5). doi:10.1103/PhysRevE.64.057104.
- Rousseeuw PJ, 1987. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *J. Comput. Appl. Math* 20, 53–65. doi:10.1016/0377-0427(87)90125-7.
- Rubinov M, Sporns O, 2010. Complex network measures of brain connectivity: uses and interpretations. *NeuroImage* 52 (3), 1059–1069. doi:10.1016/j.neuroimage.2009.10.003. [PubMed: 19819337]
- Rubinov M, Sporns O, 2011. Weight-conserving characterization of complex functional brain networks. *NeuroImage* 56 (4), 2068–2079. doi:10.1016/j.neuroimage.2011.03.069. [PubMed: 21459148]
- Santarnecchi E, Galli G, Polizzotto NR, Rossi A, Rossi S, 2014. Efficiency of weak brain connections support general cognitive functioning: Efficiency of weak and strong brain connections and intelligence. *Hum. Brain Map* 35 (9), 4566–4582. doi:10.1002/hbm.22495.
- Seabold S, Perktold J, 2010. Statsmodels: econometric and statistical modeling with python 5
- Seaborn, 2019. : Statistical data visualization. <https://seaborn.pydata.org/index.html>.
- Seguin C, van den Heuvel MP, Zalesky A, 2018. Navigation of brain networks. *Proc. Natl. Acad. Sci* 115 (24), 6297–6302. doi:10.1073/pnas.1801351115. [PubMed: 29848631]
- Shafiei G, Markello RD, de Wael RV, Bernhardt BC, Fulcher BD, Misisic B, 2020. Topographic gradients of intrinsic dynamics across neocortex. *Neuroscience* doi:10.1101/2020.07.03.186916. 7.
- Sporns O, Tononi G, Kötter R, 2005. The human connectome: a structural description of the human brain. *PLoS Comput. Biol* 1 (4), e42. doi:10.1371/journal.pcbi.0010042. [PubMed: 16201007]
- Tenenbaum JB, de Silva J, Langford JC, 2000. A global geometric framework for nonlinear dimensionality reduction. *Science* 290 (5500), 2319–2323. doi:10.1126/science.290.5500.2319. [PubMed: 11125149]
- Ungerleider LG, 1982. Two cortical visual systems. *Anal. Vis. Behav* 549–586.
- Weisstein EW, 1999. Fiedler vector <https://mathworld.wolfram.com/FiedlerVector.html>.
- Wolz R, Aljabar P, Hajnal JV, Lötjönen J, Rueckert D, 2012. Nonlinear dimensionality reduction combining MR imaging with non-imaging information. *Med. Image Anal* 16 (4), 819–830. doi:10.1016/j.media.2011.12.003. [PubMed: 22244037]

- Yarkoni T, Poldrack RA, Nichols TE, Essen DCV, Wager TD, 2011. Large-scale automated synthesis of human functional neuroimaging data. *Nat. Methods* 8 (8), 665–670. doi:10.1038/nmeth.1635. [PubMed: 21706013]
- Ye AQ, Ajilore OA, Conte G, GadElkarim J, Thomas-Ramos G, Zhan L, Yang S, Kumar A, Magin RL, Forbes AG, Leow AD, 2015. The intrinsic geometry of the human brain connectome. *Brain Inform* 2 (4), 197–210. doi:10.1007/s40708-015-0022-2. [PubMed: 27747562]
- Yeo BTT, Krienen FM, Sepulcre J, Sabuncu MR, Lashkari D, Hollinshead M, Roffman JL, Smoller JW, Zöllei L, Polimeni JR, Fischl B, Liu H, Buckner RL, 2011. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J. Neurophysiol* 106 (3), 1125–1165. doi:10.1152/jn.00338.2011. [PubMed: 21653723]
- Zhan L, Jenkins LM, Wolfson OE, GadElkarim JJ, Nocito K, Thompson PM, Ajilore OA, Chung MK, Leow AD, 2017. The significance of negative correlations in brain connectivity. *J. Comp. Neurol* 525 (15), 3251–3265. doi:10.1002/cne.24274. [PubMed: 28675490]
- Zhang J, Abiose O, Katsumi Y, Touroutoglou A, Dickerson BC, Barrett LF, 2019. Intrinsic functional connectivity is organized as three interdependent gradients. *Sci. Rep* 9 (1), 15976. doi:10.1038/s41598-019-51793-7. [PubMed: 31685830]
- Zhang S, Li Cs.R., 2012. Functional connectivity mapping of the human precuneus by resting state fMRI. *NeuroImage* 59 (4), 3548–3562. doi:10.1016/j.neuroimage.2011.11.023. [PubMed: 22116037]

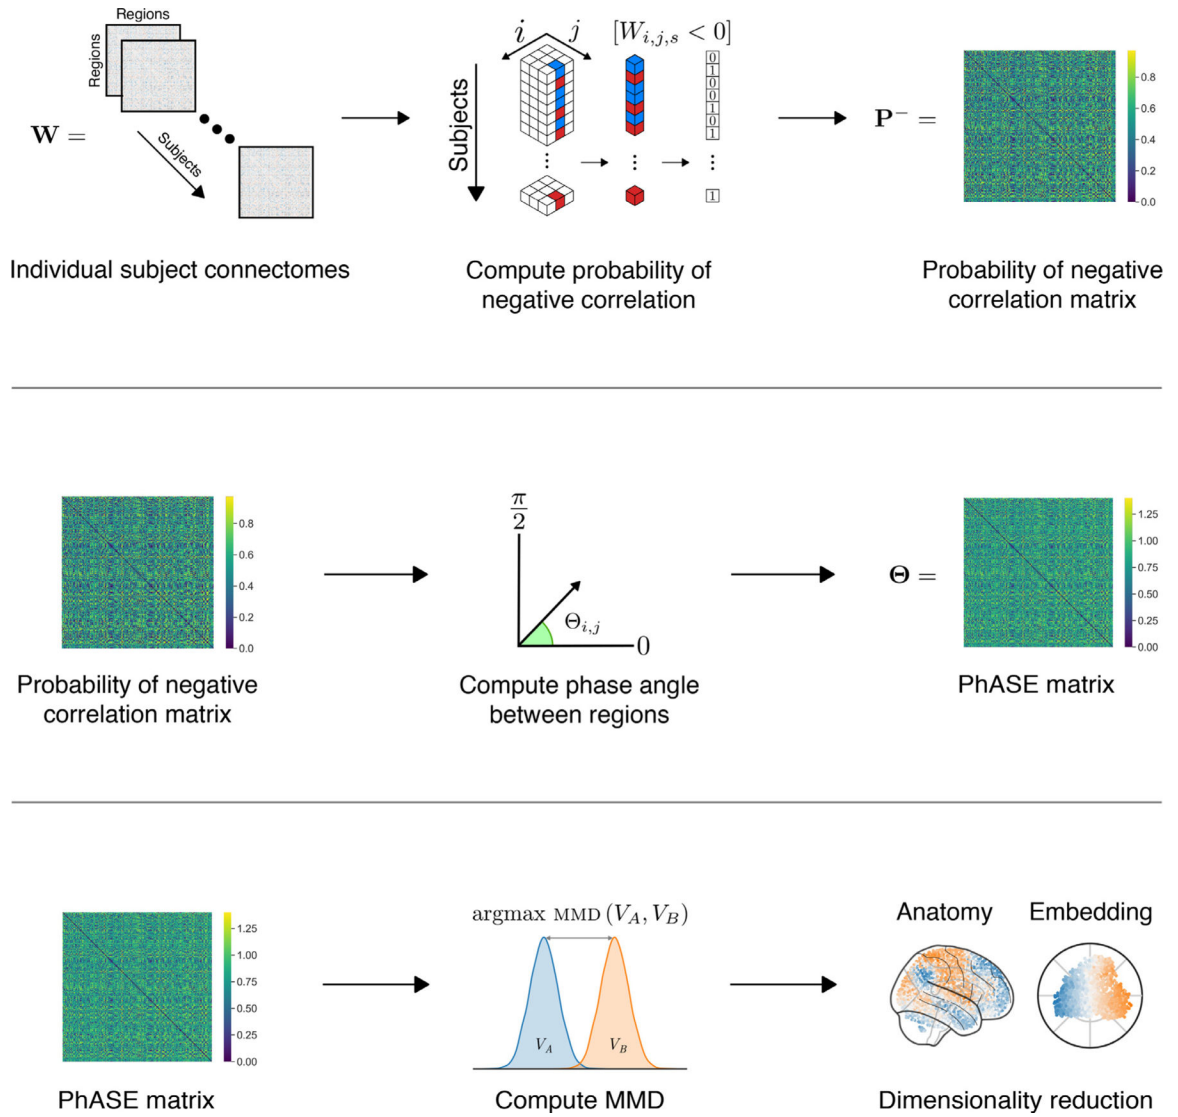


Fig. 1. rest2vec processing pipeline. (Top) The frequency of observing a negative edge between regions i and j across all subjects in the $N \times N \times S$ array \mathbf{W} of rs-fMRI connectomes is computed to form the probability of negative correlation matrix \mathbf{P}^- . (Middle) The phase angle transformation is applied to compute the phase angle spatial embedding (PhASE) matrix Θ . (Bottom) Dimensionality reduction and the maximum mean discrepancy (MMD) are used to analyze the properties of the new embedding space, where the functional connectome is represented by its intrinsic geometry.

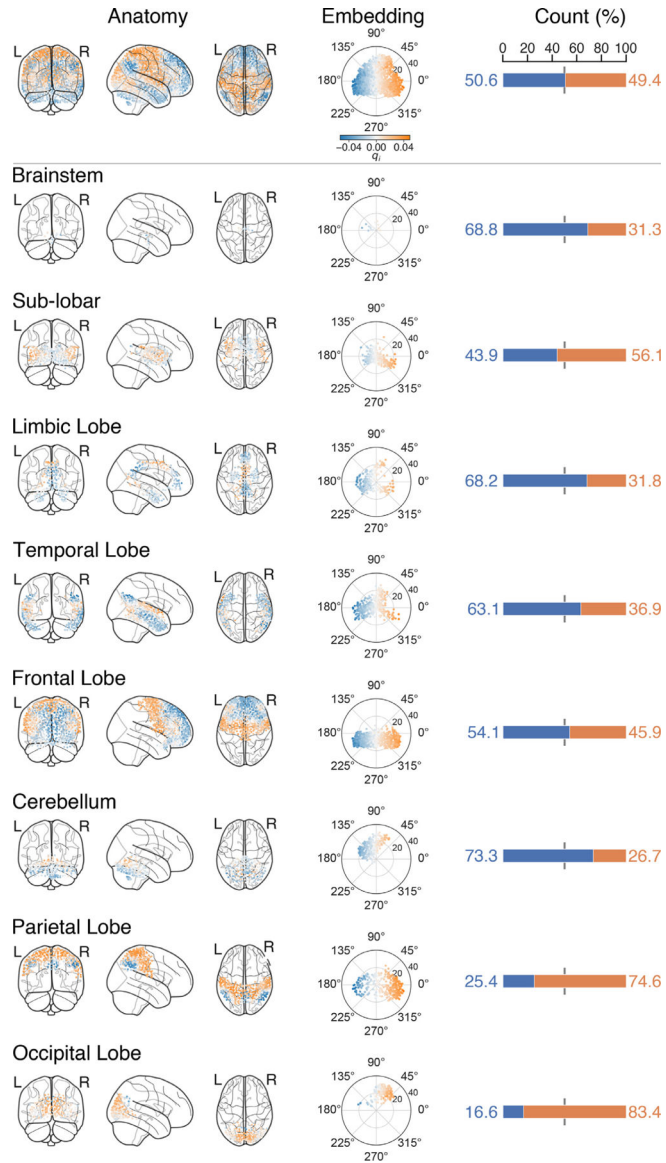


Fig. 2. Visualizing the principal rest2vec gradient. Anatomical and functional embedding of the Diez dataset faceted by anatomical lobe affiliation. (Top) Merged representations of all 2514 regions in anatomical embedding (columns 1–3), functional embedding (column 4), and the percentage of regions within each community for each lobe (column 5). (Bottom) Facet of data for each anatomical lobe. Rows are arranged from top to bottom in ascending order of median distance from the origin. Color indicates community affiliation. Vertical gray reference line for each stacked barplot indicates 50%.

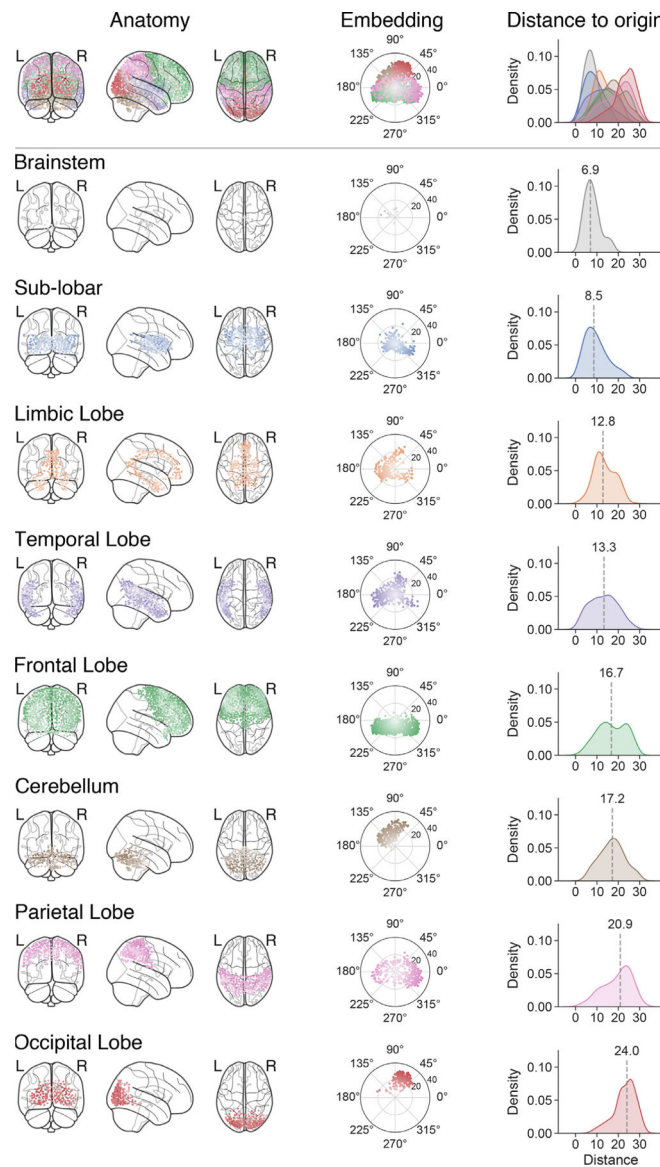


Fig. 3. Anatomical and functional embedding of the Diez dataset faceted by anatomical lobe affiliation and ranked by ascending distance to origin. (Top) Merged representations of all 2514 regions in the anatomical embedding (columns 1–3), functional embedding (column 4) and kernel density estimate of distance to origin for all regions within each lobe (column 5). (Bottom) Facet of data for each anatomical lobe. Rows are arranged from top to bottom in ascending order of median distance from the origin. Color indicates lobe affiliation. Higher saturation indicates increasing distance from the origin. Dashed gray lines in kernel density estimate plots indicate the median distance.

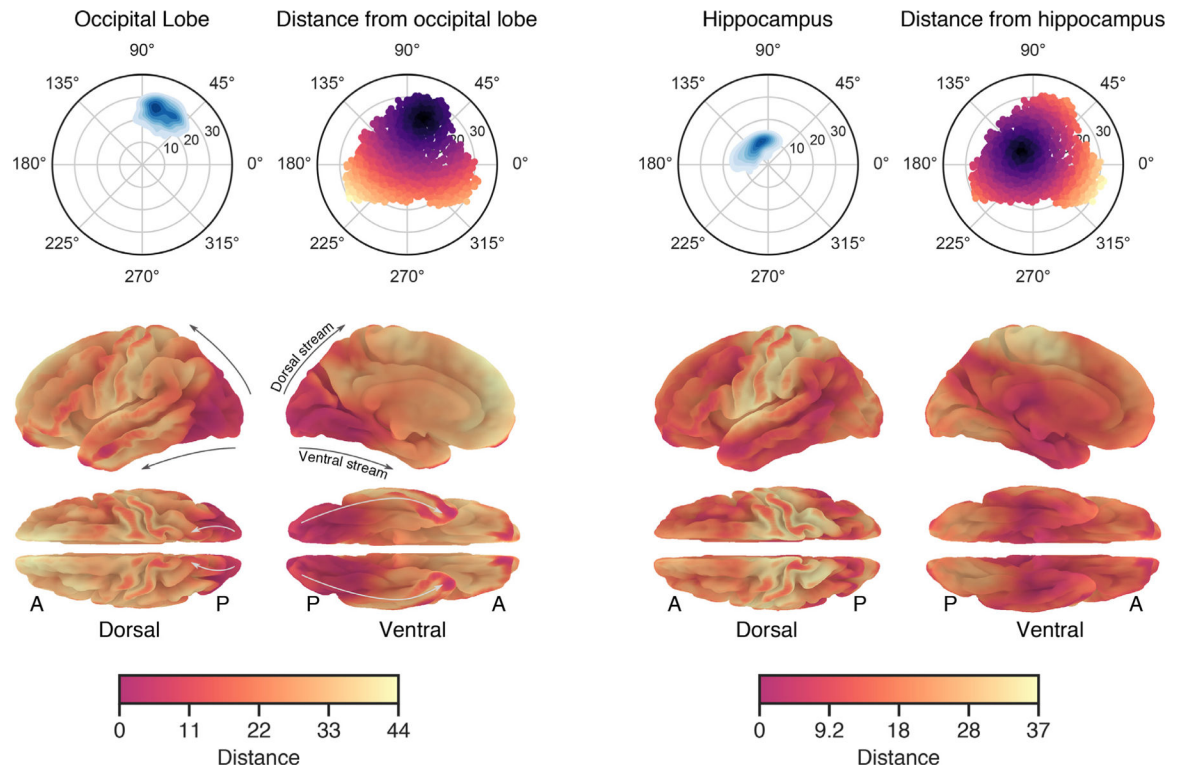


Fig. 4. Intrinsic functional distance mapping with rest2vec. (Left) Occipital lobe. (Right) Hippocampus. (Top) Kernel density estimate plots of the occipital lobe and hippocampus in the isomap embedding. The intrinsic functional distance was then computed from the mean embedding coordinate of either the occipital lobe or the hippocampus to all other regions in the brain. Darker color indicates the region is closer to the mean coordinate of the occipital lobe or hippocampus, respectively. (Bottom) Intrinsic functional distance gradient projected onto the Freesurfer pial surface template. A: anterior. P: posterior.

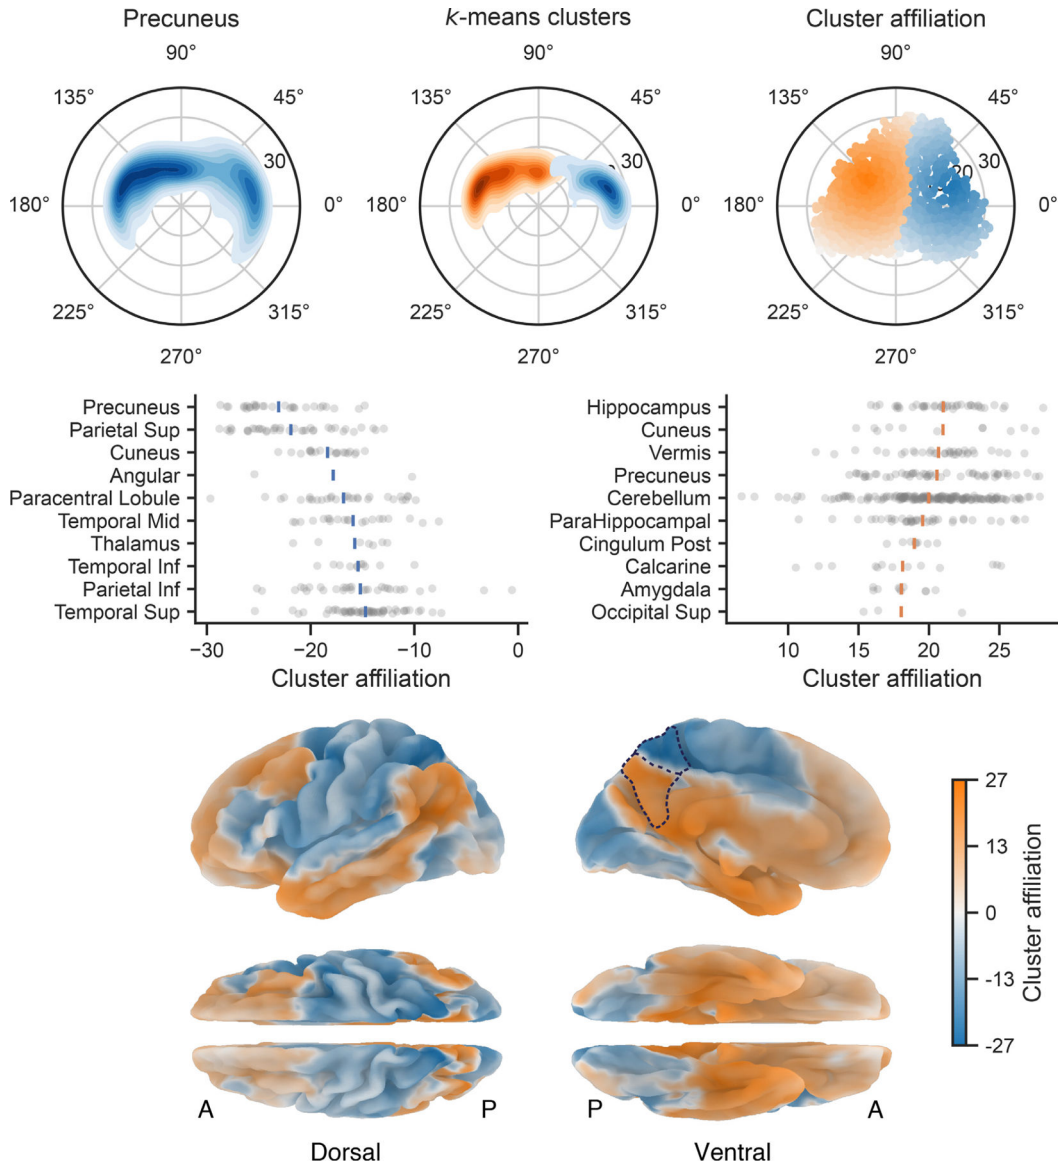



Fig. 5. Identifying subnetwork clusters within the precuneus using rest2vec. (Top, left) Kernel density estimate of the precuneus in the isomap embedding. (Top, middle) *k*-means clustering results are indicated in blue and orange. (Top, right) Cluster affiliations for all other regions based on their minimum intrinsic functional distance to their precuneus cluster centroids. Darker color indicates that region is closer to the centroid of its cluster. (Middle) Strip plot of the ten regions with the greatest mean affiliation for each cluster. Points represent individual ROI. Vertical bars indicate the mean. (Bottom) Brain surface map of cluster affiliations for the precuneus. The precuneus is outlined by a dashed line in the medial view. A: anterior. P: posterior.

F1000 dataset



Diez dataset


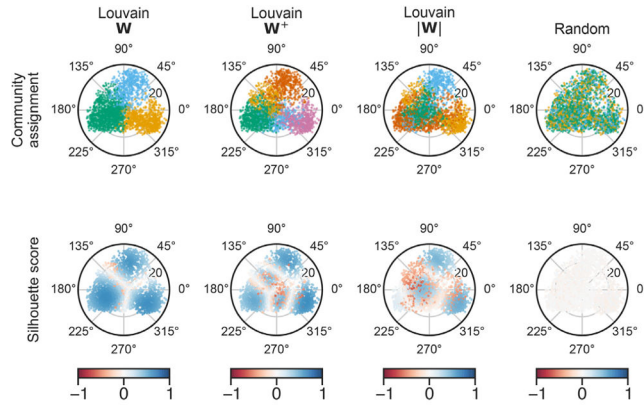


Fig. 6. Evaluating functional community mappings in the isomap embedding. (Top) For each dataset, the top row indicates the Louvain community assignment by distinct colors. The silhouette score for each ROI is shown below the respective community assignment plot. (Bottom) The overall silhouette coefficient for each connectome, computed across 100 iterations of the Louvain algorithm on each connectome. Bars represent the mean. Error bars indicate the standard deviation. Asterisks indicate Tukey’s *post hoc* pairwise comparison *p*-values < 0.001 .

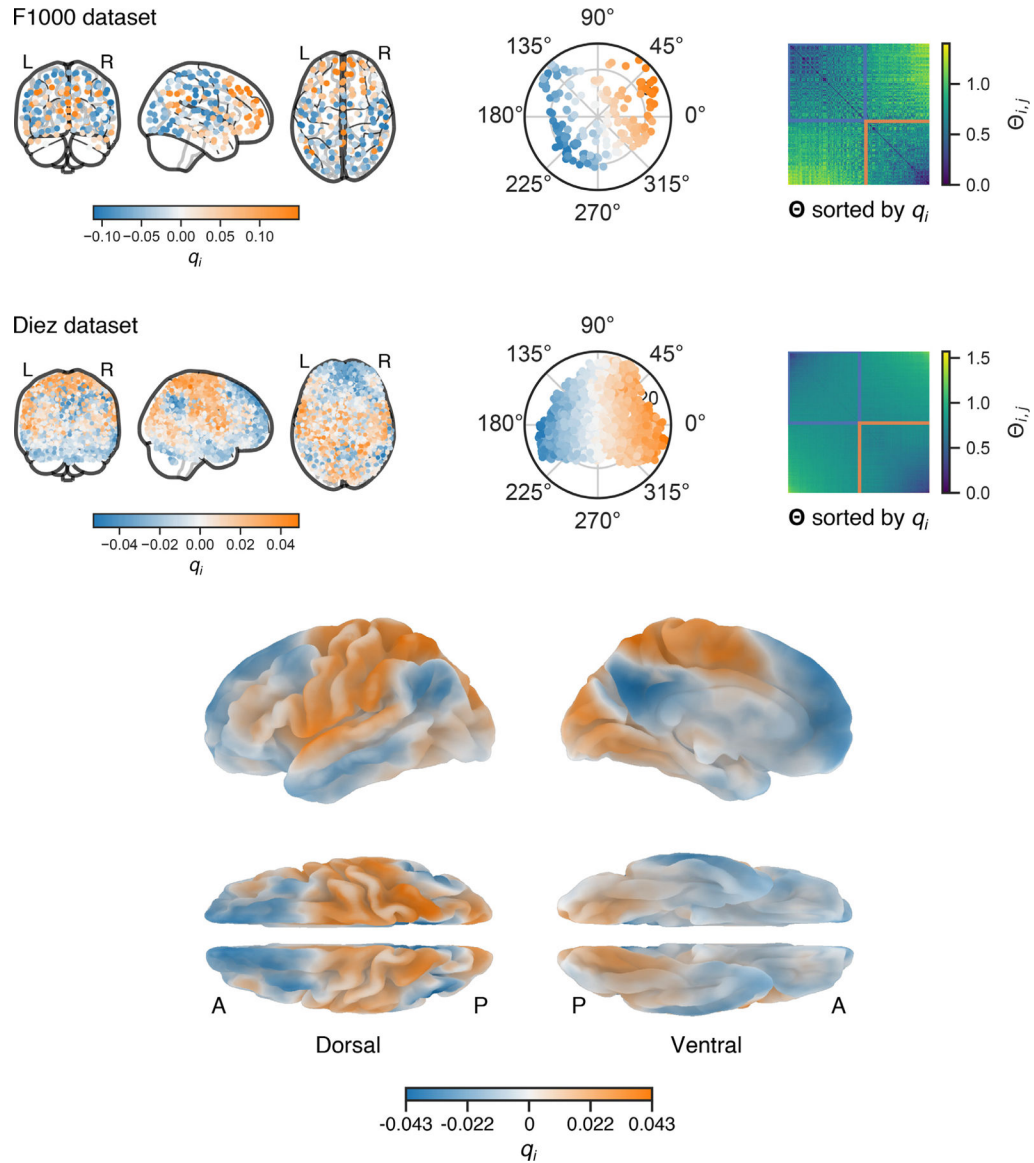


Fig. 7. Visualizing the primary eigenvector partition for the F1000 and Diez datasets. (Top) F1000 dataset. (Middle) Diez dataset. (Left column) Glass brain plots showing each ROI. Color indicates community affiliation based on the magnitude of the maximum eigenvector \mathbf{q} of the kernel similarity matrix $\tilde{\mathbf{K}}$. (Middle column) Visualization of the 2D embedding in polar coordinates. (Right column) PhASE matrix Θ sorted by the values of the maximum eigenvector \mathbf{q} . (Bottom) Brain surface map of the primary eigenvector partition gradient for the Diez dataset projected onto the Freesurfer pial surface template. Color indicates the interpolated value of q_i . A: anterior. P: posterior.

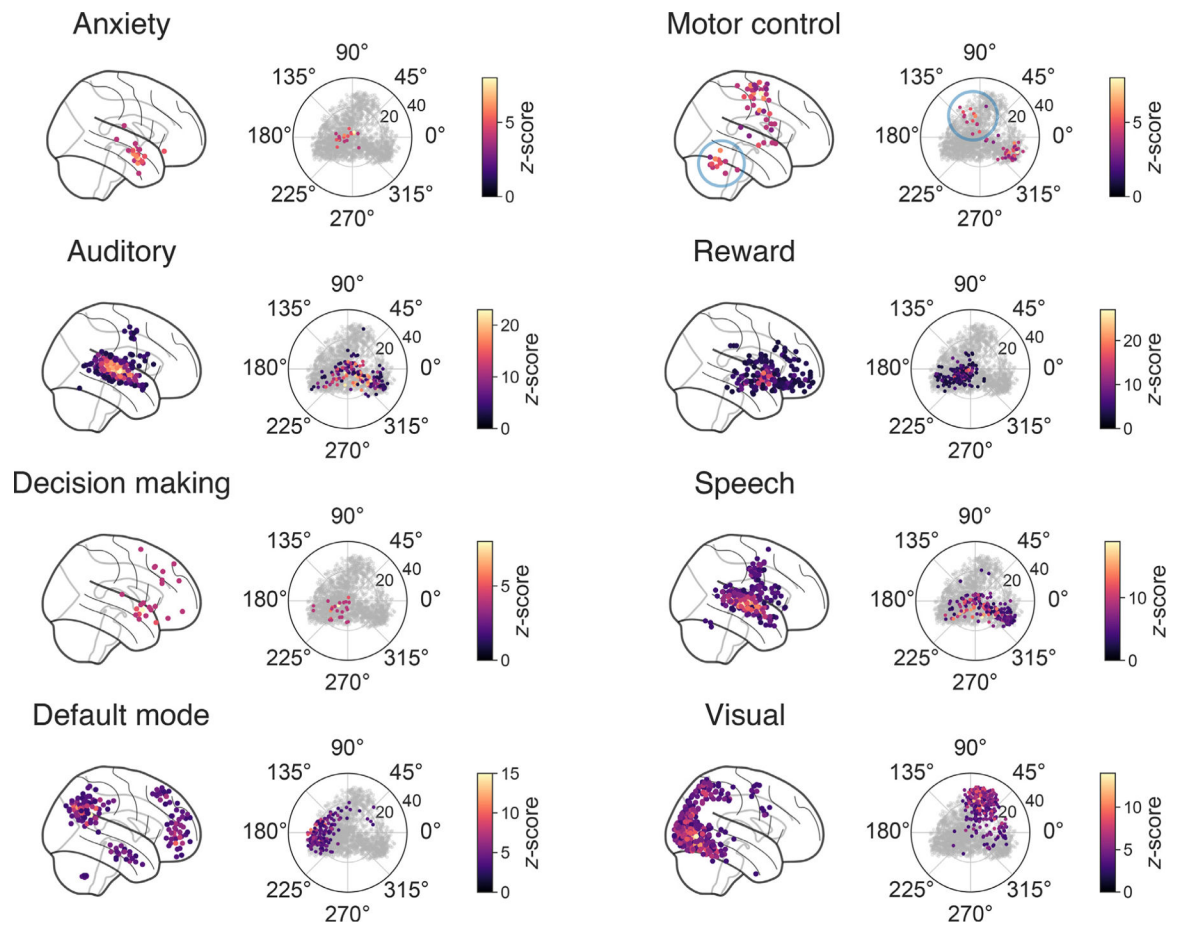


Fig. 8. Mapping Neurosynth meta-analysis terms to the rest2vec embedding. Glass brain and isomap embedding are shown for each Neurosynth term. Regions with z-scores > 0 are colored; other regions are displayed in gray. The Neurosynth data was acquired from the “association test” maps for each term.