

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

Zachery D. Morrissey^{1,2}, Liang Zhan⁵, Olusola Ajilore², and Alex D. Leow^{2,3,4}(✉)

¹Graduate Program in Neuroscience, University of Illinois at Chicago

²Dept. of Psychiatry, University of Illinois at Chicago

³Dept. of Bioengineering, University of Illinois at Chicago

⁴Dept. of Computer Science, University of Illinois at Chicago

⁵Dept. of Electrical and Computer Engineering, University of Pittsburgh

✉alexfeuillet@gmail.com

May 9 2020

1 Abstract

2 Resting-state functional magnetic resonance imaging (RS-FMRI) is widely used in
3 connectomics for studying the functional relationships between regions of the human
4 brain. RS-FMRI connectomics, however, has inherent analytical challenges, such as
5 accounting for negative correlations. In addition, functional relationships between brain
6 regions do not necessarily correspond to their anatomical distance, making the intrinsic
7 geometry of the functional connectome less well understood. Recent techniques in natural
8 language processing and machine learning, such as word2vec, have used embedding
9 methods to map high-dimensional data into meaningful vector spaces. Inspired by
10 this approach, we have developed a graph embedding pipeline, rest2vec, for studying
11 the intrinsic geometry of functional connectomes. We demonstrate how rest2vec uses
12 the phase angle spatial embedding (PHASE) method with dimensionality reduction
13 techniques to embed the functional connectome into lower dimensions. Rest2vec can
14 also be linked to the maximum mean discrepancy (MMD) metric to assign functional
15 modules of the connectome in a continuous manner, improving upon traditional binary
16 classification methods. Together, this allows for studying the functional connectome
17 such that the full range of correlative information is preserved and gives a more informed
18 understanding of the functional organization of the brain.

19 1 Introduction

20 Neuroimaging data acquired from magnetic resonance imaging (MRI) tend to be vast and
21 high-dimensional. In particular, resting-state functional MRI (RS-FMRI) produces temporal
22 snapshots of the brain's default activity in the absence of tasks, offering a window into the
23 functional macroscale organization of the brain. As computational tools have become more
24 widely available over the past two decades, researchers have applied graph theory-based
25 models to neuroimaging data to study the network properties of the brain, which has grown
26 into the field of connectomics [1]. In connectomics analyses, the brain can be represented as
27 an $N \times N$ matrix, where the rows and columns are the N brain regions of interest (ROI),
28 and the elements of the matrix represent some measure of connection between them (e.g.,
29 number of fibers, Pearson correlation of blood oxygenation level-dependent (BOLD) time
30 series). Given this volume of high-dimensional data, however, one quickly runs into the
31 "curse of dimensionality." Originally coined by Richard Bellman [2], the term refers to the
32 challenge of visualizing and analyzing high-dimensional data. Because the number of points
33 in a Cartesian space grows exponentially with increasing dimensions, high-dimensional spaces
34 become extremely sparse, an effect known as the "empty space phenomenon." Consequently,
35 this makes understanding the properties of these data more difficult, as metric comparisons
36 become less effective with increasing dimensionality [3].

37 There are a variety of dimensionality reduction techniques that address this problem.
38 These methods work by embedding a high-dimensional manifold, represented by the discrete
39 points of the data, into a lower dimension (e.g., two or three dimensions), which can then be
40 visualized. This process becomes complicated, however, if the manifold of the underlying
41 data is nonlinear, as is thought to be the case with neuroimaging data [4, 5, 6, 7]. The most
42 well-known example case of a nonlinear manifold is the 3D Swiss roll. Nonlinear dimensionality
43 reduction techniques, such as isometric mapping (isomap), solve the characteristic Swiss roll
44 problem by preserving the intrinsic geometry of nonlinear manifolds (i.e., unrolling the Swiss
45 roll) in lower-dimensional spaces [8, 4].

46 Negative correlations also remain a challenging factor in RS-FMRI connectomics, as they
47 are more difficult to interpret using network models. Simpler models generally either threshold
48 out or apply other transformations to negative correlations, such as taking the absolute value;
49 this process, however, likely removes substantive dynamics of brain connectivity [9]. Although
50 some analyses account for negative correlations, these often introduce additional parameters
51 that must be set to determine their relative contribution [9].

52 Previously, we introduced probability-associated community estimation (PACE) [10] and
53 phase angle spatial embedding (PHASE) [11] to address these challenges. These methods

54 take inspiration from the Ising model from statistical mechanics, where magnetic ions are
55 designated with in-phase or out-of-phase spin state configurations [12]. We adapted this
56 model to describe the phase relationship between regions of the brain. Here we propose a
57 novel graph embedding pipeline, rest2vec, that uses this phase angle representation with the
58 nonlinear dimensionality reduction method isomap to embed the functional connectome in
59 a lower-dimensional embedding based on its functional relationships. Doing so revealed a
60 spatial mapping of the functional organization of the brain based on its intrinsic geometry
61 when it is not constrained by neuroanatomy.

62 Additionally, we show this vectorized approach has implications for detecting functional
63 communities by linking rest2vec to the maximum mean discrepancy (MMD) metric. This
64 was originally developed by Gretton et al. [13] as a metric describing the distance between
65 probability distributions. Here, we treated the MMD as a modularity index, similar to Q -
66 based maximization methods [14], such that, when maximized, it detects the sets of brain
67 regions with the most dissimilar functional connectivity. By reformulating this connectome
68 modularity problem in a probabilistic sense, we are able to generate community assignment
69 values for each region, as opposed to a binary classification. Together, rest2vec uses nonlinear
70 dimensionality reduction and manifold learning techniques to represent the functional connec-
71 tome in its intrinsic geometry independent of neuroanatomy to improve our understanding of
72 the macroscale organization of the brain.

73 2 Methods

74 2.1 Dataset description

75 Two independent and publicly available RS-FMRI connectome datasets composed of healthy
76 subjects were used: one from the Functional 1000 (F1000) Connectomes Project [15] with
77 177 regions of interest (ROI) available through the USC Multimodal Connectivity Database
78 (<http://umcd.humanconnectomeproject.org/umcd/default/index>), and one by Diez et
79 al. [16] with 2514 ROI available through the NeuroImaging Tools & Resources Collaboratory
80 (NITRC) (<https://www.nitrc.org/projects/biocatlas/>). These are referred to as
81 the “F1000” and “Diez” datasets hereafter. The average difference in age between male
82 ($N = 426, M \pm SD = 28.7 \pm 12.7$) and female ($N = 560, M \pm SD = 27.9 \pm 12.7$) subjects
83 in the F1000 dataset was 0.83 years and was not statistically significant ($t(984) = 1.025$,
84 $p = 0.306$). The Diez dataset has 12 subjects (6 male) with a mean age of 33.5 ± 8.7 years;
85 no individual subject ages were reported. The reader can consult the references for details
86 regarding image acquisition parameters and preprocessing. For computational and network

87 analyses, Python version 3.7.3 scientific computing libraries from the Anaconda distribution
88 were used [17, 18, 19, 20, 21, 22, 23].

89 2.2 rest2vec

90 The pipeline for rest2vec is shown in [Figure 1](#). Rest2vec aims to create a graph embedding
91 of RS-FMRI connectomes by transforming positive and negative edges into N -dimensional
92 phase angle vectors that can then be represented in a low-dimensional embedding using
93 nonlinear dimensionality reduction. In brief, we first computed the probability of observing a
94 negative edge between all pairs of regions across all subjects to form the probability matrix
95 \mathbf{P}^- . This probability is then used to determine the phase angle $\Theta_{i,j}$ between regions to
96 create the phase angle spatial embedding (PHASE) matrix Θ . This process embeds the phase
97 relationship between all regions in the connectome in an N -dimensional Euclidean space and
98 transforms the values between 0 (fully co-activating) and $\pi/2$ (fully anti-activating). The
99 intrinsic functional embedding of the connectome was then visualized in two dimensions using
100 the nonlinear dimensionality reduction method isomap [8]. Finally, we use kernel functions
101 to link rest2vec to the maximum mean discrepancy MMD metric [13] to demonstrate how
102 rest2vec can be used to study functional connectome modularity. The representative matrices
103 for each step are displayed in [Figure 2](#).

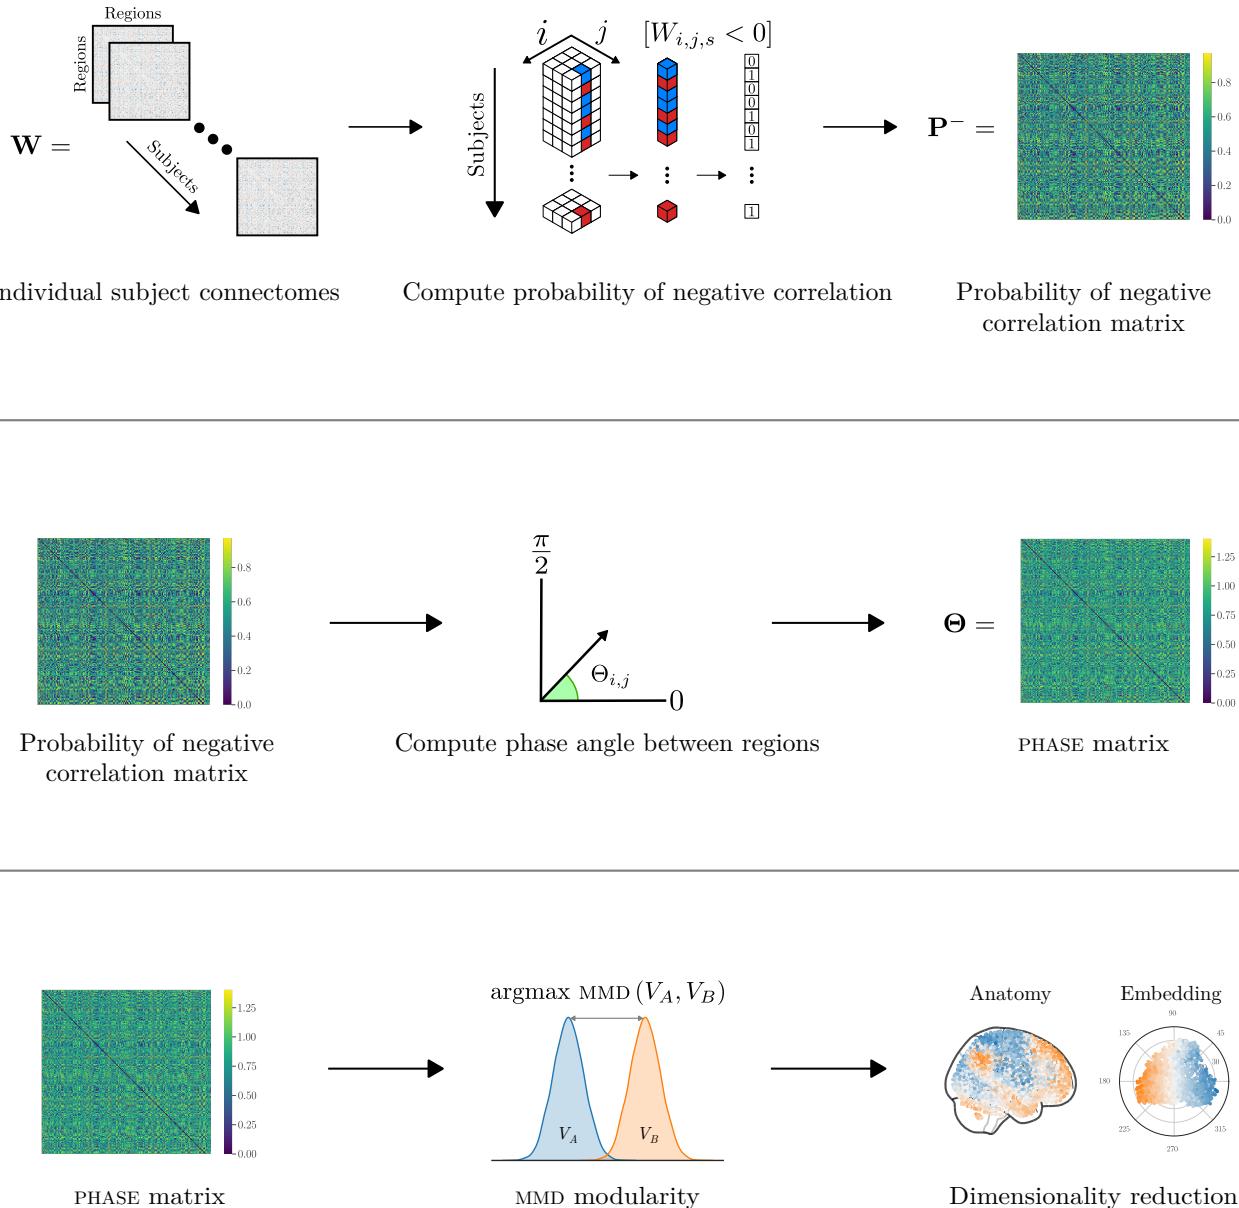


Figure 1: rest2vec processing pipeline. (Top) The frequency of observing a negative edge between regions i and j across all subjects in the $N \times N \times S$ array \mathbf{W} of RS-FMRI connectomes is computed to form the probability of negative correlation matrix \mathbf{P}^- . (Middle) The phase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

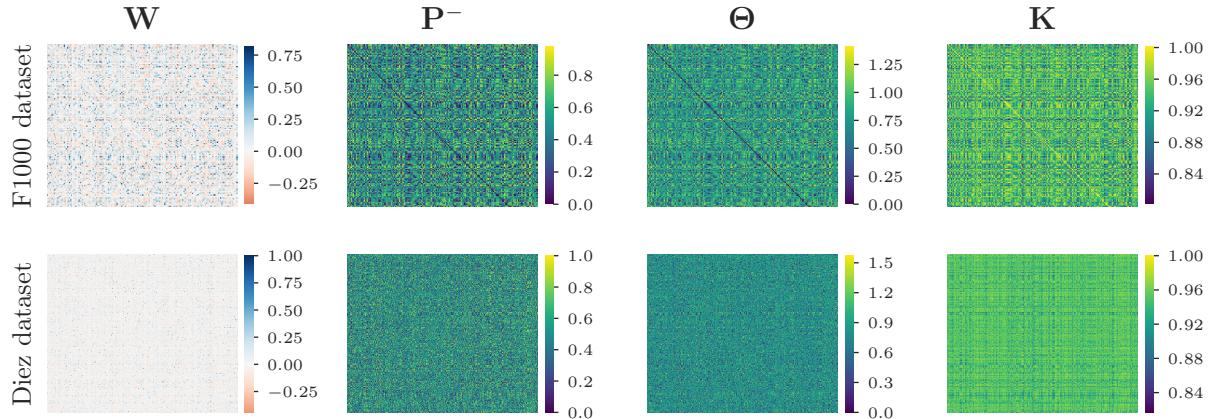


Figure 2: Representative matrices for processing steps of rest2vec pipeline for each dataset. Pearson correlation matrix \mathbf{W} , negative probability matrix \mathbf{P}^- , phase angle spatial embedding (PHASE) matrix Θ , and kernel similarity matrix \mathbf{K} are displayed.

104 2.2.1 Phase angle spatial embedding (PHASE)

105 A functional connectome derived from RS-fMRI is defined as an undirected graph $G(V, E)$,
 106 composed of a set of vertices V , i.e., brain regions of interest (ROI), and signed, weighted
 107 edges E describing the measure of connectivity between them based on their BOLD response
 108 time series. Typically, some measure of correlation, e.g., Pearson correlation, between BOLD
 109 time series is used to describe the functional connectivity between ROI.

110 Previously, we introduced probability associated community estimation (PACE) [10], and
 111 phase angle spatial embedding (PHASE) [11] for encoding resting-state fMRI connectomes
 112 based on the phase relationship between brain regions [11] to account for negative correlations
 113 in functional connectomes. We begin by briefly summarizing these procedures in the context
 114 of rest2vec.

115 Let \mathbf{W} be an $N \times N \times S$ array (i.e., a tensor) composed of $N \times N$ weighted, signed
 116 functional connectomes for N regions and S subjects. Given some weight of functional
 117 coupling between regions i and j (e.g., Pearson correlation), we define the probability of
 118 negative correlation matrix \mathbf{P}^- where each element $P_{i,j}^-$ is the probability of observing a
 119 negative edge between i and j defined as

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

120 where $W_{i,j,s}$ is the edge between regions i and j for the s th subject, and the Iverson
 121 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
 122 also follows naturally that

$$P_{i,j}^- + P_{i,j}^+ = 1. \quad (2)$$

123 One advantage of this procedure is that the probability measure defined in [Equation 1](#)
 124 can be defined by the user for their specific context. By taking advantage of this relationship,
 125 we then define the phase angle spatial embedding (PHASE) matrix Θ , where the phase angle
 126 $\Theta_{i,j}$ between regions i and j is defined as

$$\Theta_{i,j} = \arctan \sqrt{\frac{P_{i,j}^-}{P_{i,j}^+}}. \quad (3)$$

127 Thus $\Theta_{i,j} \in [0, \pi/2]$, where 0 represents a fully in-phase (co-activating) relationship
 128 and $\pi/2$ represents a fully out-of-phase (anti-activating) relationship. Each column of Θ
 129 is a vector embedding each region in an N -dimensional Euclidean space such that $\Theta_{:,i} =$
 130 $[\Theta_{i,1} \ \Theta_{i,2} \ \dots \ \Theta_{i,N}]^\top \in [0, \pi/2]^N$.

131 2.2.2 Relation of PhASE to the maximum mean discrepancy

132 Here we describe how PHASE can be linked to the maximum mean discrepancy (MMD)
 133 developed by Gretton et al. [\[13\]](#) to address the connectome modularity problem. Following
 134 the formulation defined in [\[13\]](#), consider the random variables x and y defined on a metric
 135 space \mathcal{X} equipped with the metric d , with the corresponding Borel probabilities p and q (i.e.,
 136 $x \sim p$ and $y \sim q$). Given observations $X := \{x_1, \dots, x_m\}$ and $Y := \{y_1, \dots, y_n\}$ drawn from
 137 the probability distributions p and q , $p = q$ if and only if $\mathbf{E}_x[f(x)] = \mathbf{E}_y[f(y)] \ \forall f \in C(\mathcal{X})$,
 138 where $C(\mathcal{X})$ is the space of bounded continuous functions on \mathcal{X} . Next, given a class of
 139 functions \mathcal{F} such that $f : \mathcal{X} \rightarrow \mathbb{R}$, the maximum mean discrepancy (MMD) between p and q
 140 with respect to \mathcal{F} is defined as

$$\text{MMD}[\mathcal{F}, p, q] := \sup_{f \in \mathcal{F}} (\mathbf{E}_x[f(x)] - \mathbf{E}_y[f(y)]). \quad (4)$$

141 This can be empirically estimated given X and Y as

$$\text{MMD}[\mathcal{F}, X, Y] := \sup_{f \in \mathcal{F}} \left(\frac{1}{m} \sum_{i=1}^m f(x_i) - \frac{1}{n} \sum_{i=1}^n f(y_i) \right), \quad (5)$$

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

144 To apply these definitions in the context of connectomics, we use the same definitions of
 145 x, y, p, q, X , and Y defined above to assign each region to one of the two distributions p or q .

146 Under the working assumption that the distributions of functional modules in the connectome
 147 are far apart (i.e., their within-module connections are greater than their between-module
 148 connections [24]), we thus seek to discover the arrangement of regions such that the MMD
 149 between them is maximized.

150 Using a reproducible kernel Hilbert space (RKHS), the squared form of [Equation 5](#) can be
 151 evaluated using kernel functions as

$$\begin{aligned} \text{MMD}^2[\mathcal{F}, X, Y] := & \frac{1}{m(m-1)} \sum_{i=1}^m \sum_{j \neq i}^m k(x_i, x_j) + \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j \neq i}^n k(y_i, y_j) \\ & - \frac{2}{mn} \sum_{i=1}^m \sum_{j=1}^n k(x_i, y_j). \end{aligned} \quad (6)$$

152 From [Equation 6](#), kernel functions can be used, in our case, to compute the kernel matrix
 153 \mathbf{K} where the similarity $K_{i,j}$ between regions i and j , in the case of the radial basis function
 154 (RBF) kernel k_{RBF} , is given by

$$K_{i,j} = k_{\text{RBF}}(\Theta_{i,:}, \Theta_{j,:}) = \exp \left(-\sigma \sum_{\ell=1}^N |\Theta_{i,\ell} - \Theta_{j,\ell}|^2 \right), \quad (7)$$

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

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

$$K_{i,j} = k_{\text{cos}}(\Theta_{i,:}, \Theta_{j,:}) = \frac{1}{N} \sum_{\ell=1}^N \cos(\Theta_{i,\ell} - \Theta_{j,\ell}), \quad (8)$$

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

162 **2.2.3 Using maximum mean discrepancy to address the connectome modularity
 163 problem**

164 Following the kernel definitions above, and the equation as described by [13] (with a modified
 165 notation for our purposes), let the maximum mean discrepancy (MMD) between two modules
 166 V_A and V_B be defined as

$$\text{MMD}(V_A, V_B)^2 = \frac{1}{m^2} \sum_{i,j \in V_A}^m K_{i,j} - \frac{2}{mn} \sum_{\substack{i \in V_A \\ j \in V_B}}^{m,n} K_{i,j} + \frac{1}{n^2} \sum_{i,j \in V_B}^n K_{i,j}, \quad (9)$$

167 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
168 to equal j .

169 We seek to find a partition between V_A and V_B such that [Equation 9](#) is maximized. First,
170 we can rewrite $\text{MMD}(V_A, V_B)^2 = \mathbf{y}^\top \mathbf{K} \mathbf{y}$ for $\mathbf{y} \in \mathbb{R}^{N \times 1}$ and $\mathbf{K} \in \mathbb{R}^{N \times N}$, where

$$y_i = \begin{cases} \frac{1}{m} & \text{if } i \in V_A \\ -\frac{1}{n} & \text{if } i \in V_B \end{cases}. \quad (10)$$

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

$$\text{Modularity}(V) = \underset{\substack{V_A, V_B \\ V_A \cup V_B = V \\ V_A \cap V_B = \emptyset}}{\text{argmax}} \text{MMD}(V_A, V_B)^2. \quad (11)$$

172 This maximization problem can be approximated in a simplified way by relaxing [Equation 11](#) to a Rayleigh quotient maximization problem. Letting \mathbf{y} be defined as above, where
173

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

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

$$\mathbf{v} = \left(\frac{mn}{N} \right)^{1/2} \mathbf{y}, \quad (13)$$

175 and $\|\mathbf{v}\|^2 = 1$, $\mathbf{v}^\top \mathbf{1} = 0$, where $\mathbf{1} = [1 \cdots 1]^\top$, $\mathbf{1} \in \mathbb{R}^{N \times 1}$. Then we can rewrite [Equation 11](#)
176 in terms of \mathbf{v} to define the partition that maximizes $\text{MMD}(V_A, V_B)^2$ as

$$\text{Modularity}(V) = \underset{\substack{\|\mathbf{v}\|=1 \\ \mathbf{v}^\top \mathbf{1}=0}}{\text{argmax}} \frac{N}{mn} \mathbf{v}^\top \mathbf{K} \mathbf{v}. \quad (14)$$

177 To compute $\text{MMD}(V_A, V_B)^2$ in [Equation 14](#) requires *a priori* knowledge of m and n .
178 Assuming that N is large and that the two communities V_A and V_B are approximately the
179 same size such that $|m - n| \in o(N)$, the normalization factor in [Equation 14](#) can be simplified
180 to

$$\text{Modularity}(V) \approx \frac{4}{N} \underset{\substack{\|\mathbf{v}\|=1 \\ \mathbf{v}^\top \mathbf{1}=0}}{\operatorname{argmax}} \mathbf{v}^\top \mathbf{K} \mathbf{v}. \quad (15)$$

181 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
 182 values to taking any real values such that $\mathbf{v}^* \in \mathbb{R}^N$. These relaxed constraints allow us to
 183 conveniently reframe [Equation 14](#) as a Rayleigh quotient maximization problem. We account
 184 for arbitrary origin for the Rayleigh quotient maximization by centering the kernel similarity
 185 matrix \mathbf{K} to $\tilde{\mathbf{K}} = \mathbf{C}_N \mathbf{K} \mathbf{C}_N$, where the centering matrix $\mathbf{C}_N = \mathbf{I}_N - \frac{1}{N} \mathbf{J}_N$, $\mathbf{C}_N \in \mathbb{R}^{N \times N}$,
 186 $\mathbf{I}_N \in \mathbb{R}^{N \times N}$ is the identity matrix, and $\mathbf{J}_N \in \mathbb{R}^{N \times N}$ is the ones matrix (i.e., $\mathbf{1}\mathbf{1}^\top$).

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

$$\text{Modularity}(V) \approx \frac{4}{N} \underset{\substack{\mathbf{v}^* \neq 0 \\ \mathbf{v}^* \in \mathbb{R}^N}}{\operatorname{argmax}} \frac{\mathbf{v}^{*\top} \tilde{\mathbf{K}} \mathbf{v}^*}{\mathbf{v}^{*\top} \mathbf{v}}. \quad (16)$$

190 We can then compute the mapping vector \mathbf{v}^* that maximizes the Rayleigh quotient by
 191 computing the eigenvector \mathbf{q} of $\tilde{\mathbf{K}}$ corresponding to the largest eigenvalue λ_{\max} of $\tilde{\mathbf{K}}$. Similar
 192 to the Fiedler vector in spectral clustering methods [\[25\]](#), the elements of \mathbf{q} assign both
 193 community affiliation based on its sign (+ or -) as well as magnitude. Further, \mathbf{q} can be
 194 binarized to determine discrete community labels for each region as

$$\mathbf{v}^* = \begin{cases} i \in V_A & \text{if } q_i \geq 0 \\ i \in V_B & \text{if } q_i < 0 \end{cases}, \quad \forall i = 1, 2, 3, \dots, N. \quad (17)$$

195 2.2.4 Nonlinear dimensionality reduction

196 Isomap [\[8\]](#) was used to reduce the PHASE matrix $\Theta \in \mathbb{R}^{N \times N}$ to a d -dimensional embedding
 197 $\mathbf{Y} \in \mathbb{R}^{N \times d}$, where $d < N$. Isomap is advantageous for this procedure as it is a nonlinear
 198 technique, using methods such as Dijkstra's algorithm [\[26\]](#) to compute the geodesic distances
 199 between vertices in high-dimensional space. By doing so, isomap addresses the Swiss roll
 200 problem faced by traditional linear methods such as PCA and MDS [\[8\]](#). In our case, we used
 201 $k = 12$ and $k = 50$ nearest neighbors, for the F1000 and Diez datasets, respectively, to reduce
 202 to three dimensions using the `Isomap` implementation in the Scikit-learn version 0.21.3 library
 203 [\[22\]](#). Because the isomap procedure centers data about the origin, and by [Equation 3](#) the
 204 phase angle between perfectly in-phase regions is zero, we analyzed each region's Euclidean
 205 distance to the origin in this space to observe how the phase relationship between regions

206 is preserved with respect to its low-dimensional embedding. After generating the isomap
207 embedding, the distance D_i to the origin of the isomap space $[0 \cdots 0] \in \mathbb{R}^{1 \times d}$ for the i th
208 region was calculated using the Euclidean distance

$$D_i = \|Y_{i,:}\| = \sqrt{Y_{i,1}^2 + Y_{i,2}^2 + Y_{i,3}^2 + \cdots + Y_{i,d}^2}, \quad \forall i = 1, 2, 3, \dots, N, \quad (18)$$

209 where N is equal to the number of regions. Because of its natural representation for
210 distance to the origin, the first two dimensions were transformed to polar coordinates of
211 radius r and angle θ using the polar transformation

$$r = \sqrt{x^2 + y^2} \quad (19)$$

$$\theta = \text{atan2}(y, x) \quad (20)$$

212 to visualize the functional embedding space.

213 2.3 Analyses

214 2.3.1 k -means clustering

215 k -means clustering was used to formally classify clusters for regions (such as the precuneus)
216 that had heterogeneous mappings in the isomap embedding. The k -means clustering algorithm
217 was performed using the Scikit-learn implementation [22] for $k = 2$ clusters in the isomap
218 embedding. The same seed value was used to ensure reproducible results.

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

$$a_i = \begin{cases} \max(d_{C_0}) - d_{C_0,i} & \text{if } i \in C_0 \\ d_{C_1,i} - \max(d_{C_1}) & \text{if } i \in C_1 \end{cases}, \quad (21)$$

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

227 2.3.2 Statistics

228 The StatsModels library version 0.10.1 for Python [27] was used for statistical analyses.
229 Student's independent *t*-test was used to test if there were any differences in age between
230 male and female subjects for the F1000 dataset. The ordinary least squares (OLS) method
231 was used to fit the parameters for the linear regression between isomap distance to origin
232 and phase angle.

233 2.3.3 Visualization

234 Graphics were drawn using the Matplotlib version 3.1.1 [20] and Seaborn version 0.9.0 [21]
235 libraries using Python version 3.7.3 from the Anaconda distribution [17]. Glass brain figures
236 were visualized using the `plot_connectome` function from the Nilearn version 0.6.2 library
237 [28]. Inkscape version 0.92 was used for final arrangement of some figures [29].

238 Brain surface plots were created by representing the $N \times 4$ array, consisting of the MNI
239 (x, y, z)-coordinates for all N regions, and the $N \times 1$ vector containing the data value associated
240 with each region, as a 3D volume. For brain distance maps, the intrinsic functional distance
241 vector was made by computing the Euclidean distance between the mean (x, y)-coordinates
242 of the anatomical region in the isomap embedding and all other regions. For regions that
243 had heterogeneous mapping (i.e., multiple clusters) in the isomap space, *k*-means clustering
244 was performed to calculate cluster affiliations for each region as described in §2.3.1.

245 The 3D volume containing the original data was then interpolated using a linear grid
246 interpolation and registered to the MNI template volume with 12 degrees of freedom using
247 the FLIRT tool in the FSL [30] interface from the Nipype version 1.3.0-rc1 library [31].
248 The interpolated 3D volume was mapped to the Freesurfer pial surface template using the
249 `vol_to_surf` function from the Nilearn library. The surface data was then visualized using
250 the `plot_surf_stat_map` function from the Nilearn library.

251 2.3.4 Code

252 All code used to produce the results and figures is available online via GitHub (<https://github.com/zmorrissey>) and our laboratory website (<http://brain.uic.edu/>).

254 3 Results

255 3.1 Distance in the lower-dimensional embedding preserves phase 256 angle relationships

257 After applying the rest2vec pipeline to the F1000 and Diez datasets, we sought to assess
258 how a region's lower-dimensional isomap embedding related to its PHASE vector. From
259 [Equation 3](#), lower values of $\Theta_{i,j}$ indicate a more in-phase relationship between regions. Thus we
260 hypothesized that more in-phase regions would be embedded closer to the origin of the isomap
261 space, whereas more out-of-phase regions would be embedded further from the origin. The 2-
262 norm of each N -dimensional vector of the PHASE matrix $\|\Theta_{i,:}\|$ was used as a summary measure
263 of each region's overall phase value. For each dataset, there was a statistically significant
264 positive correlation between each region's $\|\Theta_{i,:}\|$ and its distance from the origin of the 3D
265 isomap embedding (F1000 dataset: $F(1, 175) = 200.7$, $R^2 = 0.534$, $r = 0.731$, $p < 0.0001$;
266 Diez dataset: $F(1, 2512) = 533.4$, $R^2 = 0.175$, $r = 0.418$, $p < 0.0001$) ([Figure 3](#), left). This
267 pattern can be seen when the rows and columns of the PHASE matrix are sorted by ascending
268 $\|\Theta_{i,:}\|$ values, in particular for the coarser parcellation from the F1000 dataset ([Figure 3](#),
269 right). Together this suggests that regions mapped closer to the origin were more in-phase
270 with other regions, whereas more out-of-phase regions were mapped further from the origin.

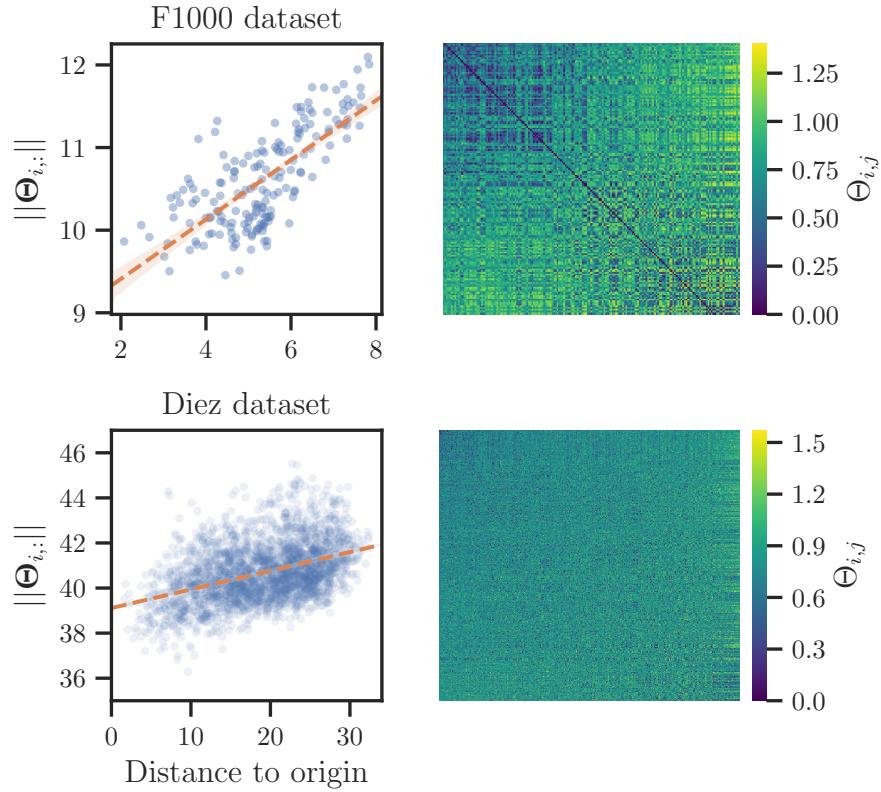


Figure 3: Relationship between phase angle and isomap embedding distance. (Left) Correlation between the 2-norm of each region's phase angle vector $\|\Theta_{i,:}\|$ and its distance to the origin of the 3D isomap embedding. Dashed orange line represents the best fit of the linear model. Shaded region around line represents the 95% confidence interval of the model. F1000 dataset: $r = 0.731$; Diez dataset: $r = 0.418$. (Right) The PHASE matrix Θ with its rows and columns sorted in ascending order by $\|\Theta_{i,:}\|$ (i.e., lowest values correspond to upper left, highest values to lower right).

271 To examine this relationship further, we faceted the anatomical and functional embeddings

272 by anatomical lobe affiliation ranked by ascending distance to the origin (Figure 4). Notably,

273 the brainstem displayed the most centrally-embedded regions (median distance = 6.9), followed

274 by (in ascending order): sub-lobar, limbic lobe, temporal lobe, frontal lobe, cerebellum, parietal

275 lobe, and occipital lobe regions. At the other extreme, the occipital lobe displayed the most

276 distant and densely clustered representation in the embedding space (median distance = 24).

277 Examination of the phase angle vectors for occipital lobe regions revealed highly in-phase

278 relationships within the occipital lobe, while regions outside the occipital lobe were mostly

279 out-of-phase (SI Figure 15). Since the occipital lobe and large portions of the parietal lobe

280 (e.g., motor cortices), and cerebellum are mapped further in the periphery, this suggests

281 that regions involved in primary sensory processing are mapped further in the periphery,

282 while regions such as the brainstem, thalamus, and heteromodal areas have more in-phase

283 relationships and are mapped closer to the origin.

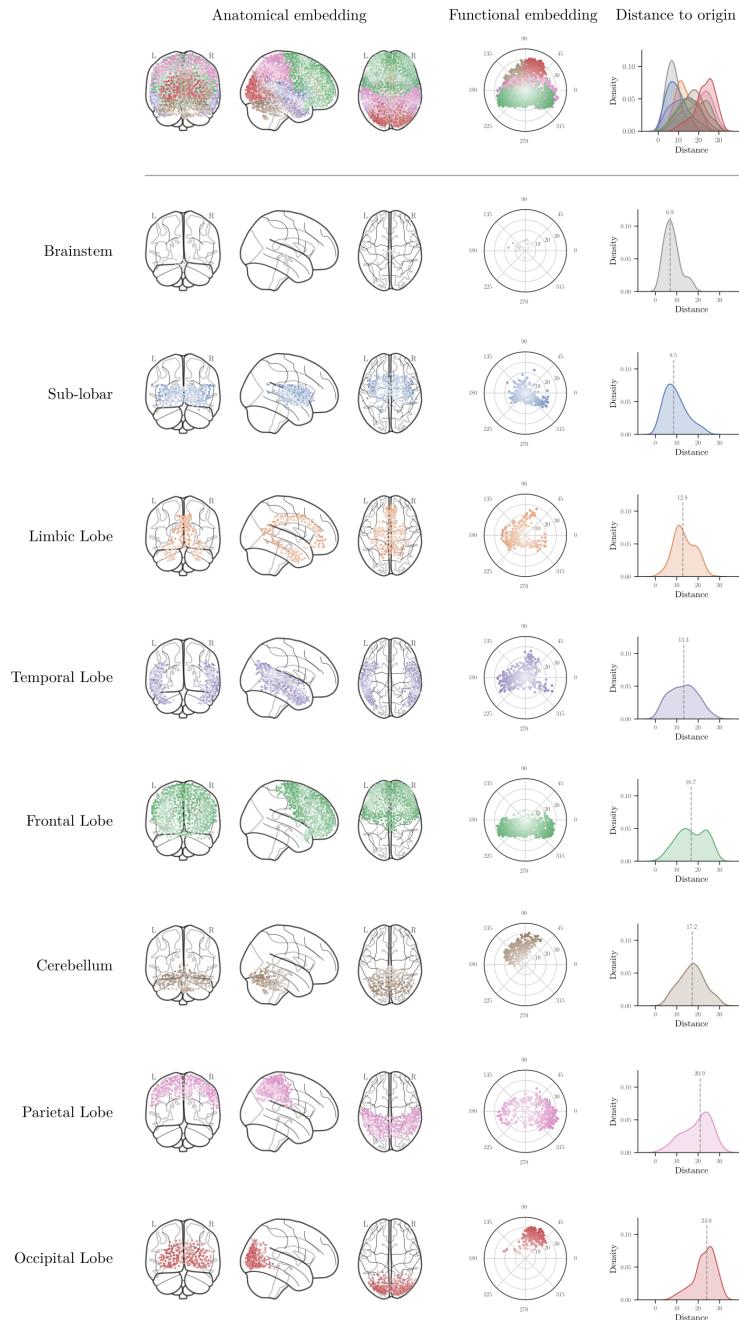


Figure 4: Anatomical and functional embedding of the Diez dataset faceted by anatomical lobe affiliation and ranked by ascending distance to origin. (Top) Merged representations of all 2514 regions in the anatomical embedding (columns 1-3), functional embedding (column 4) and kernel density estimate 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 from top to bottom. Color indicates lobe affiliation. Higher saturation indicates increasing distance from the origin. Dashed gray lines in kernel density estimate plots indicate the median distance.

284 3.2 Intrinsic functional distance can detect biologically-relevant 285 connectivity gradients

286 Given that the distance to the origin of the isomap embedding preserved phase coupling
287 characteristics across anatomical regions, we next asked if the intrinsic functional distance
288 between regions in this space could reveal biologically-relevant connectivity patterns. When
289 the intrinsic functional distance to the occipital lobe is mapped as a color gradient on the
290 brain surface, the dorsal and ventral visual streams [32, 33, 34] become apparent (Figure 5),
291 consistent with the hypothesis that distance in this embedding space preserves functionally
292 relevant information. In contrast, the hippocampus also has a relatively homogeneous cluster
293 in the isomap embedding, but has a much more distributed surface map gradient to regions
294 of the default mode network (DMN), such as the precuneus, prefrontal cortex, thalamus, and
295 inferior parietal lobule (Figure 6).

296 While certain anatomical regions showed a relatively homogeneous clustering in the isomap
297 embedding, such as the occipital lobe, others showed heterogeneous clustering patterns. Thus
298 we hypothesized that rest2vec could be used to identify functional subnetworks within
299 individual regions based on their clustering within the isomap embedding. As a test case, we
300 examined the isomap embedding pattern for the precuneus, which is known to participate
301 in different networks across its dorsal-anterior/ventral-posterior axes [35, 36]. The bivariate
302 kernel density estimate plot of the precuneus ROI in the Diez dataset appeared to indicate
303 two predominant clusters, which were formally assigned using k -means clustering (Figure 7,
304 top). A larger cluster was made that included all other regions in the Diez dataset by
305 assigning regions to the precuneus cluster they were closer to. We then measured the
306 intrinsic functional distance between each region to its precuneus cluster centroid to assign
307 an affiliation value to each region (Figure 7, top right). The brain surface map projection
308 of these data demarcated these two cluster centroids into the dorsal-anterior precuneus
309 and the ventral-posterior precuneus (Figure 7, bottom). The dorsal-anterior cluster of the
310 precuneus was most strongly affiliated with the occipital and superior parietal regions, as well
311 as the paracentral lobule, middle and superior temporal cortices, and thalamus (Figure 7,
312 middle). The ventral-posterior cluster of the precuneus was most strongly affiliated with
313 the hippocampus, cuneus, cerebellum, parahippocampal cortex, posterior cingulate cortex,
314 calcarine cortex, amygdala, and superior occipital cortices. These results suggest that rest2vec
315 can identify distinct functional networks within individual regions.

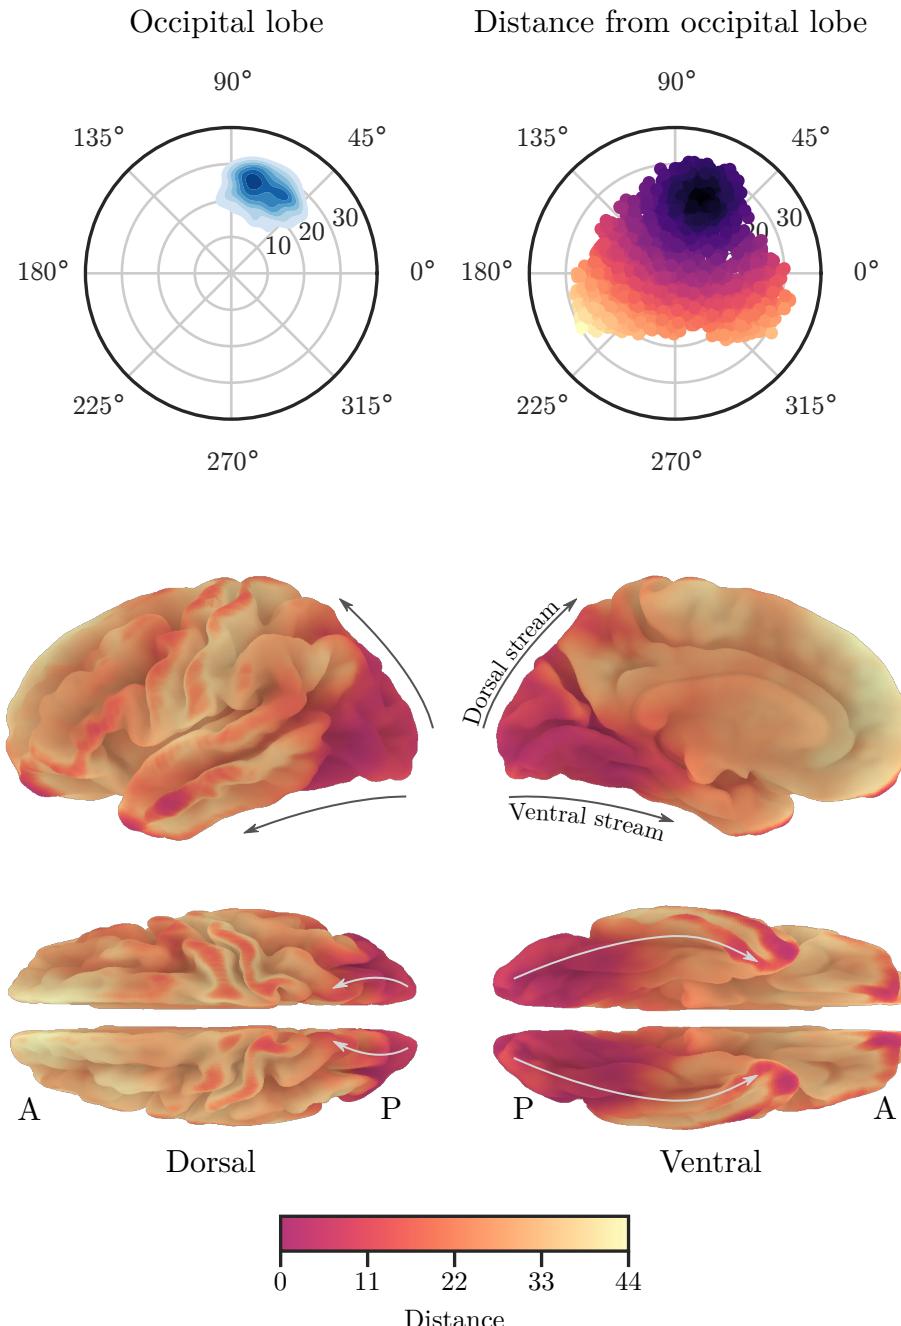


Figure 5: Occipital lobe intrinsic functional distance mapping. (Top left) Kernel density estimate plot of the occipital lobe regions in the isomap embedding. (Top right) Intrinsic functional distance to the occipital lobe for all regions. Darker color indicates the region is closer to the mean occipital lobe coordinate. (Bottom) Intrinsic functional distance to the occipital lobe projected onto the Freesurfer pial surface template. Arrows indicate the dorsal and ventral visual streams. A: anterior. P: posterior.

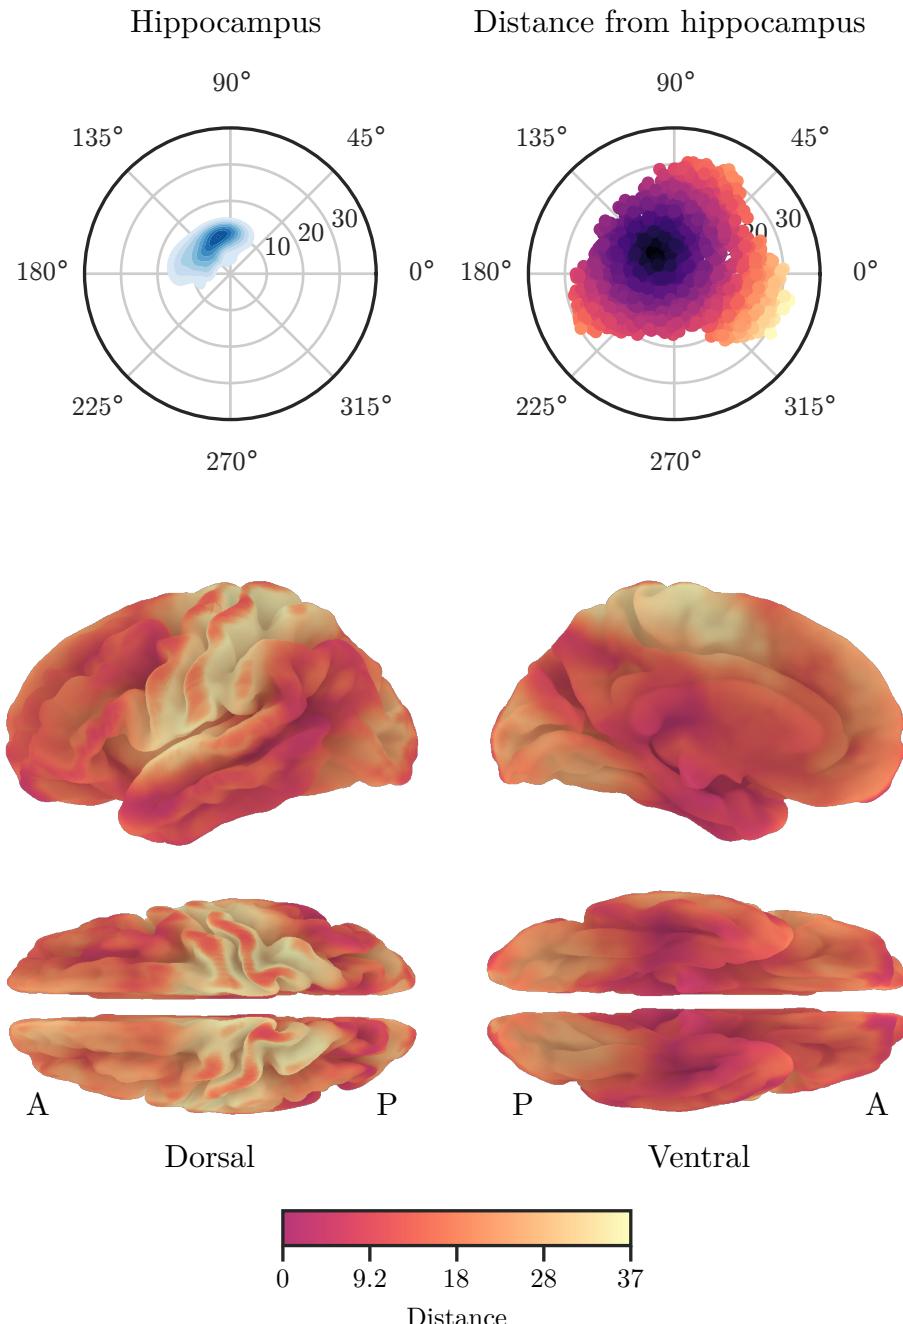


Figure 6: Hippocampus intrinsic functional distance mapping. (Top left) Kernel density estimate plot of the hippocampus regions in the isomap embedding. (Top right) Intrinsic functional distance to the hippocampus for all regions. Darker color indicates the region is closer to the mean hippocampus coordinate. (Bottom) Intrinsic functional distance to the hippocampus projected onto the Freesurfer pial surface template. A: anterior. P: posterior.

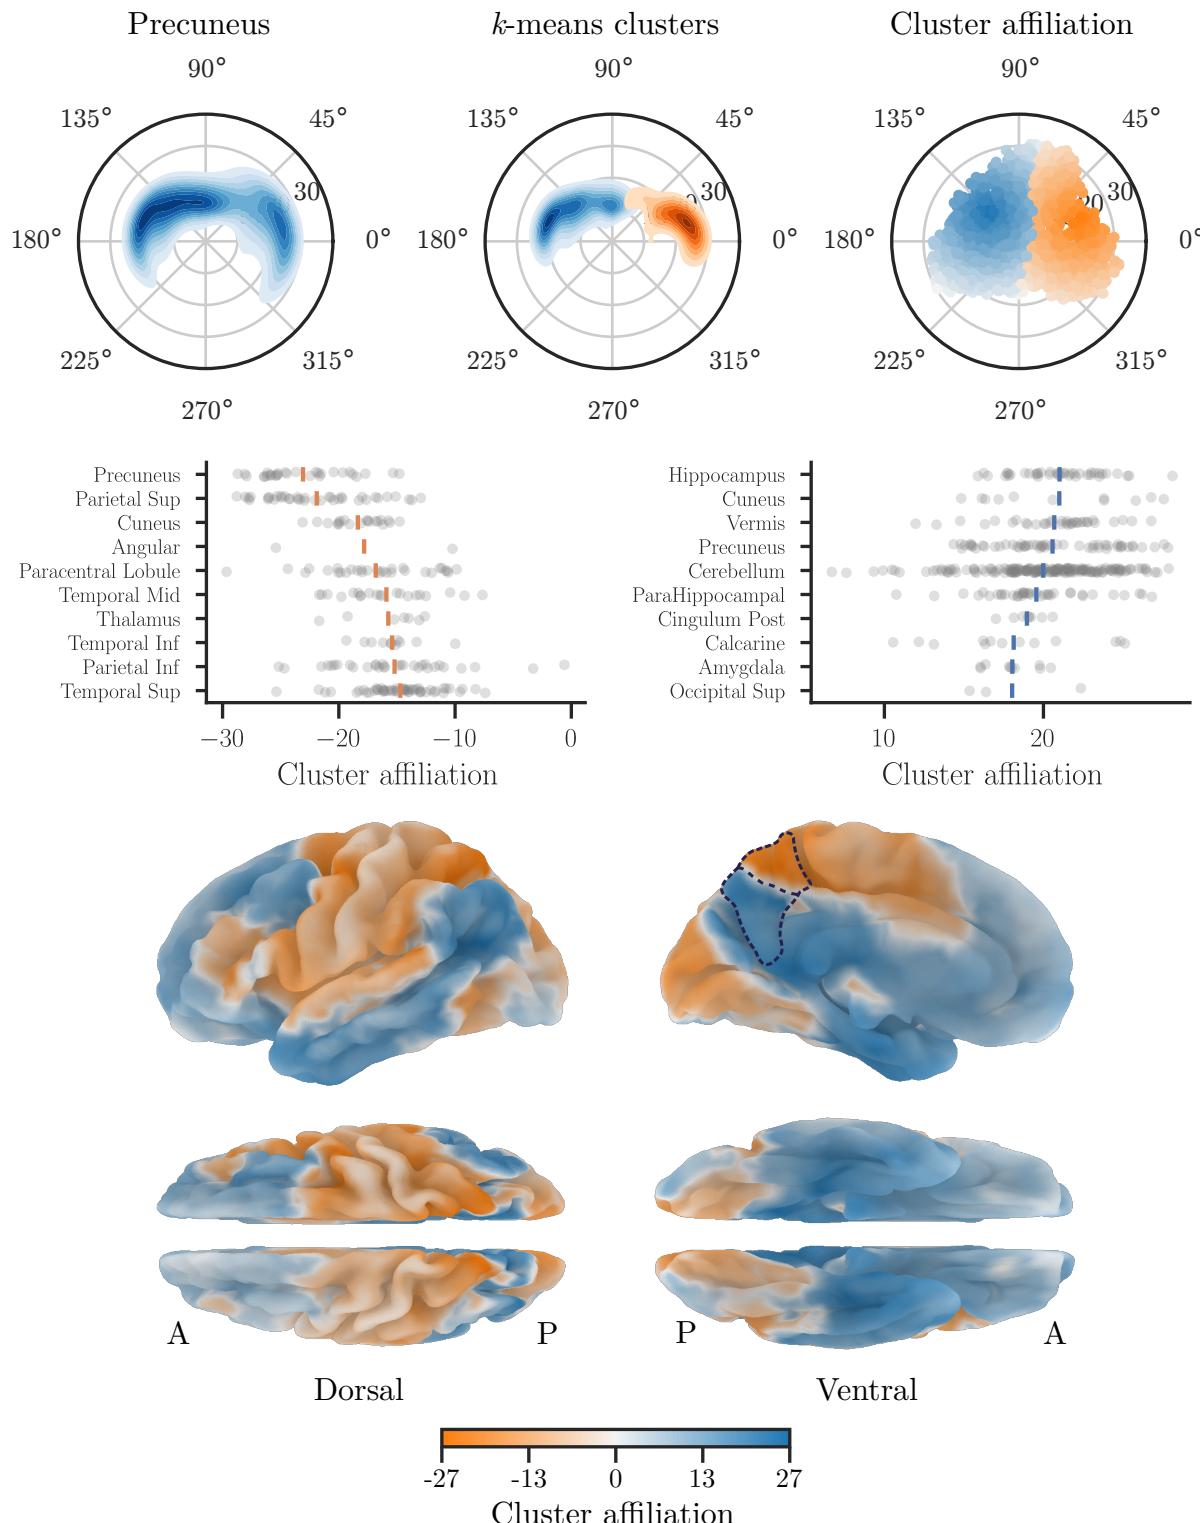


Figure 7: Identifying subnetwork clusters within the precuneus using rest2vec. (Top, left) Kernel density estimate of the precuneus in the isomap embedding. (Top, middle) *k*-means clustering results are indicated in blue and orange. (Top, right) Cluster affiliations for all other regions based on their 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.

316 3.3 Maximizing maximum mean discrepancy partitions the connec- 317 tome into putative task-positive and task-negative networks.

318 Since rest2vec could identify functionally relevant connectivity gradients within anatomical
319 lobes, we next asked if rest2vec could be used to partition RS-FMRI connectomes into functional
320 modules. To address this, we used the maximum mean discrepancy (MMD) metric developed
321 by [13] to partition the set of connectome regions V into two distributions of regions V_A and
322 V_B such that the MMD between them was maximized. A cosine kernel (Equation 8) was
323 used to compute the centered kernel similarity matrix $\tilde{\mathbf{K}}$ between all pairwise regions of Θ
324 (cf. Figure 2). Similar to spectral clustering methods, we approximated the maximum MMD
325 by reformulating the MMD to a Rayleigh quotient maximization problem (§2.2.3), where the
326 eigenvector \mathbf{q} corresponding to the maximum eigenvalue was extracted to yield the community
327 assignment vector. (Ranking of the eigenvectors of $\tilde{\mathbf{K}}$ showed the first three eigenvalues
328 account for most of the variance of the data, with the first being most dominant (Figure 8).)

329 To set the partition, the index of regions corresponding to $q_i \geq 0$ were assigned to V_A , and
330 the index of regions corresponding to $q_i < 0$ were assigned to V_B . This approach was validated
331 by iteratively evaluating the MMD across 50 threshold values of \mathbf{q} (Figure 9). The results
332 suggest that the MMD is maximized when the partition yields communities of approximately
333 equal size, which occurs for both datasets when the partition threshold for $q_i \approx 0$. Together
334 this suggests that the Rayleigh quotient maximization approximation is able to achieve an
335 accurate approximation of the global maximum MMD.

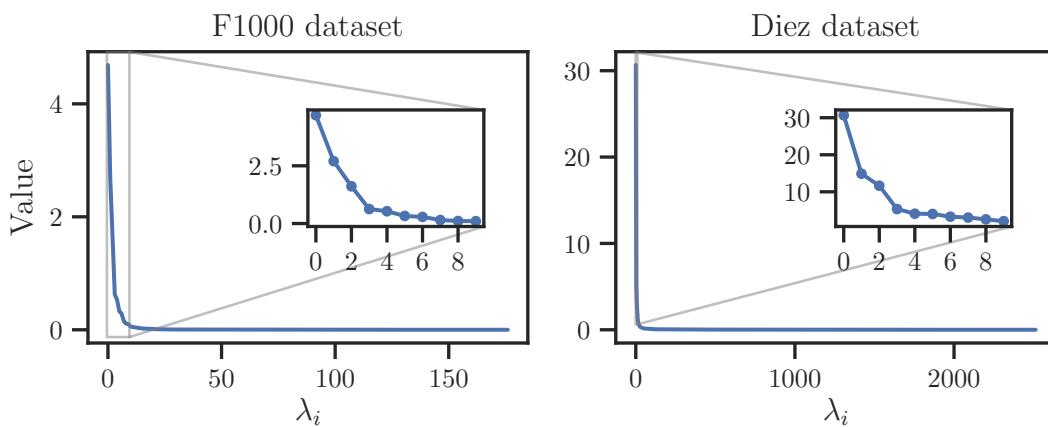


Figure 8: Eigenvalues of $\tilde{\mathbf{K}}$ for each dataset. Insets depict the first ten eigenvalues.

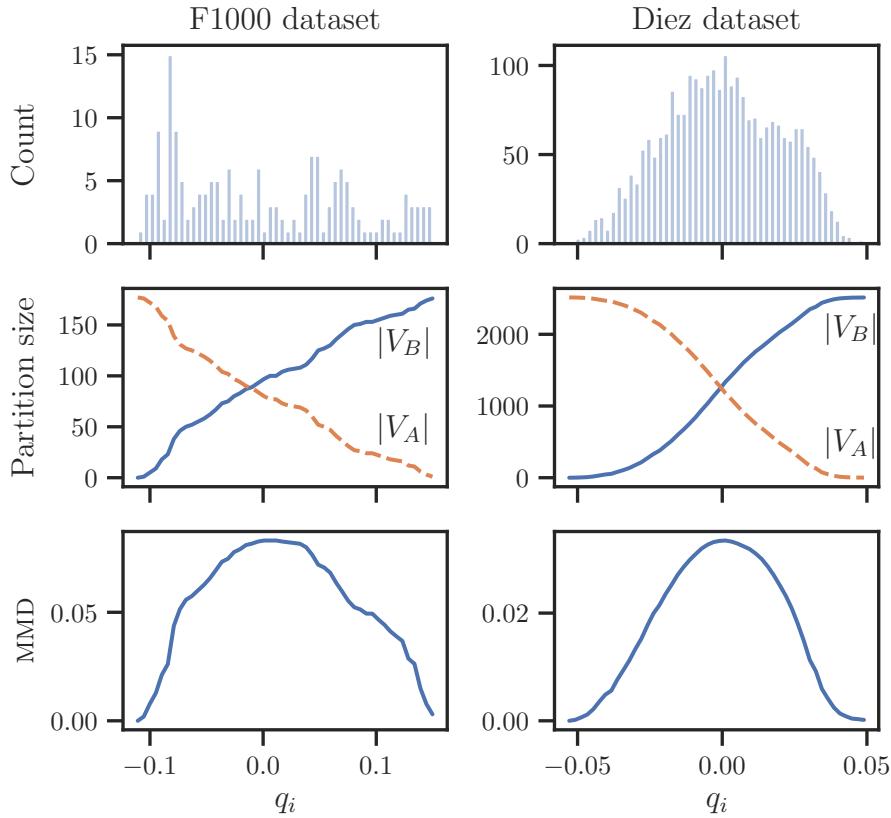


Figure 9: Evaluation of the maximum mean discrepancy (MMD) across a range of threshold values of \mathbf{q} for each dataset. (Top) Histogram of the eigenvector \mathbf{q} corresponding to the maximum eigenvalue of the centered kernel matrix $\tilde{\mathbf{K}}$. (Middle) Size of each partition as a function of threshold value. (Bottom) MMD as a function of threshold value.

336 Similar to the previous analysis in [Figure 4](#), we visualized the anatomical and functional
 337 embedding by community affiliation to see how each lobe participates in the two communities
 338 predicted by maximizing the MMD. We observed a symmetrical partition between the two
 339 communities when viewed in the functional embedding space ([Figure 10](#)). Additionally, when
 340 the magnitude and sign of q_i are mapped to a diverging colormap in the isomap space, it
 341 was observed that regions closer to the vertical axis appeared more neutral, whereas regions
 342 further from the vertical axis were polarized into either community, suggesting these regions
 343 are more strongly mapped into that community. When further examining the regions that
 344 are affiliated with each community, we observed that the partition demarcated into the
 345 putative task-positive network (TPN) and task-negative network (TNN, also called default
 346 mode network (DMN)). This relationship can be seen when the MMD eigenvector gradient
 347 is used to sort the Θ matrix for each datasets, where the resulting grid communities show
 348 out-of-phase relationships with the other community ([Figure 11](#)). This can also be observed
 349 anatomically when the eigenvector gradient is mapped to the brain surface ([Figure 12](#)).

350 When the brain is faceted by lobe affiliation, several notable patterns emerge. The frontal
351 lobe is demarcated into the prefrontal cortex (PFC) and pre-motor areas and dorsolateral
352 PFC anatomically, which are respectively situated in opposite quadrants of the functional
353 embedding. In addition, the parietal lobe is split largely into default mode network (DMN)
354 regions – including regions of the inferior parietal lobule and precuneus – and primary and
355 secondary unimodal areas, including somatosensory cortices and areas involved in visual
356 processing. Consistent with this observation, the occipital lobe has the largest proportion of
357 regions belonging to the putative TPN at 83.42%. Taken together, these data suggest that
358 maximizing the MMD in the context of vectorized connectomes is able to recover biologically-
359 relevant network characteristics, while also accounting for the presence of negative edges,
360 thereby removing heuristic steps that may bias downstream analyses as a result.

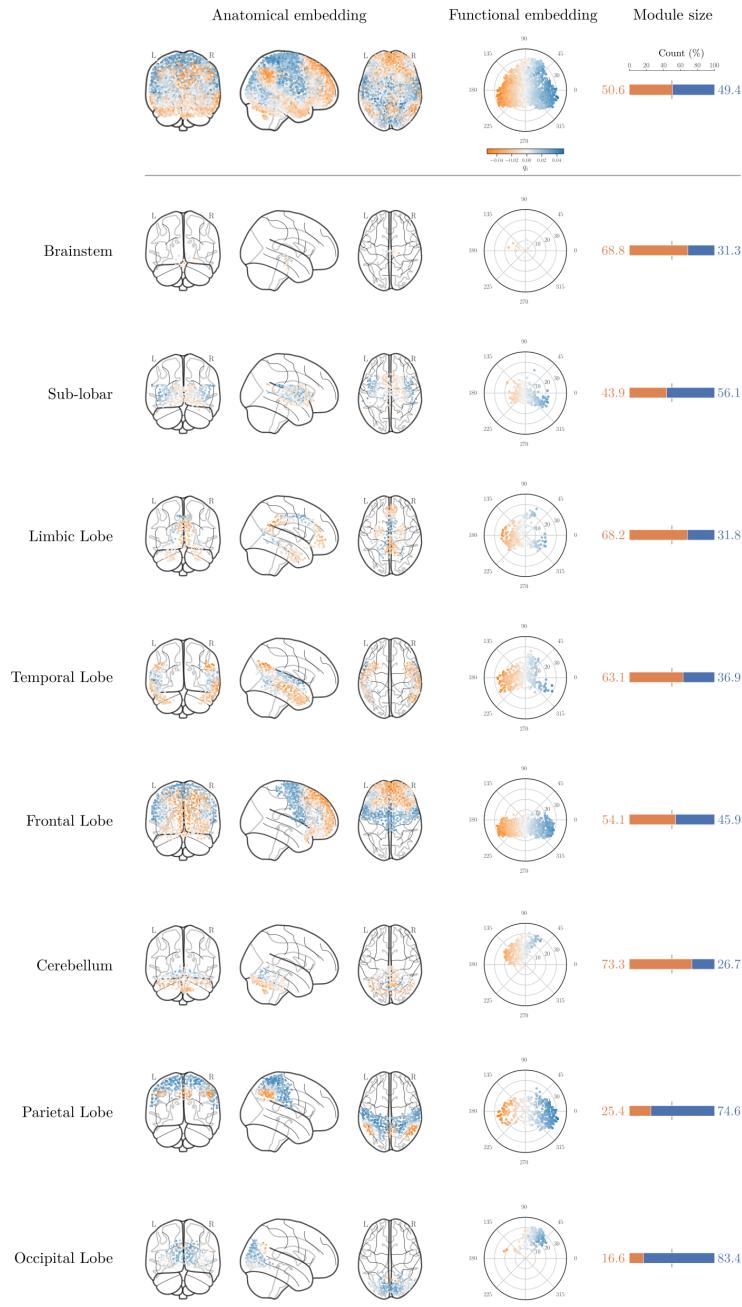


Figure 10: Anatomical and functional embedding of the Diez dataset faceted by anatomical lobe affiliation with predicted community partitions. (Top) Merged representations of all 2514 regions in anatomical embedding (columns 1-3), functional embedding (column 4), and the percentage of regions within 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%.

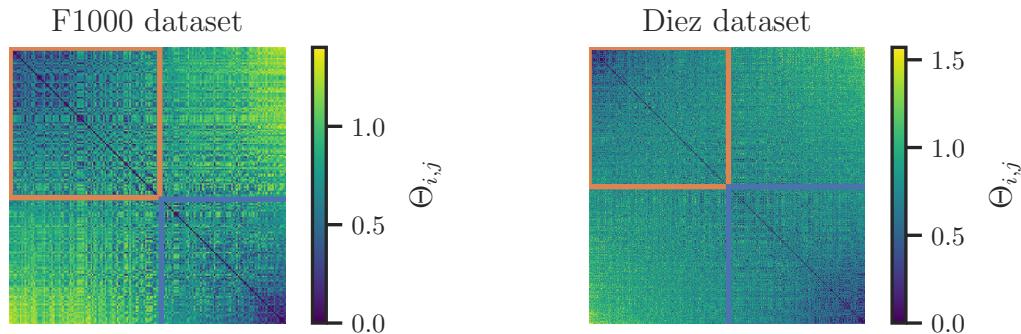


Figure 11: Community partition gradient given by the top eigenvector \mathbf{q} of the kernel similarity matrix $\tilde{\mathbf{K}}$. The PHASE matrix Θ for the F1000 and Diez datasets are shown with rows and columns sorted (in ascending order) by the elements of \mathbf{q} . Grid lines indicate the community partition (upper left and lower right).

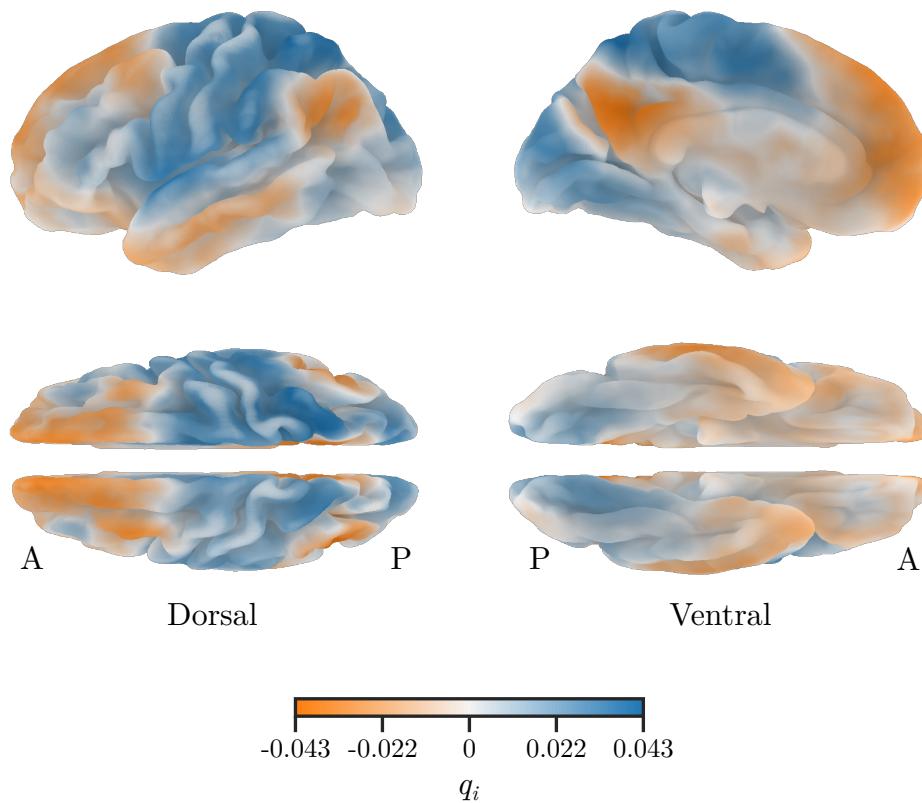


Figure 12: Brain surface map of the MMD partition gradient for the Diez dataset projected onto the Freesurfer pial surface template. Color indicates the interpolated value of q_i . A: anterior. P: posterior.

361 4 Discussion

362 In this study we presented a novel graph embedding approach for RS-FMRI connectivity using
363 rest2vec. Rest2vec improves upon current methods by using the full range of correlative
364 information and representing the functional relationships of the brain in a low-dimensional
365 embedding. Whereas many processing strategies involve arbitrary thresholds, rest2vec does
366 not involve removing any data from the functional connectome. Previous studies have
367 suggested that these negative correlations may have important – but still not fully understood
368 – biological roles [37]. While there exist variations of methods that account for negative edges,
369 such as the Louvain algorithm [14] and the Q^* -maximization method [9], the issue of deciding
370 the appropriate weight of contribution to assign to these edges still remains.

371 Previous work from our group demonstrated that using a probability-based divisive
372 approach with permutation testing could recover the hierarchical community structure of
373 RS-FMRI connectomes while preserving negative edges, which we called probability-associated
374 community estimation (PACE) [10]. In addition, our previous study [4] demonstrated how
375 nonlinear dimensionality reduction and manifold learning techniques could be used to in-
376 vestigate the intrinsic geometry of structural connectomes derived from diffusion imaging.
377 Inspired by these approaches, we sought to develop a method by which RS-FMRI functional
378 connectomes could be represented in their intrinsic geometry while also preserving negative
379 edge relationships.

380 Dimensionality reduction techniques have been previously applied to neuroimaging
381 datasets, e.g., clustering in lower dimensions to demarcate subjects belonging to different
382 clinical populations, such as healthy controls and patients. Here, rest2vec applies dimension-
383 ality reduction at the level of brain regions. Furthermore, we chose to use the isomap method
384 because it uses a geodesic distance metric for generating the lower-dimensional embedding [8].
385 By doing so, distance in the lower-dimensional embedding conveys meaningful information,
386 as opposed to other methods, such as t-SNE [38], that are stochastic and primarily meant for
387 clustering purposes.

388 In the context of functional connectivity, converting the coordinate system to a polar
389 representation was an intuitive visualization decision, as it centers the data around the
390 origin where regions with lower $\Theta_{i,j}$ values are mapped closer to the origin and regions with
391 higher $\Theta_{i,j}$ values are mapped in the periphery. Interestingly, regions with a greater number
392 of high $\Theta_{i,j}$ values (i.e., more out-of-phase relationships) tended to be unimodal and also
393 have low within-cluster $\Theta_{i,j}$ values, as seen most clearly in the occipital lobe (SI Figure 15).
394 In contrast, more centrally-embedded regions tended to be located in brainstem regions
395 (known to facilitate various sensory relay roles) and associative regions. This is reminiscent

396 of Mesulam's synaptic hierarchy model [39], where primary unimodal regions are embedded
397 at the periphery, most proximal to sensory input, with downstream synaptic connectivity
398 progressing inward towards the center to heteromodal and associative areas.

399 By using lower-dimensional embedding distance metrics, we were able to recover function-
400 ally relevant relationships. In the case of the occipital lobe, mapping the intrinsic functional
401 distance to its cluster centroid in the isomap embedding generated a gradient map in the
402 anatomical space of the dorsal and ventral visual streams [33, 34]. On the dorsal surface, the
403 gradient proximal to the occipital lobe can be seen going to the posterior parietal regions,
404 whereas on the ventral surface the proximal gradient extends from the occipital lobe to
405 the inferior temporal lobe (Figure 5). In another example, the precuneus had two primary
406 clusters in the isomap embedding. When projected onto the brain surface, these two clusters
407 demarcated the dorsal-anterior and ventral-posterior portions of the precuneus (Figure 7).
408 The dorsal-anterior gradient appeared to primarily consist of the superior parietal, somato-
409 motor, and occipital cortices. The ventral-posterior gradient appeared to be composed of the
410 posterior cingulate, parahippocampal, and superior occipital cortices and the hippocampus.

411 There is evidence for the dorsal-anterior and ventral-posterior portions of the precuneus
412 being involved in different functions. A RS-FMRI study by [36] identified the dorsal and
413 anterior portions of precuneus having stronger connectivity with areas including the occipital,
414 somatomotor, and posterior parietal cortices and the superior temporal gyri. In addition, they
415 identified the ventral precuneus as being more strongly associated with the middle frontal
416 gyrus, posterior cingulate cortex, cuneus, and calcarine sulcus. This demarcation is thought
417 to be due to the diverse roles of the precuneus. In particular, the dorsal-anterior portion
418 of the precuneus, which has strong connectivity with the occipital and superior parietal
419 cortices, is involved in processing polymodal imagery and visuospatial information, whereas
420 the ventral-posterior precuneus is thought to be more involved in episodic memory retrieval
421 [35]. While the study by [36] further subdivided the precuneus into eight clusters in their
422 study, our results were largely consistent with their observations, suggesting that rest2vec
423 can detect heterogeneous connectivity patterns within individual regions.

424 In addition to representing the intrinsic geometry of functional connectomes, we proposed
425 using the maximum mean discrepancy (MMD) method by [13] to partition the connectome
426 into maximally functionally distinct modules. The MMD was originally implemented to
427 detect how different two probability distributions were to test if they were from the same
428 population [13]. For our use case, we maximized the MMD as an objective function to find two
429 populations of brain regions such that their distributions are as distant as possible to identify
430 functional communities. One advantage is this is a vectorized approach and does not rely on
431 iterative methods. In addition, this method offers flexibility in the choice of probability and

432 kernel similarity measures used as input, and so are not limited to only Pearson correlation
433 measures.

434 When the functional connectome is represented in its intrinsic embedding using nonlinear
435 dimensionality reduction, the MMD partition elicited a strikingly symmetric representation.
436 Upon closer observation, these two communities were split approximately between the
437 canonical task-positive network (TPN) and the default mode network (DMN), consisting of
438 the precuneus, inferior parietal lobule (IPL), posterior cingulate cortex, hippocampus, and
439 areas of the prefrontal cortex (PFC), among others [40]. This initial bifurcation is consistent
440 with previous modularity studies [10], and is a validation that this embedding procedure is
441 capturing functionally-relevant characteristics. In addition, we showed lobe-specific affiliations
442 for the two communities. These results were consistent with the putative DMN/TPN split.
443 Notably, the IPL and precuneus are shown in contrast to the postcentral regions within the
444 parietal lobe; similarly, the PFC and pre-motor areas show clear boundaries. Together, these
445 results demonstrated that using this MMD approach to solve the connectome modularity
446 problem yielded reproducible and biologically-meaningful connectome partitions, and that
447 the properties of these communities can be represented using dimensionality reduction.

448 **Limitations and future directions.** In this paper, we used RS-FMRI connectomes from
449 a group of subjects in order to compute the probability of there being a negative correlation
450 between each pairwise edge between regions. While this approach led to consistent results
451 across two independent datasets, we did not assess how robust this procedure was to inter-
452 subject variability or the size of groups. In addition, while average RS-FMRI connectomes
453 yield a wealth of functional connectivity information, they are a static representation of a
454 dynamic process. Furthermore, there has been increasing emphasis on individual connectome
455 analysis with aims towards personalized medicine [41, 42]. To that end, future improvements
456 on these methods will need to incorporate dynamic as well as subject-specific analyses of
457 functional connectivity. Recent works by [43, 44] have suggested the concept of hierarchical
458 or multi-scale networks, which could lead to natural extensions of this work via subject
459 embedding spaces which are in turn composed of network embedding spaces.

460 Another limitation is we only examined the MMD partition at the first bifurcation into
461 two communities. While this proved effective as a proof-of-concept, further work will need
462 to be done to develop a hierarchical way to detect N communities with this approach. In
463 addition, more robust methods could be used for maximizing the MMD objective function to
464 avoid the possibility of local maxima to achieve better accuracy.

465 **Conclusion.** Rest2vec incorporates both positive and negative edge connectivity using a
466 model inspired by statistical mechanics to transform functional connectome data into phase
467 angle relationships. This representation of the connectome can be combined with nonlinear
468 dimensionality reduction techniques to represent the intrinsic geometry of the functional
469 connectome in a lower-dimensional embedding. Together, these methods allow for a vectorized
470 approach to investigate the functional relationships of RS-FMRI brain connectivity data. In
471 addition, we connected rest2vec to the maximum mean discrepancy metric to demonstrate
472 how rest2vec can be used to address the modularity problem as a kernel two-sample test.
473 In summary, we presented a RS-FMRI connectome graph embedding technique that uses
474 nonlinear dimensionality reduction and statistical learning methods to create a low-dimensional
475 representation of the intrinsic geometry of the functional connectome.

476 **Acknowledgements.** This study was supported by NIH AG057468 to ZDM, NIH
477 AG056782 to LZ and ADL. The authors declare no conflicts of interest.

478 5 Appendix: Supplemental information

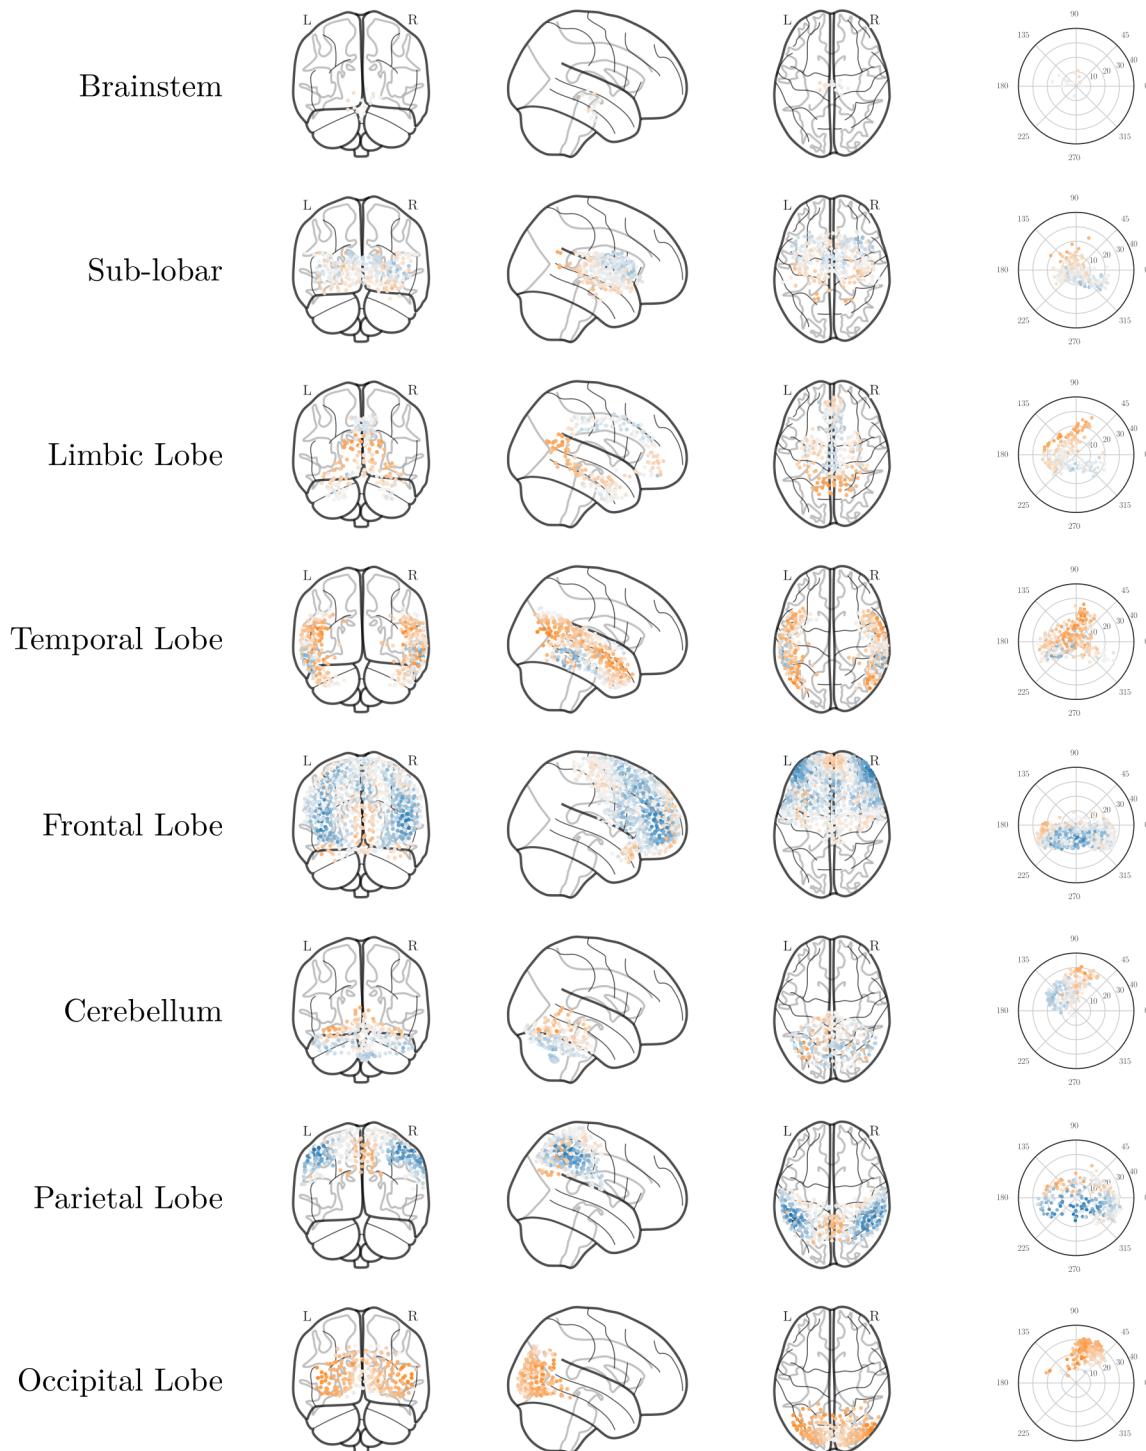


Figure 13: Anatomical and functional embedding of the Diez dataset faceted by anatomical lobe affiliation. Color indicates the partition predicted by the eigenvector of $\tilde{\mathbf{K}}$ corresponding to the second-highest eigenvalue.

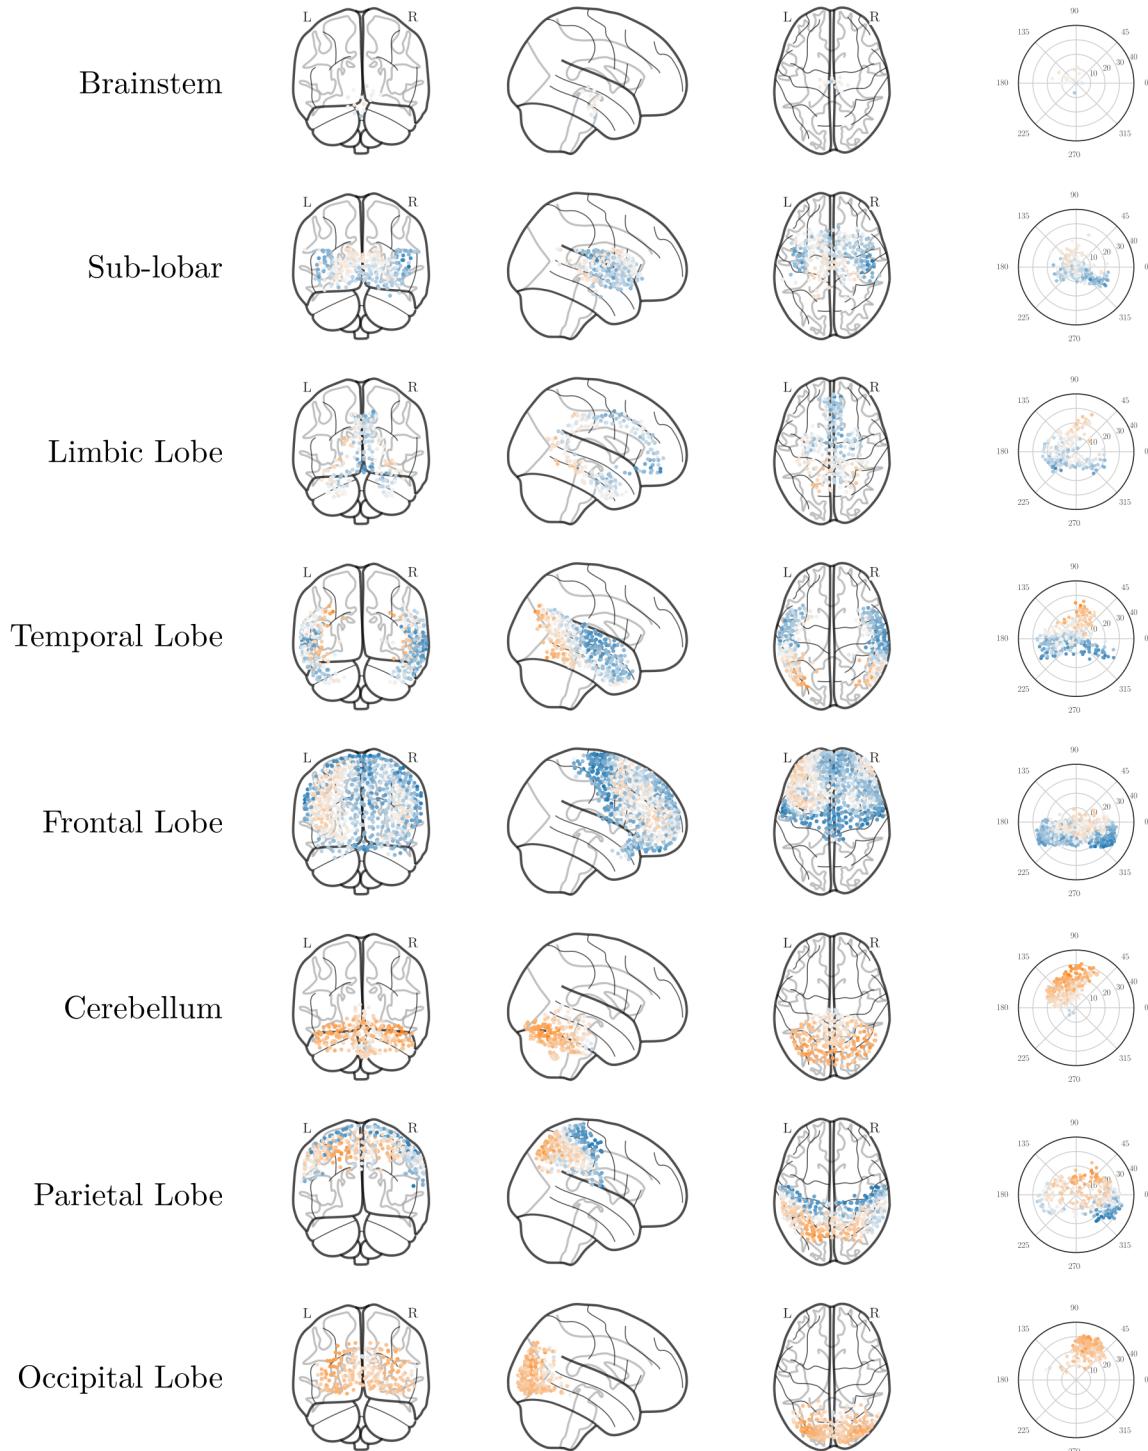


Figure 14: Anatomical and functional embedding of the Diez dataset faceted by anatomical lobe affiliation. Color indicates the partition predicted by the eigenvector of $\tilde{\mathbf{K}}$ corresponding to the third-highest eigenvalue.

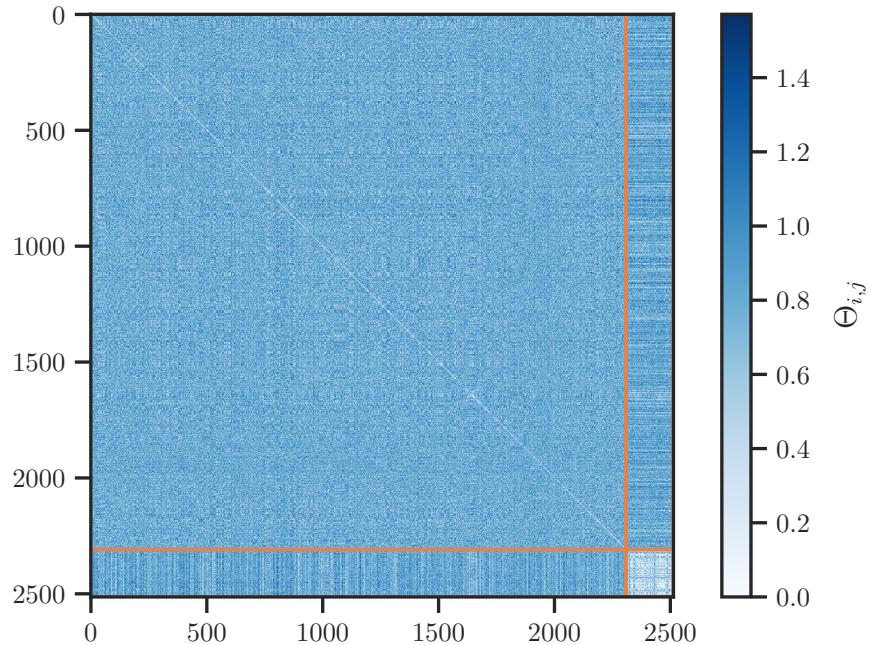


Figure 15: Θ for the Diez dataset with rows and columns arranged in a diagonal grid into non-occipital lobe regions (top left block) and occipital lobe regions (bottom right). Boundary between blocks is indicated by the orange line.

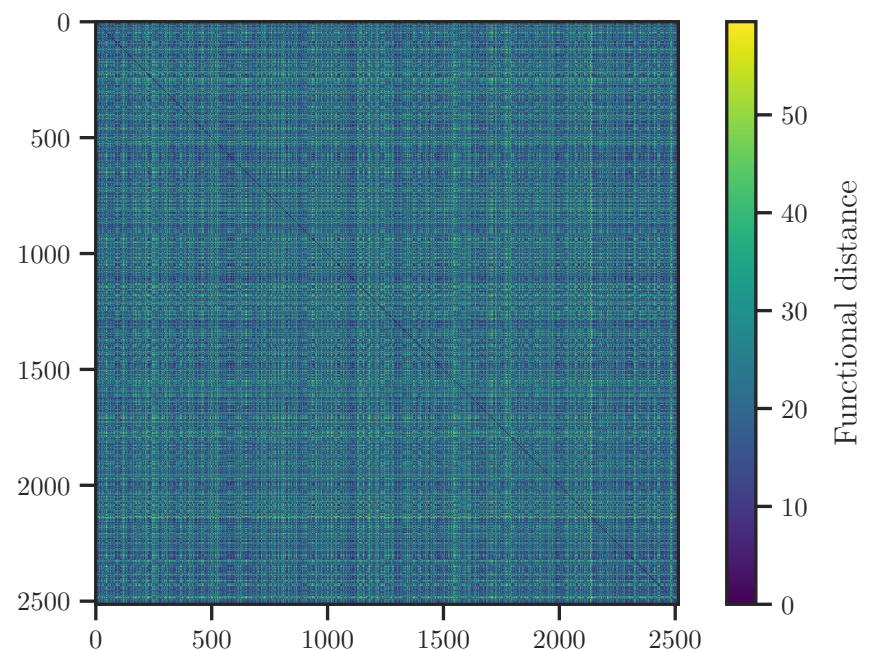


Figure 16: Pairwise functional distance for the Diez dataset.

479 References

480 [1] O. Sporns, G. Tononi, and R. Kötter, "The Human Connectome: A Structural Description of
481 the Human Brain," *PLoS Computational Biology*, vol. 1, no. 4, p. e42, 2005. [Online]. Available:
482 <http://dx.plos.org/10.1371/journal.pcbi.0010042>

483 [2] R. E. Bellman, *Adaptive Control Processes: A Guided Tour*. Princeton university press, 2015.

484 [3] J. A. Lee and M. Verleysen, *Nonlinear Dimensionality Reduction*. Springer Science & Business Media,
485 2007.

486 [4] A. Q. Ye, O. A. Ajilore, G. Conte, J. GadElkarim, G. Thomas-Ramos, L. Zhan, S. Yang,
487 A. Kumar, R. L. Magin, A. G. Forbes, and A. D. Leow, "The intrinsic geometry of the human
488 brain connectome," *Brain Informatics*, vol. 2, no. 4, pp. 197–210, Dec. 2015. [Online]. Available:
489 <http://link.springer.com/10.1007/s40708-015-0022-2>

490 [5] J. McClurkin, L. Optican, B. Richmond, and T. Gawne, "Concurrent processing and complexity of
491 temporally encoded neuronal messages in visual perception," *Science*, vol. 253, no. 5020, pp. 675–677,
492 Aug. 1991. [Online]. Available: <http://www.sciencemag.org/cgi/doi/10.1126/science.1908118>

493 [6] R. Wolz, P. Aljabar, J. V. Hajnal, J. Lötjönen, and D. Rueckert, "Nonlinear dimensionality reduction
494 combining MR imaging with non-imaging information," *Medical Image Analysis*, vol. 16, no. 4, pp.
495 819–830, May 2012. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S1361841511001733>

496 [7] S. Gerber, T. Tasdizen, P. Thomas Fletcher, S. Joshi, and R. Whitaker, "Manifold modeling for brain
497 population analysis," *Medical Image Analysis*, vol. 14, no. 5, pp. 643–653, Oct. 2010. [Online]. Available:
498 <https://linkinghub.elsevier.com/retrieve/pii/S1361841510000617>

499 [8] J. B. Tenenbaum, J. de Silva, and J. C. Langford, "A Global Geometric Framework for Nonlinear
500 Dimensionality Reduction," *Science*, vol. 290, no. 5500, pp. 2319–2323, Dec. 2000. [Online]. Available:
501 <http://www.sciencemag.org/cgi/doi/10.1126/science.290.5500.2319>

502 [9] M. Rubinov and O. Sporns, "Weight-conserving characterization of complex functional brain
503 networks," *NeuroImage*, vol. 56, no. 4, pp. 2068–2079, Jun. 2011. [Online]. Available:
504 <http://linkinghub.elsevier.com/retrieve/pii/S105381191100348X>

505 [10] L. Zhan, L. M. Jenkins, O. E. Wolfson, J. J. GadElkarim, K. Nocito, P. M. Thompson, O. A.
506 Ajilore, M. K. Chung, and A. D. Leow, "The significance of negative correlations in brain
507 connectivity," *Journal of Comparative Neurology*, vol. 525, no. 15, pp. 3251–3265, Oct. 2017, cites:
508 Zhansignificancenegativecorrelations2017. [Online]. Available: <http://doi.wiley.com/10.1002/cne.24274>

509 [11] Z. Morrissey, L. Zhan, H. Lee, J. Keiriz, A. Forbes, O. Ajilore, A. Leow, and M. Chung,
510 "Phase Angle Spatial Embedding (PhASE): A Kernel Method for Studying the Topology of the
511 Human Functional Connectome," in *Medical Image Computing and Computer Assisted Intervention –*
512 *MICCAI 2018*, A. F. Frangi, J. A. Schnabel, C. Davatzikos, C. Alberola-López, and G. Fichtinger,
513 Eds. Cham: Springer International Publishing, 2018, vol. 11072, pp. 367–374. [Online]. Available:
514 http://link.springer.com/10.1007/978-3-030-00931-1_42

515 [12] A. Pękalski, "Ising model on a small world network," *Physical Review E*, vol. 64, no. 5, Oct. 2001.
516 [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevE.64.057104>

517 [13] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, “A kernel two-sample test,”
518 *Journal of Machine Learning Research*, vol. 13, no. Mar, pp. 723–773, 2012.

519 [14] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre, “Fast unfolding of communities in large
520 networks,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2008, no. 10, p. P10008,
521 Oct. 2008. [Online]. Available: <http://stacks.iop.org/1742-5468/2008/i=10/a=P10008?key=crossref.46968f6ec61eb8f907a760be1c5ace52>

523 [15] B. B. Biswal, M. Mennes, X.-N. Zuo, S. Gohel, C. Kelly, S. M. Smith, C. F. Beckmann, J. S. Adelstein,
524 R. L. Buckner, S. Colcombe, A.-M. Dogonowski, M. Ernst, D. Fair, M. Hampson, M. J. Hoptman, J. S.
525 Hyde, V. J. Kiviniemi, R. Kotter, S.-J. Li, C.-P. Lin, M. J. Lowe, C. Mackay, D. J. Madden, K. H.
526 Madsen, D. S. Margulies, H. S. Mayberg, K. McMahon, C. S. Monk, S. H. Mostofsky, B. J. Nagel, J. J.
527 Pekar, S. J. Peltier, S. E. Petersen, V. Riedl, S. A. R. B. Rombouts, B. Rypma, B. L. Schlaggar,
528 S. Schmidt, R. D. Seidler, G. J. Siegle, C. Sorg, G.-J. Teng, J. Veijola, A. Villringer, M. Walter,
529 L. Wang, X.-C. Weng, S. Whitfield-Gabrieli, P. Williamson, C. Windischberger, Y.-F. Zang, H.-Y.
530 Zhang, F. X. Castellanos, and M. P. Milham, “Toward discovery science of human brain function,”
531 *Proceedings of the National Academy of Sciences*, vol. 107, no. 10, pp. 4734–4739, Mar. 2010. [Online].
532 Available: <http://www.pnas.org/cgi/doi/10.1073/pnas.0911855107>

533 [16] I. Diez, P. Bonifazi, I. Escudero, B. Mateos, M. A. Muñoz, S. Stramaglia, and J. M. Cortes, “A novel
534 brain partition highlights the modular skeleton shared by structure and function,” *Scientific Reports*,
535 vol. 5, no. 1, Sep. 2015. [Online]. Available: <http://www.nature.com/articles/srep10532>

536 [17] “Anaconda Software Distribution,” Anaconda, Dec. 2018. [Online]. Available: <https://anaconda.com/>

537 [18] T. E. Oliphant, *A Guide to NumPy*. Trelgol Publishing USA, 2006, vol. 1.

538 [19] W. McKinney, “Data Structures for Statistical Computing in Python,” p. 6, 2010.

539 [20] J. D. Hunter, “Matplotlib: A 2D Graphics Environment,” *Computing in Science & Engineering*, vol. 9,
540 no. 3, pp. 90–95, 2007. [Online]. Available: <http://ieeexplore.ieee.org/document/4160265/>

541 [21] “Seaborn: Statistical data visualization.” [Online]. Available: <https://seaborn.pydata.org/index.html>

542 [22] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer,
543 R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duch-
544 esnay, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp.
545 2825–2830, 2011.

546 [23] F. Perez and B. E. Granger, “IPython: A System for Interactive Scientific Computing,”
547 *Computing in Science & Engineering*, vol. 9, no. 3, pp. 21–29, 2007. [Online]. Available:
548 <http://ieeexplore.ieee.org/document/4160251/>

549 [24] S. Fortunato, “Community detection in graphs,” *Physics Reports*, vol. 486, no. 3–5, pp. 75–174, Feb.
550 2010. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0370157309002841>

551 [25] E. W. Weisstein, “Fiedler Vector,” from MathWorld—A Wolfram Web Resource. [Online]. Available:
552 <https://mathworld.wolfram.com/FiedlerVector.html>

553 [26] E. W. Dijkstra, “A Note on Two Problems in Connexion with Graphs,” p. 3, 1959.

554 [27] S. Seabold and J. Perktold, “Statsmodels: Econometric and Statistical Modeling with Python,” p. 5,
555 2010.

556 [28] A. Abraham, F. Pedregosa, M. Eickenberg, P. Gervais, A. Mueller, J. Kossaifi, A. Gramfort, B. Thirion,
557 and G. Varoquaux, "Machine learning for neuroimaging with scikit-learn," *Frontiers in Neuroinformatics*,
558 vol. 8, 2014. [Online]. Available: <http://journal.frontiersin.org/article/10.3389/fninf.2014.00014/abstract>

559 [29] M. Albert, J. A. Andler, T. Bah, P. Barbry-Blot, J.-F. Barraud, C. Barton, B. Baxter,
560 J. Beard, J. Bintz, A. Biro, N. Bishop, J. L. Blocher, H. Böck, T. Boczkowski, H. Bohre,
561 Boldewyn, D. Borgmann, B. Bouclet, H. Breuer, G. Broberg, C. Brown, M. Brubaker, L. Bruno,
562 B. (brynn@inkscapecommunity.com), N. Buculei, B. Byak, P. Caclin, I. Caldwell, G. Carmichael,
563 E. Catmur, C. Celorio, J. A. Cenoz, J. Ceuppens, Z. Chyla, A. Clausen, J. Cliff, K. Cook, B. Cromwell,
564 R. Crosbie, J. Cruz, A. De-Cooman, M. Dereyynski, D. Díaz, B. Dilly, L. Doolittle, N. Dufour, T. Dwyer,
565 M. V. Dziumanenko, M. Eberl, J. Engelen, M. Erdelyi, U. Erikson, N. Falzon, S. Faubel, F. Felfe,
566 A. Fitzsimon, E. Flick, M. Floryan, B. Fowler, Fred, C. Gemy, S. Giannini, O. Gondouin, T. Gould,
567 T. de Greef, M. Grosberg, K. D. Gussem, B. Harrington, D. Harvey, A. A. Heckert, C. Hetherington,
568 and J. Hirth, "Inkscape." [Online]. Available: <https://inkscape.org/%0070>

569 [30] M. Jenkinson, C. F. Beckmann, T. E. Behrens, M. W. Woolrich, and S. M. Smith,
570 "FSL," *NeuroImage*, vol. 62, no. 2, pp. 782–790, Aug. 2012. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S1053811911010603>

571 [31] K. Gorgolewski, C. D. Burns, C. Madison, D. Clark, Y. O. Halchenko, M. L. Waskom,
572 and S. S. Ghosh, "Nipype: A Flexible, Lightweight and Extensible Neuroimaging Data
573 Processing Framework in Python," *Frontiers in Neuroinformatics*, vol. 5, 2011. [Online]. Available:
574 <http://journal.frontiersin.org/article/10.3389/fninf.2011.00013/abstract>

575 [32] L. G. Ungerleider, "Two cortical visual systems," *Analysis of visual behavior*, pp. 549–586, 1982.

576 [33] M. A. Goodale, "Transforming vision into action," *Vision Research*, vol. 51, no. 13, pp. 1567–1587, Jul.
577 2011. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0042698910003743>

578 [34] M. A. Goodale, A. D. Milner *et al.*, "Separate visual pathways for perception and action," 1992.

579 [35] A. E. Cavanna and M. R. Trimble, "The precuneus: A review of its functional anatomy and behavioural
580 correlates," *Brain*, vol. 129, no. 3, pp. 564–583, Mar. 2006. [Online]. Available: <http://academic.oup.com/brain/article/129/3/564/390904/The-precuneus-a-review-of-its-functional-anatomy>

581 [36] S. Zhang and C.-s. R. Li, "Functional connectivity mapping of the human precuneus by
582 resting state fMRI," *NeuroImage*, vol. 59, no. 4, pp. 3548–3562, Feb. 2012. [Online]. Available:
583 <https://linkinghub.elsevier.com/retrieve/pii/S1053811911013048>

584 [37] M. Rubinov and O. Sporns, "Complex network measures of brain connectivity: Uses and
585 interpretations," *NeuroImage*, vol. 52, no. 3, pp. 1059–1069, Sep. 2010. [Online]. Available:
586 <http://linkinghub.elsevier.com/retrieve/pii/S105381190901074X>

587 [38] L. van der Maaten and G. Hinton, "Visualizing data using t-SNE," *Journal of machine learning research*,
588 vol. 9, no. Nov, pp. 2579–2605, 2008.

589 [39] M. Mesulam, "From sensation to cognition," *Brain*, vol. 121, no. 6, pp. 1013–1052, Jun. 1998. [Online].
590 Available: <https://academic.oup.com/brain/article-lookup/doi/10.1093/brain/121.6.1013>

591 [40] R. L. Buckner, J. R. Andrews-Hanna, and D. L. Schacter, "The Brain's Default Network: Anatomy,
592 Function, and Relevance to Disease," *Annals of the New York Academy of Sciences*, vol. 1124, no. 1, pp.
593 1–38, Mar. 2008. [Online]. Available: <http://doi.wiley.com/10.1196/annals.1440.011>

596 [41] E. S. Finn, X. Shen, D. Scheinost, M. D. Rosenberg, J. Huang, M. M. Chun, X. Papademetris, and
597 R. T. Constable, “Functional connectome fingerprinting: Identifying individuals using patterns of brain
598 connectivity,” *Nature Neuroscience*, vol. 18, no. 11, pp. 1664–1671, Nov. 2015. [Online]. Available:
599 <http://www.nature.com/articles/nn.4135>

600 [42] O. Miranda-Dominguez, B. D. Mills, S. D. Carpenter, K. A. Grant, C. D. Kroenke, J. T. Nigg, and D. A.
601 Fair, “Connectotyping: Model Based Fingerprinting of the Functional Connectome,” *PLoS ONE*, vol. 9,
602 no. 11, p. e111048, Nov. 2014. [Online]. Available: <http://dx.plos.org/10.1371/journal.pone.0111048>

603 [43] R. F. Betzel and D. S. Bassett, “Multi-scale brain networks,” *NeuroImage*, vol. 160, pp. 73–83, Oct.
604 2017. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S1053811916306152>

605 [44] M. Gosak, R. Markovič, J. Dolenšek, M. Slak Rupnik, M. Marhl, A. Stožer, and M. Perc, “Network
606 science of biological systems at different scales: A review,” *Physics of Life Reviews*, Nov. 2017. [Online].
607 Available: <http://linkinghub.elsevier.com/retrieve/pii/S1571064517301501>