

LEARNING A COMPACT, PARCEL-INDEPENDENT REPRESENTATION OF THE fMRI FUNCTIONAL CONNECTIVITY

Anonymous authors

Paper under double-blind review

ABSTRACT

Functional connectivity in functional magnetic resonance imaging (fMRI) data is often calculated at the level of area parcels. Given the data’s low-dimensional nature, we posit a substantial degree of redundancy in these representations. Moreover, establishing correspondence across different individuals poses a significant challenge in that framework. We hypothesize that learning a compact representation of the functional connectivity data without losing the essential structure of the original data is possible. Our analysis, based on various performance benchmarks, indicates that the pre-computed mapping to low-dimensional latent space learned from the functional connectivity of one dataset generalizes well to another with both linear and non-linear autoencoder-based methods. Notably, the latent space learned using a variational autoencoder represents the data more effectively than linear methods at lower dimensions (2 dimensions). However, at higher dimensions (32 dimensions), the differences between linear and nonlinear dimensionality reduction methods diminish, rendering the performance comparable to the parcel space representation with 333 dimensions. Our findings highlight the potential of employing an established transformation to obtain a low-dimensional latent representation in future functional connectivity research, thereby solving the correspondence problem across parcel definitions, promoting reproducibility, and supporting open science objectives.

1 INTRODUCTION

Spontaneous brain activity, as measured using the blood-oxygen-level dependent (BOLD) signals estimated with functional magnetic resonance imaging (fMRI), demonstrates a high correspondence between functionally homotopic brain regions (Biswal et al., 1995). Researchers have measured resting-state functional connectivity to delineate the functional organizations of the brain (Bijsterbosch et al., 2020; Eickhoff et al., 2015). Many studies omitted the spatial topography and described the brain organization as abstract nodes and edges, giving the pairwise “functional connectome” matrix and the associated network graph representations (Bassett & Sporns, 2017; Bullmore & Sporns, 2009; Cheirdaris, 2023; Rubinov & Sporns, 2010; Sporns et al., 2004). Many node definitions exist and vary in number from less than 100 to 1000 area parcels in an atlas (Arslan et al., 2018; Craddock et al., 2012; Gordon et al., 2016; Schaefer et al., 2018; Shen et al., 2013). In addition to the lack of consensus in node definition, recent work has demonstrated that the optimal definition of nodes can vary across individuals (Gordon et al., 2017b) and populations (Han et al., 2018; Myers et al., 2024). Secondly, recent studies showed that only a few dimensions can capture the majority of variance in functional connectivity data (Bolt et al., 2022; Gotts et al., 2020; Margulies et al., 2016; Snyder et al., 2022), potentially due to the high spatial correlation in the cerebral cortex (Pang et al., 2023; Shinn et al., 2023). Therefore, we explored the possibility of learning a pre-computed transformation that maps the functional connectivity seed maps (spatial patterns of connectivity from one node to the rest of the brain) to a low-dimensional latent representation with little compromise in preserving the key information for downstream analyses.

We obtained low-dimensional embeddings (2, 4, 32 dimensions) from the high-dimensional functional connectivity seed map (59412 dimensions) data using a few variants of autoencoders (Rumelhart et al., 1986) including the conventional autoencoder (Rumelhart et al., 1986), vari-

054 ational autoencoder(Kingma & Welling, 2022), β -variational autoencoder (β -VAE)(Higgins et al.,
055 2017) , and adversarial autoencoder(Makhzani et al., 2016). We also compared the results with lin-
056 ear dimensionality reduction methods using Principal Component Analysis (PCA) and Independent
057 Component Analysis (ICA). We explored the extent to which key features in the original data could
058 be retained using these low-dimensional embeddings, and how it varied across dimensionality re-
059 duction methods and the number of dimensions. We adopted multiple performance metrics and used
060 the parcel connectome with the Gordon parcellation (Gordon et al., 2016) of 333 nodes (a.k.a. 333
061 dimensions) as a reference benchmark.

062 2 RELATED WORK

063 Our current work is most similar to the representation of functional connectivity into principal gra-
064 dients with diffusion embedding and other non-deep-learning-based methods (Langs et al., 2015;
065 2010; Margulies et al., 2016; Vos de Wael et al., 2020). However, learning the low-dimensional
066 manifold in individuals and aligning them requires matched parcels across subjects using Procrustes
067 analysis (Langs et al., 2015) and can be computationally expensive. Even more detrimentally, if the
068 manifold spaces (e.g. across development or clinical conditions) differ substantially, the alignments
069 may provide output that is not meaningful (Vos de Wael et al., 2020). Furthermore, this approach
070 does not provide a backward projection to the original space from the embeddings to give an intu-
071 itive demonstration of a walk in the latent space. More discussion about challenges in establishing
072 correspondence across subjects is provided in Appendix A.1.

073 There are also numerous other deep-learning approaches to capture fMRI activity and connectivity
074 with advanced deep neural network architectures(Caro et al., 2023; Zhang et al., 2020; Zhao et al.,
075 2020; Zuo et al., 2023; Ryali et al., 2024; Qiang et al., 2021). However, many dealt with data in
076 an abstract node format and did not respect the spatial topography in functional connectivity in the
077 anatomical space. Moreover, most prior work tested the model with performance on a specific task
078 (e.g. disease classification(Qiang et al., 2021; Zhao et al., 2020)) or sex classification (Ryali et al.,
079 2024) rather than a range of metrics as we suggested here. The majority of them use data samples as
080 the complete functional connectome (one per subject) rather than the functional connectivity seed
081 maps from each node, as we propose here. Moreover, to our knowledge, only limited prior research
082 has investigated the use of an unsupervised method to embed fMRI data into a latent space (Caro
083 et al., 2023; Kim et al., 2021) for general-purpose visualization and analysis, or demonstrated the
084 effect of traversing the latent space. None of these studies have applied such methods to functional
085 connectivity seed maps or used an extremely low-dimensional bottleneck (2, 4, and 32 dimensions)
086 as we do here, and only one tried to generalize to a different dataset(Caro et al., 2023).

087 Here, we aim to find a pre-computed transformation that maps new functional connectivity seed
088 map data onto a low-dimensional space for visualization and analysis. This approach is easy to
089 apply and addresses the correspondence dilemma across node definitions. Furthermore, analyses
090 such as clustering can benefit from the computational advantage of a low-dimensional embedding
091 space (Langs et al., 2016; Zhakubayev & Hamerly, 2022) (more discussion on this in Appendix
092 A.2).

093 3 METHODS

094 3.1 NEUROIMAGING DATA

095 We used resting-state fMRI data from the Washington University 120 (WU120) (Power et al., 2014)
096 and the Human Connectome Project (HCP) (Glasser et al., 2016; Van Essen et al., 2012b) datasets
097 for training and testing the models. Both datasets were collected from young adult subjects (19-35
098 years) while they were asked to fixate on a center cross on the screen in a 3-Tesla MRI scanner.
099 Procedures were then applied to normalize intensity, correct for motion in the scanner, and trans-
100 form the data onto a standard 32k-fsLR surface (Van Essen et al., 2012a). Details are available in
101 the Appendix A.3. The WU120 dataset contains one resting-state session per subject for 120 sub-
102 jects. Out of the 120 subjects from WU120, 100 subjects were selected as the training data, 10 as
103 the validation data, and 10 as the test data. The HCP dataset contains 94 subjects from unrelated
104 families, with two resting-state sessions for each subject. We used both sessions (Rest1 and Rest2)
105 in the HCP dataset as the test data.

3.1.1 PREPARATION OF THE FUNCTIONAL CONNECTOME DATA

First, to make the functional connectome comparable across different nodal definitions, we used a functional connectivity seed map, where the values at each location in the cortical vertex represent its connectivity with a given seed region. This was calculated as the Pearson’s correlation between the BOLD time series of the seed region and the BOLD time series at each vertex in the cerebral cortex surface ($N = 59412$ in the standard 32k-fsLR surface). This results in a matrix of N_{nodes} seed maps, each with a dimension of 59412×1 . To encourage diversity and variability in the training data while keeping computational demand manageable, we used the cortical vertices as seed regions ($N_{\text{nodes}} = 59412$ per subject) and randomly sampled 10% of the seed maps for each subject (total training samples = 591200). In the test data, we used individual parcels from the Gordon parcellation (Gordon et al., 2016) as seed regions ($N_{\text{nodes}} = 333$ per subject) for each subject.

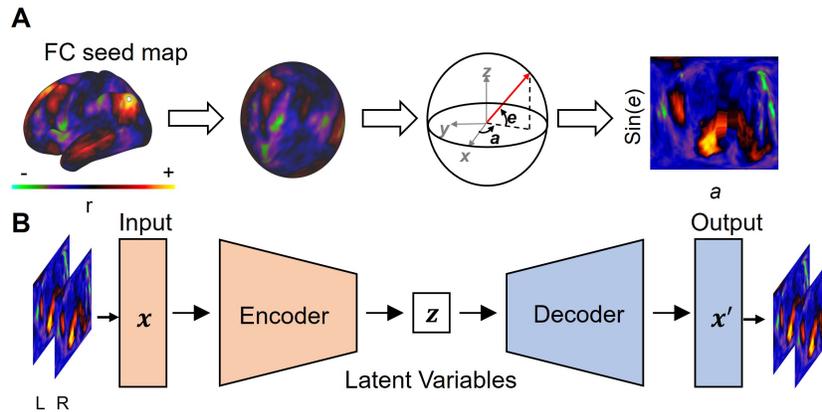


Figure 1: **Geometric reformatting and model architecture.** A) Geometric reformatting. The cortical distribution of fMRI activity is converted onto a spherical surface and then to an image by evenly resampling the spherical surface with respect to $\sin(e)$ and a , where e and a indicate elevation and azimuth, respectively. B) Architecture of VAE. Simplified block diagram of VAE model (upper panel in B). An encoder network samples latent variables given an input image under the inference model while a decoder network generates a genuine input image from under the generative model. Details of VAE model (bottom panel in B). The encoder and decoder networks both contain 5 convolutional layers and 1 fully-connected layer. Each input is a pair of 2-D images (left and right hemispheres). Adapted with permission from Kim, J., Zhang, Y., Han, K., Wen, Z., Choi, M., & Liu, Z. (2021). Representation learning of resting state fMRI with variational autoencoder. *NeuroImage*, 241, 118423. Copyright 2021 by Elsevier Inc.

For the autoencoder-based models, the input was passed on as a pair of 2-D images (left and right hemispheres) to take advantage of the convolutional layers of the model (Figure 1A, see details in Appendix A.4). This transformation would result in a small distortion such that the reconstructed 59412×1 seed maps from this 192×192 grid explain 99.9% variance from the original seed maps. To ensure fairness across algorithms, we passed the reconstructed seed maps (with distortion) instead of the original seed map from the 192×192 grid to PCA and ICA.

3.2 DIMENSIONALITY REDUCTION METHODS

3.2.1 AUTOENCODER-BASED METHODS

Autoencoders are neural networks designed to encode the input into a compressed representation, and then decode it back to a reconstructed input similar to the original one. The variational autoencoders (VAE) (Kingma & Welling, 2022) are especially useful in obtaining a smooth, continuous latent space for generating new data. The β -VAE variant (Higgins et al., 2017) is most effective in disentangling latent generative factors from images. Unlike t-SNE and diffusion maps, the AE-based models feature an intrinsic encoder for easy new data embedding and a decoder for data reconstruction. We adopted the same model architecture as described in (Kim et al., 2021; 2023; 2024) with five convolutional layers and one fully-connected layer in the encoder, and one fully-connected layer

and five convolutional layers in the decoder. The encoder transformed an fMRI seed map (a pair of left and right hemisphere images formatted to two 192 x 192 grids) into a probabilistic distribution of N latent variables (or a deterministic N latent variables in the case of a conventional autoencoder). Each convolutional layer conducted linear convolutions followed by rectifying the outputs as described by (Nair & Hinton, 2010). The first layer utilized 8×8 convolutions on inputs from each hemisphere and combined the results. Subsequent layers, from the second to the fifth, applied 4×4 convolutions to this combined output. Circular padding was employed at the azimuth boundaries, while zero padding was used at the elevation boundaries. A fully connected layer applied linear weighting to generate the mean and standard deviation for the distribution of each latent variable. The decoder replicated this structure in reverse, reconnecting the layers to recreate the fMRI seed map from a sample latent variable. Additional details on the model design and various autoencoder-based models can be found in Appendix A. 5 The VAE variant was optimized to reconstruct the input while constraining the distribution of every latent variable to be close to an independent and standard normal distribution. This is achieved by optimizing the encoding parameters, ϕ , and the decoding parameters, θ , to minimize the loss function below:

$$L(\phi, \theta | x) = \|x - x'\|_2 + \beta \cdot D_{KL} [N(\mu_z, \sigma_z) \| N(0, I)] \quad (1)$$

where x is the input data from both hemispheres, x' is the reconstructed data, and $N(\mu_z, \sigma_z)$ is the posterior distribution, $N(0, I)$ is the prior distribution. D_{KL} measures the Kullback-Leibler (KL) divergence between the posterior and prior distributions, and β is a hyperparameter balancing the two terms in the loss function. A $\beta < 1$ places less emphasis on the KL divergence and focuses more on reconstruction, while a $\beta > 1$ places a higher emphasis on KL divergence, enforcing stricter regularization of the latent space. When $\beta = 0$, the loss function only depends on the reconstruction error and is thus the conventional autoencoder. Instead of using the KL-divergence loss for regularization, the adversarial autoencoder used a separate discriminator network (3-layers, with the first two with a leaky ReLU activation function with a negative slope angle of 0.2 and the last one with a sigmoid activation function) for regularization (Makhzani et al., 2016).

Models were trained with stochastic gradient descent with a batch size of 128, initial learning rate of 1×10^{-4} , and 50 epochs with random data selection in each batch. An Adam optimizer (Kingma & Ba, 2014) was implemented, and the learning rate decayed by a factor of 10 every 20 epochs for the AE and VAE models. Final hyperparameters including the number of latent dimensions and the beta value were determined by the trade-off between KL divergence and reconstruction loss on the validation data (See Appendix A.6). The model was trained in Python 3.8 using PyTorch (v2.1.2+cu118) using a server with an NVIDIA A100 GPU (40 GB memory). A brief demonstration of model performance variability using different subsets of training data is provided in Appendix A.7.

3.2.2 LINEAR DIMENSIONALITY REDUCTION METHODS

We chose principal component analysis (PCA) and independent component analysis (ICA) as alternative linear dimensionality reduction methods. PCA uses Singular Value Decomposition of the data, keeping only the most significant singular vectors to project the data to a lower dimensional space. ICA attempts to decompose the data into a set of independent spatial maps. For memory management, we performed PCA with incremental PCA (IPCA) from the Scikit-learn package (v1.3.2) in Python 3.8 (Golub & Loan, 2013; Ross et al., 2008). IPCA builds a low-rank approximation for the input data using a memory amount that is independent of the number of samples. ICA was performed on the reduced data using the first 100 PCs with FastICA in Scikit-learn.

3.3 ASSESSING THE QUALITY OF EMBEDDINGS

3.3.1 RECONSTRUCTION PERFORMANCE

Reconstruction performance was calculated with η^2 , namely, the fraction of variance in the original seed map accounted for by variance in the reconstructed seed map on a point-by-point basis (Cohen et al., 2008). η^2 ranged from 0 (no similarity) to 1 (identical) and is formally defined by:

$$\eta^2 = 1 - \frac{SS_{\text{Within}}}{SS_{\text{Total}}} = 1 - \frac{\sum_{i=1}^n [(a_i - m_i)^2 + (b_i - m_i)^2]}{\sum_{i=1}^n [(a_i - \bar{M})^2 + (b_i - \bar{M})^2]} \quad (2)$$

where a_i and b_i represent the values at position i in maps a and b , respectively. μ_i is the mean value of the two images at position i , $(a_i + b_i)/2$. \bar{M} is the grand mean value across the mean image m .

Self-connectivity at all positions was excluded. This similarity matrix is sensitive to the difference in a and b in scales and offsets. For convenience, we used the inversely formatted seed maps from the 2D image in Figure 1A as the reference ground truth data ($\eta^2 > 0.99$ to the original data) to be compared with the reconstructed data from latent representations. Note that perfect reconstruction is not necessarily desirable because the original data might be noisy, and the reconstructed data could potentially represent a “denoised” version of the data. We calculated two additional reference measures: the first one is the η^2 between the ground truth data in HCP Rest1 and HCP Rest2 of the same subject, which provides the noise ceiling of the reconstruction; the second one is the η^2 between each subject to the rest 93 subjects in each session (averaged across the Rest1 and Rest2 sessions), which provides the null baseline of group average.

3.3.2 SEPARATION BY CANONICAL FUNCTIONAL NETWORKS

One key feature of functional connectivity is that nodes within the same functional network tend to possess similar seed maps (Yeo et al., 2011). The 333 parcels in the Gordon parcellation were grouped into 12 functional networks and one additional group (named “None”) of 47 parcels which covered the low-SNR regions and cannot be confidently grouped into any of the 12 functional networks and omitted from this analysis (Gordon et al., 2016). We evaluated the segregation of seed maps from different functional networks with the silhouette index (SI) (Rousseeuw, 1987; Yeo et al., 2011), calculated as:

$$SI(i) = \frac{b_i - a_i}{\max(a_i, b_i)} \quad (3)$$

where b_i is the mean Euclidean distance between the current parcel i to the parcels in the best alternative network, and a_i is the mean within-network Euclidean distance of the latent. A 95% confidence interval was calculated by bootstrapping the individuals 1000 times.

3.3.3 INTRASUBJECT AND INTERSUBJECT VARIABILITY

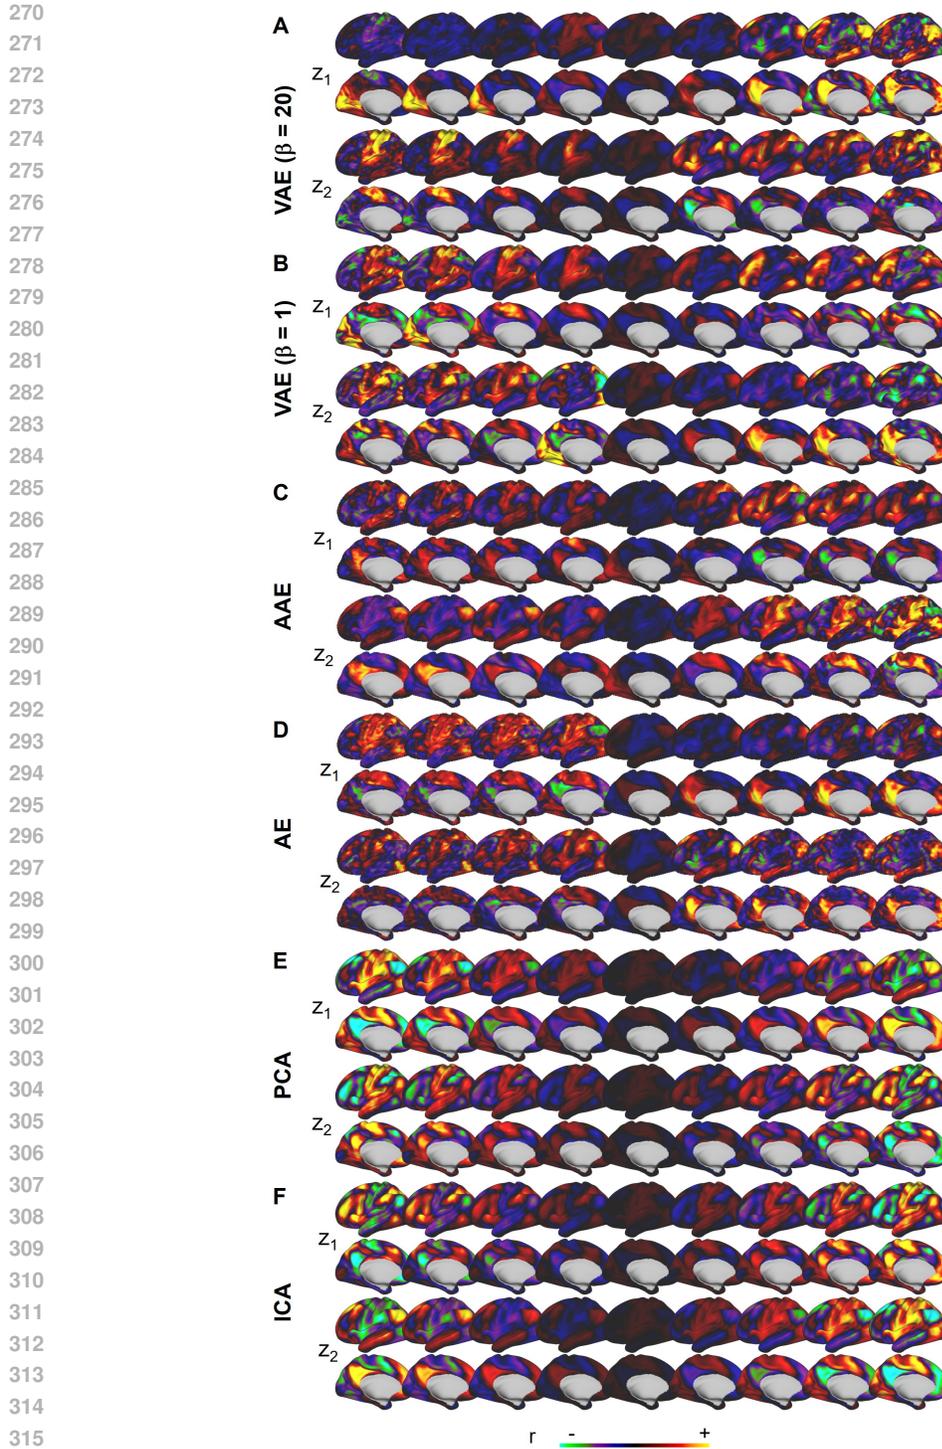
To capture individual differences in functional network assignments (Gordon et al., 2017a;b), we ran a k-means clustering algorithm on the latent embeddings with the cluster centroid initialized to be the center of each of the 12 Gordon networks (Gordon et al., 2016). 12 clusters were optimized by minimizing Euclidean distances between latent embeddings (dimensions = 2, 4, 32) within a cluster. Intersubject variability was calculated as the average normalized Hamming distance (a.k.a. proportion of mismatches, with 0 indicating perfect match) from each subject to other subjects; the intrasubject variability was calculated as the average normalized Hamming distance between the two resting sessions of the same subject. The latent representations that capture the most individual differences need to exhibit low intrasubject reliability while simultaneously maintaining a high intersubject variability. We calculated the variability signal-to-noise ratio (vSNR) to quantify the relative magnitude of the two sources of variability (Langs et al., 2016). Specifically:

$$vSNR = \frac{\text{Intersubject variability}}{\text{Intrasubject variability}} - 1 \quad (4)$$

4 RESULTS

4.1 SINGLE LATENT TRAVERSAL

To understand how the changes in the magnitude of each latent dimension affect the reconstructed seed map appearance, we plotted the single latent traversals when the number of dimensions = 2 (Figure 2). We varied the magnitude of one latent dimension while keeping the other latent dimension fixed at zero. We observed patterns reminiscent of sensorimotor networks and association networks (Margulies et al., 2016; Sydnor et al., 2021), as well as task-positive to task-negative networks (Buckner et al., 2008; Fox et al., 2005; Raichle, 2015). Notably, reconstructed maps resembling distinct somatomotor and visual networks were only observable in VAE ($\beta = 20$). We also observed a gradual transition from somatomotor hand to somatomotor mouth networks when varying a single dimension in VAE ($\beta = 20$). The conventional AE has similar profiles in its first and second latent dimensions, suggesting an ineffective disentangling of generative factors. In summary, a single latent dimension in VAE ($\beta = 20$) seems to disentangle the most subtle details in functional connectivity.



317 Figure 2: Single latent traversal. The reconstructed seed maps from latent values of one dimension
318 varied in equal steps from one end to the other end and the other dimension was fixed to 0. A) VAE
319 ($\beta = 20$), B) VAE ($\beta = 1$), C) AAE, D) AE, E) PCA, F) ICA.

320 4.2 RECONSTRUCTION PERFORMANCE

321

322 Next, we examined the reconstruction performance for parcel seed maps (functional connectivity
323 from an area parcel defined from an atlas (Gordon et al., 2016)) to examine how effectively different
dimensionality reduction methods retain information. 10 test subjects in the WU120 dataset and

both Rest1 and Rest2 sessions for the 94 subjects in the HCP dataset were chosen to test the out-of-sample and out-of-distribution generalization, respectively. Figure 3A shows reconstructed seed maps from one example parcel (parcel 15 in the medial visual cortex). Overall, the mean reconstruction performance across all parcels in each individual was similar across methods, with AE-based latent representations providing better reconstruction performance when the latent representation had only 2 dimensions (Figure 3B) and linear methods providing marginally better reconstruction performance at 32 dimensions (Figure 3D). In all cases, the reconstruction performance from the latent representations was on average higher than the null baseline, suggesting that the latent representations captured individual-specific features in addition to group-average features. With 32 dimensions, the reconstruction performance was approaching the noise ceiling (for the normalized reconstruction performance relative to the individuals’ noise ceiling see Appendix A.9).

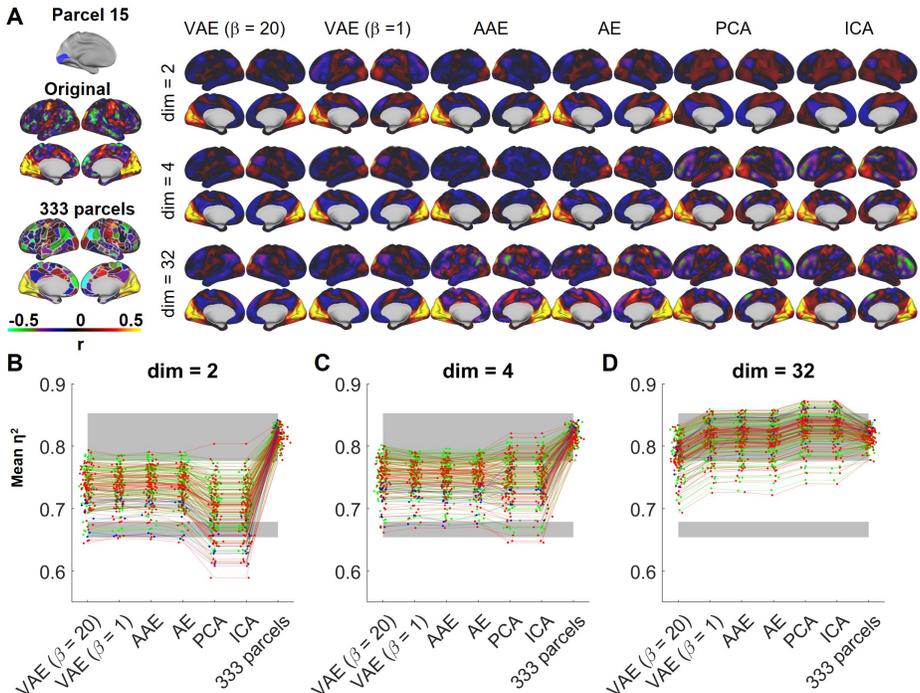


Figure 3: **Reconstruction performance in the test set.** A) Visualization of the original and reconstructed seed maps from parcel 15 in an example subject (subject 111). B-D) The reconstruction performance for all subjects in the HCP and WU120 test datasets. Reconstruction Green line: HCP Rest1, Red line: HCP Rest2, Blue line: WU120. Gray shaded area: mean and standard deviation of the noise ceiling and the null baseline. N.B.: 333 parcels always have 333 dimensions and were repeatedly displayed in all three panels as a reference. VAE = variational autoencoder. AAE = adversarial autoencoder. AE = autoencoder. PCA = principal component analysis. ICA = independent component analysis.

4.3 SEPARATION BY CANONICAL FUNCTIONAL NETWORKS

Seed maps from area parcels belonging to the same functional network should be similar to each other. We asked whether the functional network clusters defined in adult group-average data (Gordon et al., 2016) were preserved in the low-dimensional embeddings (examples for dimension = 2 in Figure 4). We quantified the distinguishability from the best alternative network for each of the 333 parcels with the silhouette index (SI) for the average latent embeddings of Rest1 sessions seed maps across 94 HCP subjects. First, we found the separation by functional networks in the parcellated functional connectome (dimension = 333, SI = 0.158, see Appendix A.10). Even at dimension =2, functional network segregation in the latent space was evident across all dimensionality reduction methods, particularly in the separation of sensorimotor networks from association networks (with the exception of AE). However, the distinction within the sensorimotor networks (visual, somato-

378 motor, auditory) was far more pronounced in the VAE and AAE latent spaces, especially with VAE
 379 ($\beta = 20$). Another interesting observation was that the default mode network (colored red in Figure
 380 4A) appeared to separate into two distinct clusters for the VAE ($\beta = 20$) latent space (See Appendix
 381 A.11). At two dimensions, the average SI for VAE ($\beta = 20$) was 0.086 (95% CI: [0.066, 0.099]),
 382 higher than the VAE ($\beta = 1$) (0.044, 95% CI: [0.015, 0.063]), AAE (0.035, 95% CI: [0.006, 0.054]),
 383 AE (0.014, [-0.018, 0.030]), PCA (7×10^{-4} , 95% CI: [-0.016, 0.009]) and ICA (-0.010, 95% CI:
 384 [-0.028, 0.001]). When the latent dimensions were 4 or 32, the average SI can sometimes be higher
 385 than that in the parcel connectome (Table 2). Among the models tested, VAE ($\beta = 20$) with four di-
 386 mensions (SI = 0.213, 95% CI: [0.197, 0.221]) best separates the parcels according to the functional
 387 network labels (A.10). Notably, functional connectivity networks have hierarchical organization
 388 (Betzel & Bassett, 2017; Urchs et al., 2019), and the current functional network labels (Figure 4A)
 389 only represent one popular scale of investigation. The reduction in SI from dimension = 4 to di-
 390 mension = 32 for some methods (A.10) might be due to the more pronounced separation of network
 391 sub-components at a relatively higher-dimension latent space.

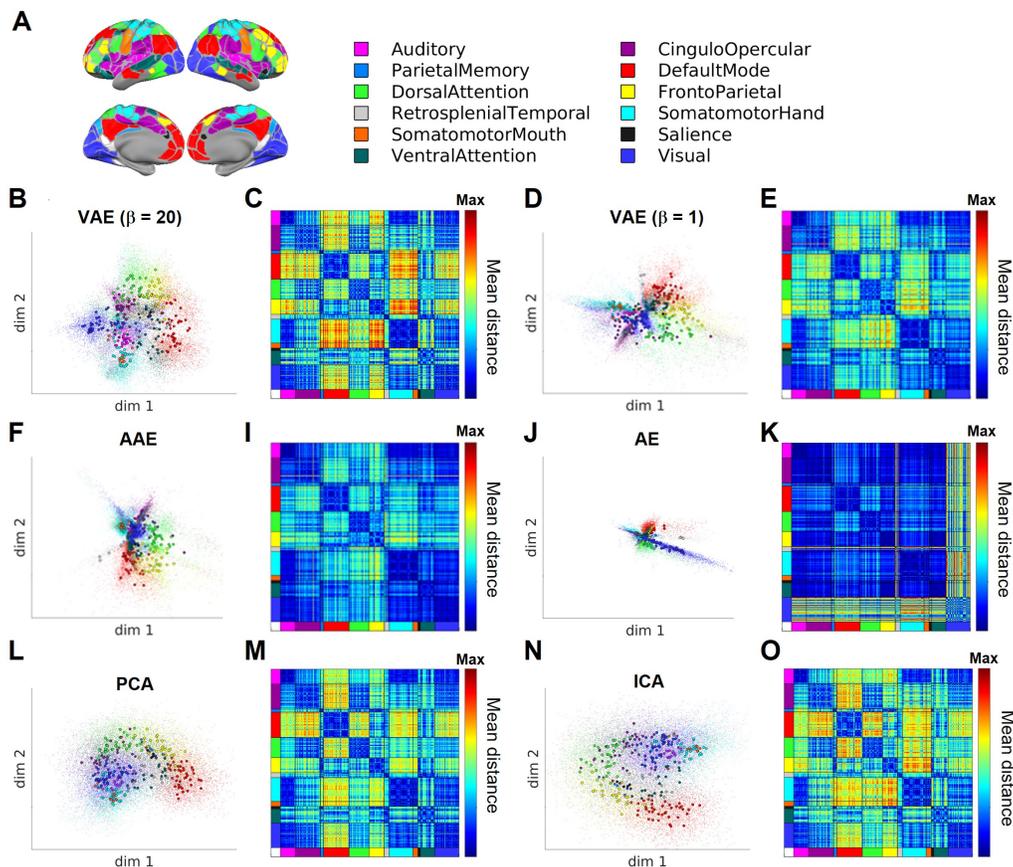


Figure 4: **Separation by canonical functional networks.** A) Gordon network assignments. B/D/F/J/L/N) The functional connectome 2-dimensional embeddings. Small dots represent data from individual subjects and large circles represent data averaged across 94 subjects. C/E/I/K/M/O) The mean Euclidean distance between the latent representations of the average across 94 subjects. VAE = variational autoencoder. AAE = adversarial autoencoder. AE = autoencoder. PCA = principal component analysis. ICA = independent component analysis.

4.4 INTRASUBJECT AND INTERSUBJECT VARIABILITY

Despite the largely consistent topography of functional networks (Damoiseaux et al., 2006; Gratton et al., 2018), individual differences in functional network assignment are evident in prior studies (Gordon et al., 2017b;a; Gratton et al., 2018). We tested how well the latent representations preserve

these individual differences. First, we examined the ability of the low-dimensional representation to capture the intersubject variability while preserving the similarity of intrasubject data in individual parcels. We used an example parcel (parcel 213) with a large variability. We found that the relative positions of the embeddings for different subjects were consistent across the two resting scan sessions (Rest1 and Rest2) in the VAE ($\beta = 20$) latent space (Figure 5). The two subjects at the extremes demonstrated very distinct seed map profiles (Figure 5). Next, we applied a k-means clus-

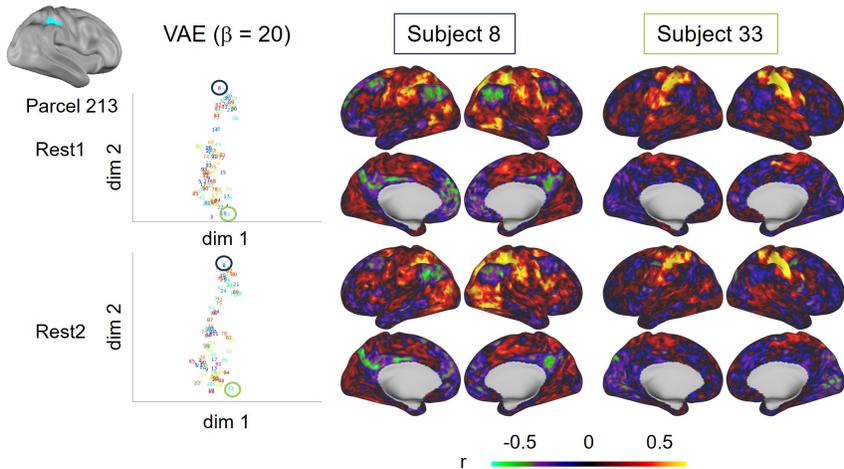


Figure 5: **Interindividual variability of an example parcel across sessions across 94 subjects in the VAE ($\beta = 20$) latent representation, at dimension = 2.** Each data point is a subject indicated by the numbers 1 to 94. The seed maps from parcel 213 for subject 8 and subject 33 in Rest1 and Rest2 were shown as an example.

tering with $k = 12$ to the low-dimensional embeddings for the 94 subjects in HCP across Rest1 and Rest2 sessions. We calculated the intrasubject and intersubject variability using Hamming distance (proportion of mismatch in the 333 parcel assignments) and vSNR as described in section 3.3.3. We found that independent of the dimensionality reduction method used, the intersubject variability was higher than the intrasubject variability by about 50%, suggestive of the preservation of individual specificity in the functional connectome in the latent spaces (See Table 3 in Appendix A.12). Across all subjects, the within-subject network assignment was more similar than the between-network subject assignment (Figure 6, also see Appendix A.12). vSNR was among the highest for VAE ($\beta = 20$) across 2, 4 and 32 dimensions. Furthermore, we tested whether different sessions of the same individual can be successfully identified from a list of subjects and found that the identification accuracy using the latent embeddings at various dimensionalities closely approximated the fingerprinting accuracy using the parcel connectome, especially when the scan length was long enough (> 20 min). One exception is the AE embeddings which had a low identification accuracy (Details in Appendix A.13). Similarly, the prediction performance of age and sex from the different embeddings was similar and increased with the number of dimensions (Appendix A.14).

5 DISCUSSION

This work presents the learning of mappings of functional connectivity seed maps from vertex space to a low-dimensional latent space using data from one young adult dataset. Among different performance benchmarks, including the reconstruction performance, separation by functional networks, and the retention of individual specificity, the linear dimensionality-reduction methods PCA and ICA achieved similar performance to the nonlinear AE-based dimensionality-reduction methods, especially at higher dimensions. This is consistent with prior research finding comparable functional connectivity gradients with PCA, Laplacian eigenmaps, and diffusion mapping (Vos de Wael et al., 2020), suggesting that the functional connectivity data likely has a predominantly linear structure. Other supervised machine learning approaches also found comparable performance with deep neural networks and kernel regression (He et al., 2018), suggesting that more complex models are not always better. However, with only 2 dimensions, the VAE ($\beta = 20$) latent embeddings best and

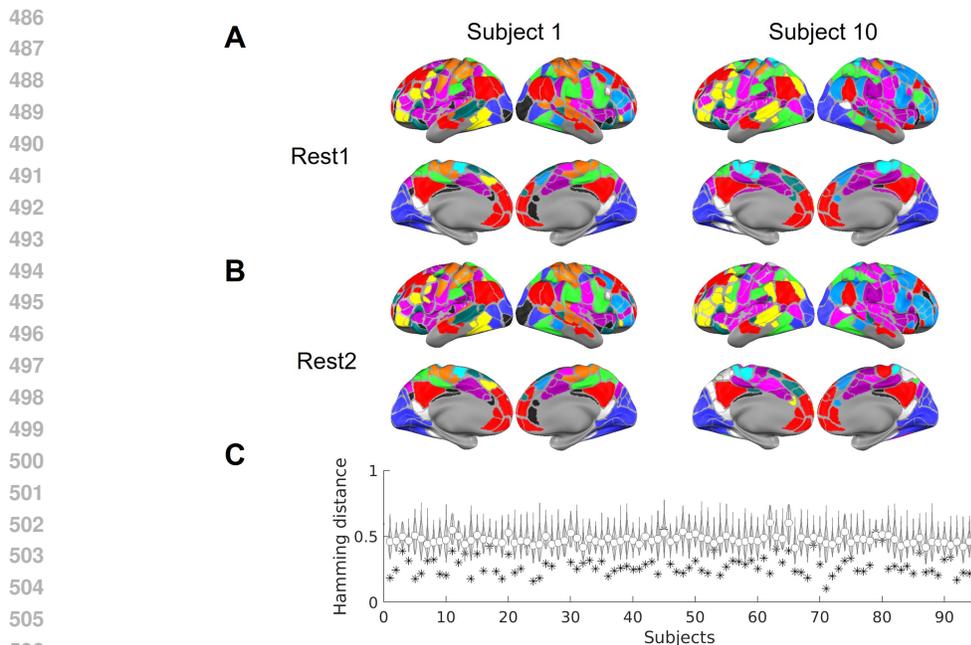


Figure 6: **Individual-specific network using k-means clustering with 12 networks (VAE beta = 20, dimension = 32).** The functional network assignment for Rest1 and Rest2 sessions based on k-means clustering on latent representations across of all parcels in all subjects and sessions for A) subject 1 and B) subject 10. C) The normalized Hamming distance for between-subject network assignments (violins) and within-subject network assignments (asterisks).

separates the fine nuisances across functional networks. It may make a useful visualization tool for simultaneously comparing the seed map profiles across sessions, individuals, and states independent of the definition of parcels. This could enable analyses such as community detection and behavioral phenotype prediction to be conducted in a much more manageable space, and allow efficient data sharing for big data initiatives (Horien et al., 2021).

We infer that our learned mappings will apply to new unseen datasets based on two lines of reasoning: 1) we demonstrated that the mappings were generalizable beyond the original training data, despite the differences in acquisition procedures and subjects; 2) extensive prior literature has observed low dimensionality in functional connectivity data (Gotts et al., 2020; Margulies et al., 2016; Snyder et al., 2022) with prominent patterns conserved across development (Dong et al., 2021; Xia et al., 2022). However, future studies could explicitly test this assumption in other populations, especially in developmental and clinical cohorts.

6 LIMITATIONS AND FUTURE WORK

We acknowledge that our training data is a small sample with a narrow profile of demographics. The training dataset could be extended to incorporate diverse demographics, acquisition parameters, developmental stages, etc. However, with the continuous, regularized latent space learned by a VAE, it is reasonable to assume that even if the exact seed map does not exist in the original dataset, as long as they are dominated by the same generative factors (as expected by the constraints of physical and biological properties), the latent space representation mapped using our existing model would still retain key information of the new dataset. Additionally, the VAE latent space lacks interpretability, unlike other methods that explore harmonic modes/eigenmodes in the brain based on geometric anatomy. Therefore, our model is a descriptive/phenomenological model rather than a mechanistic model. Moreover, our framework targets resting-state functional connectivity analyses and overlooks the temporal information in the neural data. Similar methods of representational learning of temporal information in the neural data exist in the literature (Caro et al., 2023; Kim et al., 2021), and we explained our motivation for embedding the functional connectivity data instead of time series data with more detailed comparison in appendix A.15.

REFERENCES

- Salim Arslan, Sofia Ira Ktena, Antonios Makropoulos, Emma C. Robinson, Daniel Rueckert, and Sarah Parisot. Human brain mapping: A systematic comparison of parcellation methods for the human cerebral cortex. *NeuroImage*, 170:5–30, April 2018. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2017.04.014. URL <https://www.sciencedirect.com/science/article/pii/S1053811917303026>.
- D. S. Bassett, N. F. Wymbs, M. A. Porter, P. J. Mucha, J. M. Carlson, and S. T. Grafton. Dynamic reconfiguration of human brain networks during learning. *Proceedings of the National Academy of Sciences*, 108(18):7641–7646, May 2011. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1018985108. URL <http://www.pnas.org/cgi/doi/10.1073/pnas.1018985108>.
- Danielle S Bassett and Olaf Sporns. Network neuroscience. *Nature neuroscience*, 20(3):353–364, February 2017. ISSN 1097-6256. doi: 10.1038/nn.4502. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5485642/>.
- Richard F. Betzel and Danielle S. Bassett. Multi-scale brain networks. *NeuroImage*, 160:73–83, October 2017. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2016.11.006. URL <https://www.sciencedirect.com/science/article/pii/S1053811916306152>.
- Richard F. Betzel, Lisa Byrge, Ye He, Joaquín Goñi, Xi-Nian Zuo, and Olaf Sporns. Changes in structural and functional connectivity among resting-state networks across the human lifespan. *NeuroImage*, 102:345–357, November 2014. ISSN 10538119. doi: 10.1016/j.neuroimage.2014.07.067. URL <https://linkinghub.elsevier.com/retrieve/pii/S1053811914006508>.
- Richard F. Betzel, Maxwell A. Bertolero, Evan M. Gordon, Caterina Gratton, Nico U. F. Dosenbach, and Danielle S. Bassett. The community structure of functional brain networks exhibits scale-specific patterns of inter- and intra-subject variability. *NeuroImage*, 202, November 2019. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2019.07.003. URL <https://www.sciencedirect.com/science/article/pii/S1053811919305658>.
- Janine Bijsterbosch, Samuel J. Harrison, Saad Jbabdi, Mark Woolrich, Christian Beckmann, Stephen Smith, and Eugene P. Duff. Challenges and future directions for representations of functional brain organization. *Nature Neuroscience*, 23(12):1484–1495, December 2020. ISSN 1546-1726. doi: 10.1038/s41593-020-00726-z. URL <https://www.nature.com/articles/s41593-020-00726-z>.
- Bharat Biswal, F. Zerrin Yetkin, Victor M. Haughton, and James S. Hyde. Functional connectivity in the motor cortex of resting human brain using echo-planar mri. *Magnetic Resonance in Medicine*, 34(4):537–541, 1995. ISSN 1522-2594. doi: 10.1002/mrm.1910340409. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.1910340409>.
- Taylor Bolt, Jason S. Nomi, Danilo Bzdok, Jorge A. Salas, Catie Chang, B. T. Thomas Yeo, Lucina Q. Uddin, and Shella D. Keilholz. A parsimonious description of global functional brain organization in three spatiotemporal patterns. *Nature Neuroscience*, 25(8):1093–1103, August 2022. ISSN 1097-6256, 1546-1726. doi: 10.1038/s41593-022-01118-1. URL <https://www.nature.com/articles/s41593-022-01118-1>.
- Randy L. Buckner, Jessica R. Andrews-Hanna, and Daniel L. Schacter. The brain’s default network. *Annals of the New York Academy of Sciences*, 1124(1):1–38, 2008. ISSN 1749-6632. doi: 10.1196/annals.1440.011. URL <https://onlinelibrary.wiley.com/doi/abs/10.1196/annals.1440.011>.
- Ed Bullmore and Olaf Sporns. Complex brain networks: Graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, 10(3):186–198, March 2009. ISSN 1471-003X, 1471-0048. doi: 10.1038/nrn2575. URL <http://www.nature.com/articles/nrn2575>.
- Gregory C. Burgess, Sridhar Kandala, Dan Nolan, Timothy O. Laumann, Jonathan D. Power, Babatunde Adeyemo, Michael P. Harms, Steven E. Petersen, and Deanna M. Barch. Evaluation

- 594 of denoising strategies to address motion-correlated artifacts in resting-state functional mag-
595 netic resonance imaging data from the human connectome project. *Brain Connectivity*, 6(9):
596 669–680, November 2016. ISSN 2158-0014. doi: 10.1089/brain.2016.0435. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5105353/>.
- 598 Josue Ortega Caro, Antonio H. de O. Fonseca, Christopher Averill, Syed A. Rizvi, Matteo Rosati,
599 James L. Cross, Prateek Mittal, Emanuele Zappala, Daniel Levine, Rahul M. Dhodapkar, Chadi G.
600 Abdallah, and David van Dijk. Brainlm: A foundation model for brain activity recordings,
601 September 2023. URL [https://www.biorxiv.org/content/10.1101/2023.09.](https://www.biorxiv.org/content/10.1101/2023.09.12.557460v1)
602 [12.557460v1](https://www.biorxiv.org/content/10.1101/2023.09.12.557460v1).
- 603 Dionysios G. Cheirdaris. Graph theory-based approach in brain connectivity modeling and
604 alzheimer’s disease detection. In Panagiotis Vlamos (ed.), *GeNeDis 2022*, pp. 49–58,
605 Cham, 2023. Springer International Publishing. ISBN 978-3-031-31982-2. doi: 10.1007/
606 978-3-031-31982-2_5.
- 607 Rastko Ciric, Daniel H. Wolf, Jonathan D. Power, David R. Roalf, Graham L. Baum, Kosha Ru-
608 parel, Russell T. Shinohara, Mark A. Elliott, Simon B. Eickhoff, Christos Davatzikos, Ruben C.
609 Gur, Raquel E. Gur, Danielle S. Bassett, and Theodore D. Satterthwaite. Benchmarking of
610 participant-level confound regression strategies for the control of motion artifact in studies of
611 functional connectivity. *NeuroImage*, 154:174–187, July 2017. ISSN 1053-8119. doi: 10.
612 1016/j.neuroimage.2017.03.020. URL [https://www.sciencedirect.com/science/](https://www.sciencedirect.com/science/article/pii/S1053811917302288)
613 [article/pii/S1053811917302288](https://www.sciencedirect.com/science/article/pii/S1053811917302288).
- 614 Alexander L. Cohen, Damien A. Fair, Nico U. F. Dosenbach, Francis M. Miezin, Donna Dierker,
615 David C. Van Essen, Bradley L. Schlaggar, and Steven E. Petersen. Defining functional areas
616 in individual human brains using resting functional connectivity mri. *NeuroImage*, 41(1):45–57,
617 May 2008. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2008.01.066. URL [https://www.](https://www.sciencedirect.com/science/article/pii/S1053811908001171)
618 [sciencedirect.com/science/article/pii/S1053811908001171](https://www.sciencedirect.com/science/article/pii/S1053811908001171).
- 619 R. Cameron Craddock, G. Andrew James, Paul E. Holtzheimer III, Xiaoping P. Hu, and Helen S.
620 Mayberg. A whole brain fmri atlas generated via spatially constrained spectral clustering. *Human*
621 *Brain Mapping*, 33(8):1914–1928, 2012. ISSN 1097-0193. doi: 10.1002/hbm.21333. URL
622 <https://onlinelibrary.wiley.com/doi/abs/10.1002/hbm.21333>.
- 623 J. S. Damoiseaux, S. a. R. B. Rombouts, F. Barkhof, P. Scheltens, C. J. Stam, S. M. Smith, and C. F.
624 Beckmann. Consistent resting-state networks across healthy subjects. *Proceedings of the National*
625 *Academy of Sciences*, 103(37):13848–13853, September 2006. ISSN 0027-8424, 1091-6490. doi:
626 10.1073/pnas.0601417103. URL <https://www.pnas.org/content/103/37/13848>.
- 627 Manlio de Domenico. Multilayer modeling and analysis of human brain networks. *GigaScience*,
628 6(5):1, February 2017. doi: 10.1093/gigascience/gix004. URL [https://pmc.ncbi.nlm.](https://pmc.ncbi.nlm.nih.gov/articles/PMC5437946/)
629 [nih.gov/articles/PMC5437946/](https://pmc.ncbi.nlm.nih.gov/articles/PMC5437946/).
- 630 Hao-Ming Dong, Daniel S. Margulies, Xi-Nian Zuo, and Avram J. Holmes. Shifting gradi-
631 ents of macroscale cortical organization mark the transition from childhood to adolescence.
632 *Proceedings of the National Academy of Sciences*, 118(28):e2024448118, July 2021. doi:
633 10.1073/pnas.2024448118. URL [https://www.pnas.org/doi/abs/10.1073/pnas.](https://www.pnas.org/doi/abs/10.1073/pnas.2024448118)
634 [2024448118](https://www.pnas.org/doi/abs/10.1073/pnas.2024448118).
- 635 Simon B. Eickhoff, Bertrand Thirion, Gaël Varoquaux, and Danilo Bzdok. Connectivity-based par-
636 cellation: Critique and implications. *Human Brain Mapping*, 36(12):4771–4792, 2015. ISSN
637 1097-0193. doi: 10.1002/hbm.22933. URL [https://onlinelibrary.wiley.com/](https://onlinelibrary.wiley.com/doi/abs/10.1002/hbm.22933)
638 [doi/abs/10.1002/hbm.22933](https://onlinelibrary.wiley.com/doi/abs/10.1002/hbm.22933).
- 639 Damien A Fair. Correction of respiratory artifacts in mri head motion estimates. *NeuroImage*, pp.
640 17, 2020.
- 641 Emily S. Finn, Xilin Shen, Dustin Scheinost, Monica D. Rosenberg, Jessica Huang, Marvin M.
642 Chun, Xenophon Papademetris, and R. Todd Constable. Functional connectome fingerprint-
643 ing: Identifying individuals using patterns of brain connectivity. *Nature Neuroscience*, 18
644 (11):1664–1671, November 2015. ISSN 1546-1726. doi: 10.1038/nn.4135. URL [https://www.nature.com/articles/nn.4135!](https://www.nature.com/articles/nn.4135)
645 [https://www.nature.com/articles/nn.4135!](https://www.nature.com/articles/nn.4135)
646 [https://www.nature.com/articles/nn.4135!](https://www.nature.com/articles/nn.4135)
647 [https://www.nature.com/articles/nn.4135!](https://www.nature.com/articles/nn.4135)

- 648 Bruce Fischl. Freesurfer. *NeuroImage*, 62(2):774–781, August 2012. ISSN 1053-8119. doi: 10.
649 1016/j.neuroimage.2012.01.021. URL [https://www.sciencedirect.com/science/
650 article/pii/S1053811912000389](https://www.sciencedirect.com/science/article/pii/S1053811912000389).
651
- 652 Bruce Fischl, David H. Salat, Evelina Busa, Marilyn Albert, Megan Dieterich, Christian Hasel-
653 grove, Andre van der Kouwe, Ron Killiany, David Kennedy, Shuna Klaveness, Albert Montillo,
654 Nikos Makris, Bruce Rosen, and Anders M. Dale. Whole brain segmentation: Automated la-
655 beling of neuroanatomical structures in the human brain. *Neuron*, 33(3):341–355, January 2002.
656 ISSN 0896-6273. doi: 10.1016/S0896-6273(02)00569-X. URL [https://www.cell.com/
657 neuron/abstract/S0896-6273\(02\)00569-X](https://www.cell.com/neuron/abstract/S0896-6273(02)00569-X).
- 658 Michael D. Fox, Abraham Z. Snyder, Justin L. Vincent, Maurizio Corbetta, David C. Van Essen,
659 and Marcus E. Raichle. The human brain is intrinsically organized into dynamic, anticorrelated
660 functional networks. *Proceedings of the National Academy of Sciences*, 102(27):9673–9678, July
661 2005. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.0504136102. URL [https://pnas.
662 org/doi/full/10.1073/pnas.0504136102](https://pnas.org/doi/full/10.1073/pnas.0504136102).
- 663 Karl J. Friston, Steven Williams, Robert Howard, Richard S. J. Frackowiak, and Robert Turner.
664 Movement-related effects in fmri time-series. *Magnetic Resonance in Medicine*, 35(3):346–355,
665 1996. ISSN 1522-2594. doi: 10.1002/mrm.1910350312. URL [https://onlinelibrary.
666 wiley.com/doi/abs/10.1002/mrm.1910350312](https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.1910350312).
- 667 Matthew F. Glasser, Stamatiou N. Sotiropoulos, J. Anthony Wilson, Timothy S. Coalson, Bruce
668 Fischl, Jesper L. Andersson, Junqian Xu, Saad Jbabdi, Matthew Webster, Jonathan R. Polimeni,
669 David C. Van Essen, and Mark Jenkinson. The minimal preprocessing pipelines for the hu-
670 man connectome project. *NeuroImage*, 80:105–124, October 2013. ISSN 1053-8119. doi: 10.
671 1016/j.neuroimage.2013.04.127. URL [https://www.sciencedirect.com/science/
672 article/pii/S1053811913005053](https://www.sciencedirect.com/science/article/pii/S1053811913005053).
- 673 Matthew F. Glasser, Timothy S. Coalson, Emma C. Robinson, Carl D. Hacker, John Harwell, Essa
674 Yacoub, Kamil Ugurbil, Jesper Andersson, Christian F. Beckmann, Mark Jenkinson, Stephen M.
675 Smith, and David C. Van Essen. A multi-modal parcellation of human cerebral cortex. *Nature*,
676 536(7615):171–178, August 2016. ISSN 0028-0836, 1476-4687. doi: 10.1038/nature18933.
677 URL <http://www.nature.com/articles/nature18933>.
678
- 679 Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. JHU Press, February 2013. ISBN
680 978-1-4214-0859-0.
- 681 Evan M. Gordon, Timothy O. Laumann, Babatunde Adeyemo, Jeremy F. Huckins, William M.
682 Kelley, and Steven E. Petersen. Generation and evaluation of a cortical area parcellation from
683 resting-state correlations. *Cerebral Cortex*, 26(1):288–303, January 2016. ISSN 1047-3211,
684 1460-2199. doi: 10.1093/cercor/bhu239. URL [https://academic.oup.com/cercor/
685 article-lookup/doi/10.1093/cercor/bhu239](https://academic.oup.com/cercor/article-lookup/doi/10.1093/cercor/bhu239).
- 686 Evan M. Gordon, Timothy O. Laumann, Babatunde Adeyemo, and Steven E. Petersen. Individual
687 variability of the system-level organization of the human brain. *Cerebral Cortex*, 27(1):386–399,
688 January 2017a. ISSN 1047-3211. doi: 10.1093/cercor/bhv239. URL [https://doi.org/
689 10.1093/cercor/bhv239](https://doi.org/10.1093/cercor/bhv239).
- 690 Evan M. Gordon, Timothy O. Laumann, Adrian W. Gilmore, Dillan J. Newbold, Deanna J. Greene,
691 Jeffrey J. Berg, Mario Ortega, Catherine Hoyt-Drazen, Caterina Gratton, Haoxin Sun, Jacque-
692 line M. Hampton, Rebecca S. Coalson, Annie L. Nguyen, Kathleen B. McDermott, Joshua S.
693 Shimony, Abraham Z. Snyder, Bradley L. Schlaggar, Steven E. Petersen, Steven M. Nelson,
694 and Nico U.F. Dosenbach. Precision functional mapping of individual human brains. *Neuron*,
695 95(4), August 2017b. ISSN 08966273. doi: 10.1016/j.neuron.2017.07.011. URL [https:
696 //linkinghub.elsevier.com/retrieve/pii/S089662731730613X](https://linkinghub.elsevier.com/retrieve/pii/S089662731730613X).
697
- 698 Evan M. Gordon, Timothy O. Laumann, Scott Marek, Ryan V. Raut, Caterina Gratton, Dillan J.
699 Newbold, Deanna J. Greene, Rebecca S. Coalson, Abraham Z. Snyder, Bradley L. Schlaggar,
700 Steven E. Petersen, Nico U. F. Dosenbach, and Steven M. Nelson. Default-mode network streams
701 for coupling to language and control systems. *Proceedings of the National Academy of Sci-
ences of the United States of America*, 117(29):17308–17319, July 2020. ISSN 0027-8424. doi:

- 702 10.1073/pnas.2005238117. URL [https://www.ncbi.nlm.nih.gov/pmc/articles/](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7382234/)
703 [PMC7382234/](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7382234/).
704
- 705 Stephen J. Gotts, Adrian W. Gilmore, and Alex Martin. Brain networks, dimensionality, and
706 global signal averaging in resting-state fmri: Hierarchical network structure results in low-
707 dimensional spatiotemporal dynamics. *NeuroImage*, 205:116289, January 2020. ISSN 1053-
708 8119. doi: 10.1016/j.neuroimage.2019.116289. URL [https://www.ncbi.nlm.nih.gov/](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6919311/)
709 [pmc/articles/PMC6919311/](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6919311/).
- 710 Caterina Gratton, Timothy O. Laumann, Ashley N. Nielsen, Deanna J. Greene, Evan M. Gordon,
711 Adrian W. Gilmore, Steven M. Nelson, Rebecca S. Coalson, Abraham Z. Snyder, Bradley L.
712 Schlaggar, Nico U. F. Dosenbach, and Steven E. Petersen. Functional brain networks are dom-
713 inated by stable group and individual factors, not cognitive or daily variation. *Neuron*, 98(2):
714 439–452.e5, April 2018. ISSN 0896-6273. doi: 10.1016/j.neuron.2018.03.035. URL <https://www.sciencedirect.com/science/article/pii/S0896627318302411>.
715
- 716 David S. Grayson and Damien A. Fair. Development of large-scale functional networks from
717 birth to adulthood: A guide to the neuroimaging literature. *NeuroImage*, 160:15–31, Octo-
718 ber 2017. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2017.01.079. URL [https://www.](https://www.sciencedirect.com/science/article/pii/S1053811917301027)
719 [sciencedirect.com/science/article/pii/S1053811917301027](https://www.sciencedirect.com/science/article/pii/S1053811917301027).
720
- 721 Liang Han, Neil K Savalia, Micaela Y Chan, Phillip F Agres, Anupama S Nair, and Gagan S Wig.
722 Functional parcellation of the cerebral cortex across the human adult lifespan. *Cerebral Cortex*, 28
723 (12):4403–4423, December 2018. ISSN 1047-3211. doi: 10.1093/cercor/bhy218. URL <https://doi.org/10.1093/cercor/bhy218>.
724 <https://doi.org/10.1093/cercor/bhy218>.
725
- 726 Tong He, Ru Kong, Avram J. Holmes, Mert R. Sabuncu, Simon B. Eickhoff, Danilo Bzdok, Ji-
727ashi Feng, and B. T. Thomas Yeo. Is deep learning better than kernel regression for functional
728 connectivity prediction of fluid intelligence? In *2018 International Workshop on Pattern Recog-
729 nition in Neuroimaging (PRNI)*, pp. 1–4, June 2018. doi: 10.1109/PRNI.2018.8423958. URL
730 <https://ieeexplore.ieee.org/document/8423958/?arnumber=8423958>.
- 731 Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick,
732 Shakir Mohamed, and Alexander Lerchner. β -vae: Learning basic visual concepts with a con-
733 strained variational framework. In *ICLR*, 2017.
734
- 735 K. Hocke and N. Kämpfer. Gap filling and noise reduction of unevenly sampled data by means
736 of the lomb-scargle periodogram. *Atmospheric Chemistry and Physics*, 9(12):4197–4206, June
737 2009. ISSN 1680-7316. doi: 10.5194/acp-9-4197-2009. URL [https://acp.copernicus.](https://acp.copernicus.org/articles/9/4197/2009/)
738 [org/articles/9/4197/2009/](https://acp.copernicus.org/articles/9/4197/2009/).
- 739 Corey Horien, Stephanie Noble, Abigali S. Greene, Kangjoo Lee, Daniel S. Barron, Siyuan Gao,
740 David O’Connor, Mehraveh Salehi, Javid Dadashkarimi, Xilin Shen, Evelyn M. R. Lake, R. Todd
741 Constable, and Dustin Scheinost. A hitchhiker’s guide to working with large, open-source
742 neuroimaging datasets. *Nature human behaviour*, 5(2):185–193, February 2021. ISSN 2397-
743 3374. doi: 10.1038/s41562-020-01005-4. URL [https://www.ncbi.nlm.nih.gov/](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7992920/)
744 [pmc/articles/PMC7992920/](https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7992920/).
745
- 746 Brittany R. Howell, Martin A. Styner, Wei Gao, Pew-Thian Yap, Li Wang, Kristine Baluyot, Essa
747 Yacoub, Geng Chen, Taylor Potts, Andrew Salzwedel, Gang Li, John H. Gilmore, Joseph Piven,
748 J. Keith Smith, Dinggang Shen, Kamil Ugurbil, Hongtu Zhu, Weili Lin, and Jed T. Ellison. The
749 unc/umn baby connectome project (bcp): An overview of the study design and protocol develop-
750 ment. *NeuroImage*, 185:891–905, January 2019. ISSN 1053-8119. doi: 10.1016/j.neuroimage.
751 2018.03.049. URL [https://www.sciencedirect.com/science/article/pii/](https://www.sciencedirect.com/science/article/pii/S1053811918302593)
752 [S1053811918302593](https://www.sciencedirect.com/science/article/pii/S1053811918302593).
- 753 Jung-Hoon Kim, Yizhen Zhang, Kuan Han, Zheyu Wen, Minkyu Choi, and Zhongming Liu. Rep-
754 resentation learning of resting state fmri with variational autoencoder. *NeuroImage*, 241:118423,
755 November 2021. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2021.118423. URL <https://www.sciencedirect.com/science/article/pii/S1053811921006984>.

- 756 Jung-Hoon Kim, Josepheen De Asis-Cruz, Dhineshvikram Krishnamurthy, and Catherine Limperopoulou. Toward a more informative representation of the fetal–neonatal brain connectome using variational autoencoder. *eLife*, 12:e80878, May 2023. ISSN 2050-084X. doi: 10.7554/eLife.80878. URL <https://doi.org/10.7554/eLife.80878>.
- 760 Jung-Hoon Kim, Josepheen De Asis-Cruz, and Catherine Limperopoulou. Separating group- and individual-level brain signatures in the newborn functional connectome: A deep learning approach. *NeuroImage*, 299:120806, October 2024. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2024.120806. URL <https://www.sciencedirect.com/science/article/pii/S1053811924003033>.
- 766 Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- 768 Diederik P. Kingma and Max Welling. Auto-encoding variational bayes, December 2022. URL <http://arxiv.org/abs/1312.6114>.
- 771 Jack L. Lancaster, Thomas G. Glass, Bhujanga R. Lankipalli, Hunter Downs, Helen Mayberg, and Peter T. Fox. A modality-independent approach to spatial normalization of tomographic images of the human brain. *Human Brain Mapping*, 3(3):209–223, 1995. ISSN 1097-0193. doi: 10.1002/hbm.460030305. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/hbm.460030305>.
- 776 Georg Langs, Polina Golland, Yanmei Tie, Laura Rigolo, and Alexandra J. Golby. Functional geometry alignment and localization of brain areas. *Advances in neural information processing systems*, 1:1225–1233, 2010. ISSN 1049-5258. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4010233/>.
- 780 Georg Langs, Polina Golland, and Satrajit S. Ghosh. Predicting activation across individuals with resting-state functional connectivity based multi-atlas label fusion. In Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro Frangi (eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pp. 313–320, Cham, 2015. Springer International Publishing. ISBN 978-3-319-24571-3. doi: 10.1007/978-3-319-24571-3_38.
- 785 Georg Langs, Danhong Wang, Polina Golland, Sophia Mueller, Ruiqi Pan, Mert R. Sabuncu, Wei Sun, Kuncheng Li, and Hesheng Liu. Identifying shared brain networks in individuals by decoupling functional and anatomical variability. *Cerebral Cortex*, 26(10):4004–4014, October 2016. ISSN 1047-3211. doi: 10.1093/cercor/bhv189. URL <https://doi.org/10.1093/cercor/bhv189>.
- 791 Timothy O. Laumann, Evan M. Gordon, Babatunde Adeyemo, Abraham Z. Snyder, Sung Jun Joo, Mei-Yen Chen, Adrian W. Gilmore, Kathleen B. McDermott, Steven M. Nelson, Nico U.F. Dosenbach, Bradley L. Schlaggar, Jeanette A. Mumford, Russell A. Poldrack, and Steven E. Petersen. Functional system and areal organization of a highly sampled individual human brain. *Neuron*, 87(3):657–670, August 2015. ISSN 08966273. doi: 10.1016/j.neuron.2015.06.037. URL <https://linkinghub.elsevier.com/retrieve/pii/S0896627315006005>.
- 797 Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, Ian Goodfellow, and Brendan Frey. Adversarial autoencoders. ICLR, May 2016. doi: 10.48550/arXiv.1511.05644. URL <http://arxiv.org/abs/1511.05644>.
- 800 Daniel S. Margulies, Satrajit S. Ghosh, Alexandros Goulas, Marcel Falkiewicz, Julia M. Huentenburg, Georg Langs, Gleb Bezgin, Simon B. Eickhoff, F. Xavier Castellanos, Michael Petrides, Elizabeth Jefferies, and Jonathan Smallwood. Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proceedings of the National Academy of Sciences*, 113(44):12574–12579, November 2016. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1608282113. URL <https://pnas.org/doi/full/10.1073/pnas.1608282113>.
- 806 F. M. Miezin, L. Maccotta, J. M. Ollinger, S. E. Petersen, and R. L. Buckner. Characterizing the hemodynamic response: Effects of presentation rate, sampling procedure, and the possibility of ordering brain activity based on relative timing. *NeuroImage*, 11(6):735–759, June 2000. ISSN 1053-8119. doi: 10.1006/nimg.2000.0568. URL <https://www.sciencedirect.com/science/article/pii/S1053811900905688>.

- 810 Anish Mitra, Abraham Z. Snyder, Enzo Tagliazucchi, Helmut Laufs, Jed Elison, Robert W. Emerson,
811 Mark D. Shen, Jason J. Wolff, Kelly N. Botteron, Stephen Dager, Annette M. Estes, Alan
812 Evans, Guido Gerig, Heather C. Hazlett, Sarah J. Paterson, Robert T. Schultz, Martin A. Styner,
813 Lonnie Zwaigenbaum, The IBIS Network, Bradley L. Schlaggar, Joseph Piven, John R. Pruett Jr,
814 and Marcus Raichle. Resting-state fmri in sleeping infants more closely resembles adult sleep
815 than adult wakefulness. *PLOS ONE*, 12(11):e0188122, November 2017. ISSN 1932-6203.
816 doi: 10.1371/journal.pone.0188122. URL [https://journals.plos.org/plosone/
817 article?id=10.1371/journal.pone.0188122](https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0188122).
- 818 Lucille A. Moore, Robert J. M. Hermosillo, Eric Feczko, Julia Moser, Sanju Koirala, Madeleine C.
819 Allen, Claudia Buss, Greg Conan, Anthony C. Juliano, Mollie Marr, Oscar Miranda-Dominguez,
820 Michael Mooney, Michael Myers, Jerod Rasmussen, Cynthia E. Rogers, Christopher D. Smyser,
821 Kathy Snider, Chad Sylvester, Elina Thomas, Damien A. Fair, and Alice M. Graham. Towards
822 personalized precision functional mapping in infancy. *Imaging Neuroscience*, 2:1–20, May 2024.
823 ISSN 2837-6056. doi: 10.1162/imag_a.00165. URL [https://doi.org/10.1162/imag_
824 a_00165](https://doi.org/10.1162/imag_a_00165).
- 825 Sarah Feldt Muldoon and Danielle S. Bassett. Network and multilayer network approaches
826 to understanding human brain dynamics. *Philosophy of Science*, 83(5):710–720, De-
827 cember 2016. ISSN 0031-8248, 1539-767X. doi: 10.1086/687857. URL [https:
828 //www.cambridge.org/core/journals/philosophy-of-science/article/
829 network-and-multilayer-network-approaches-to-understanding-human-brain-dynamics/
830 9F495DE23816D53D4C99BF8840B610E4](https://www.cambridge.org/core/journals/philosophy-of-science/article/network-and-multilayer-network-approaches-to-understanding-human-brain-dynamics/9F495DE23816D53D4C99BF8840B610E4).
- 831 Michael J Myers, Alyssa K Labonte, Evan M Gordon, Timothy O Laumann, Jiaxin C Tu, Muriah D
832 Wheelock, Ashley N Nielsen, Rebecca F Schwarzlose, M Catalina Camacho, Dimitrios Alex-
833 opoulos, Barbara B Warner, Nandini Raghuraman, Joan L Luby, Deanna M Barch, Damien A
834 Fair, Steven E Petersen, Cynthia E Rogers, Christopher D Smyser, and Chad M Sylvester. Func-
835 tional parcellation of the neonatal cortical surface. *Cerebral Cortex*, 34(2):bhae047, February
836 2024. ISSN 1460-2199. doi: 10.1093/cercor/bhae047. URL [https://doi.org/10.1093/
837 cercor/bhae047](https://doi.org/10.1093/cercor/bhae047).
- 838 Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In
839 *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pp. 807–814,
840 2010.
- 841 James C. Pang, Kevin M. Aquino, Marianne Oldehinkel, Peter A. Robinson, Ben D. Fulcher,
842 Michael Breakspear, and Alex Fornito. Geometric constraints on human brain function. *Nature*,
843 618(7965):566–574, June 2023. ISSN 1476-4687. doi: 10.1038/s41586-023-06098-1. URL
844 <https://www.nature.com/articles/s41586-023-06098-1>.
- 845 Jonathan D. Power, Kelly A. Barnes, Abraham Z. Snyder, Bradley L. Schlaggar, and Steven E.
846 Petersen. Spurious but systematic correlations in functional connectivity mri networks arise from
847 subject motion. *NeuroImage*, 59(3):2142–2154, February 2012. ISSN 1053-8119. doi: 10.
848 1016/j.neuroimage.2011.10.018. URL [https://www.sciencedirect.com/science/
849 article/pii/S1053811911011815](https://www.sciencedirect.com/science/article/pii/S1053811911011815).
- 850 Jonathan D. Power, Anish Mitra, Timothy O. Laumann, Abraham Z. Snyder, Bradley L. Schlaggar,
851 and Steven E. Petersen. Methods to detect, characterize, and remove motion artifact in resting state
852 fmri. *NeuroImage*, 84:320–341, January 2014. ISSN 1053-8119. doi: 10.1016/j.neuroimage.
853 2013.08.048. URL [https://www.sciencedirect.com/science/article/pii/
854 S1053811913009117](https://www.sciencedirect.com/science/article/pii/S1053811913009117).
- 855 Maria Grazia Puxeddu, Joshua Faskowitz, Richard F. Betzel, Manuela Petti, Laura Astolfi,
856 and Olaf Sporns. The modular organization of brain cortical connectivity across the hu-
857 man lifespan. *NeuroImage*, 218:116974, September 2020. ISSN 10538119. doi: 10.1016/
858 j.neuroimage.2020.116974. URL [https://linkinghub.elsevier.com/retrieve/
859 pii/S1053811920304602](https://linkinghub.elsevier.com/retrieve/pii/S1053811920304602).
- 860 Ning Qiang, Qinglin Dong, Fangfei Ge, Hongtao Liang, Bao Ge, Shu Zhang, Yifei Sun, Jie Gao,
861 and Tianming Liu. Deep variational autoencoder for mapping functional brain networks. *IEEE
862 Transactions on Cognitive and Developmental Systems*, 13(4):841–852, December 2021. ISSN
863

- 864 2379-8939. doi: 10.1109/TCDS.2020.3025137. URL [https://ieeexplore.ieee.org/
865 document/9204760/?arnumber=9204760](https://ieeexplore.ieee.org/document/9204760/?arnumber=9204760).
- 866
- 867 Marcus E. Raichle. The brain’s default mode network. *Annual Review of Neuroscience*, 38(1):
868 433–447, 2015. doi: 10.1146/annurev-neuro-071013-014030. URL [https://doi.org/10.
869 1146/annurev-neuro-071013-014030](https://doi.org/10.1146/annurev-neuro-071013-014030).
- 870 David A. Ross, Jongwoo Lim, Ruei-Sung Lin, and Ming-Hsuan Yang. Incremental learning
871 for robust visual tracking. *International Journal of Computer Vision*, 77(1-3):125–141, May
872 2008. ISSN 0920-5691, 1573-1405. doi: 10.1007/s11263-007-0075-7. URL [http://link.
873 springer.com/10.1007/s11263-007-0075-7](http://link.springer.com/10.1007/s11263-007-0075-7).
- 874
- 875 Peter J. Rousseeuw. Silhouettes: A graphical aid to the interpretation and validation of cluster
876 analysis. *Journal of Computational and Applied Mathematics*, 20:53–65, November 1987. ISSN
877 0377-0427. doi: 10.1016/0377-0427(87)90125-7. URL [https://www.sciencedirect.
878 com/science/article/pii/0377042787901257](https://www.sciencedirect.com/science/article/pii/0377042787901257).
- 879 Mikail Rubinov and Olaf Sporns. Complex network measures of brain connectivity: Uses and inter-
880 pretations. *NeuroImage*, 52(3):1059–1069, September 2010. ISSN 10538119. doi: 10.1016/
881 j.neuroimage.2009.10.003. URL [https://linkinghub.elsevier.com/retrieve/
882 pii/S105381190901074X](https://linkinghub.elsevier.com/retrieve/pii/S105381190901074X).
- 883 D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning internal representations by error
884 propagation. In *Parallel Distributed Processing: Explorations in the Microstructure of Cognition,
885 Vol. 1: Foundations*, pp. 318–362. MIT Press, Cambridge, MA, USA, January 1986. ISBN 978-
886 0-262-68053-0.
- 887
- 888 Srikanth Ryali, Yuan Zhang, Carlo De Los Angeles, Kaustubh Supekar, and Vinod Menon. Deep
889 learning models reveal replicable, generalizable, and behaviorally relevant sex differences in hu-
890 man functional brain organization. *Proceedings of the National Academy of Sciences*, 121(9):
891 e2310012121, February 2024. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.2310012121.
892 URL <https://pnas.org/doi/10.1073/pnas.2310012121>.
- 893 Alexander Schaefer, Ru Kong, Evan M Gordon, Timothy O Laumann, Xi-Nian Zuo, Avram J
894 Holmes, Simon B Eickhoff, and B T Thomas Yeo. Local-global parcellation of the human
895 cerebral cortex from intrinsic functional connectivity mri. *Cerebral Cortex (New York, NY)*,
896 28(9):3095–3114, September 2018. ISSN 1047-3211. doi: 10.1093/cercor/bhx179. URL
897 <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6095216/>.
- 898 Benjamin A. Seitzman, Caterina Gratton, Scott Marek, Ryan V. Raut, Nico U.F. Dosenbach,
899 Bradley L. Schlaggar, Steven E. Petersen, and Deanna J. Greene. A set of functionally-defined
900 brain regions with improved representation of the subcortex and cerebellum. *NeuroImage*, 206:
901 116290, February 2020. ISSN 10538119. doi: 10.1016/j.neuroimage.2019.116290. URL
902 <https://linkinghub.elsevier.com/retrieve/pii/S105381191930881X>.
- 903
- 904 X. Shen, F. Tokoglu, X. Papademetris, and R. T. Constable. Groupwise whole-brain parcella-
905 tion from resting-state fmri data for network node identification. *NeuroImage*, 82:403–415,
906 November 2013. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2013.05.081. URL <https://www.sciencedirect.com/science/article/pii/S1053811913005818>.
- 907
- 908 Maxwell Shinn, Amber Hu, Laurel Turner, Stephanie Noble, Katrin H. Preller, Jie Lisa Ji, Flora
909 Moujaes, Sophie Achard, Dustin Scheinost, R. Todd Constable, John H. Krystal, Franz X. Vol-
910 lenweider, Daeyeol Lee, Alan Anticevic, Edward T. Bullmore, and John D. Murray. Func-
911 tional brain networks reflect spatial and temporal autocorrelation. *Nature Neuroscience*, 26(5):
912 867–878, May 2023. ISSN 1546-1726. doi: 10.1038/s41593-023-01299-3. URL <https://www.nature.com/articles/s41593-023-01299-3>.
- 913
- 914 Joshua S. Siegel, Anish Mitra, Timothy O. Laumann, Benjamin A. Seitzman, Marcus Raichle, Mau-
915 rizio Corbetta, and Abraham Z. Snyder. Data quality influences observed links between func-
916 tional connectivity and behavior. *Cerebral Cortex*, 27(9):4492–4502, September 2017. ISSN
917 1047-3211. doi: 10.1093/cercor/bhw253. URL [https://doi.org/10.1093/cercor/
bhw253](https://doi.org/10.1093/cercor/bhw253).

- 918 Stephen M. Smith, Mark Jenkinson, Mark W. Woolrich, Christian F. Beckmann, Timothy E. J.
919 Behrens, Heidi Johansen-Berg, Peter R. Bannister, Marilena De Luca, Ivana Drobnyak, David E.
920 Flitney, Rami K. Niazy, James Saunders, John Vickers, Yongyue Zhang, Nicola De Stefano,
921 J. Michael Brady, and Paul M. Matthews. Advances in functional and structural mr image
922 analysis and implementation as fsl. *NeuroImage*, 23:S208–S219, January 2004. ISSN 1053-
923 8119. doi: 10.1016/j.neuroimage.2004.07.051. URL <https://www.sciencedirect.com/science/article/pii/S1053811904003933>.
- 925 Abraham Z. Snyder, Tomoyuki Nishino, Joshua S. Shimony, Eric J. Lenze, Julie Loebach Wetherell,
926 Michelle Voegtle, J. Philip Miller, Michael D. Yingling, Daniel Marcus, Jenny Gurney, Jer-
927 rel Rutlin, Drew Scott, Lisa Eyler, and Deanna Barch. Covariance and correlation analy-
928 sis of resting state functional magnetic resonance imaging data acquired in a clinical trial of
929 mindfulness-based stress reduction and exercise in older individuals. *Frontiers in Neuroscience*,
930 16:825547, March 2022. ISSN 1662-453X. doi: 10.3389/fnins.2022.825547. URL <https://www.frontiersin.org/articles/10.3389/fnins.2022.825547/full>.
- 932 Olaf Sporns, Dante R. Chialvo, Marcus Kaiser, and Claus C. Hilgetag. Organization, devel-
933 opment and function of complex brain networks. *Trends in Cognitive Sciences*, 8(9):418–
934 425, September 2004. ISSN 1364-6613. doi: 10.1016/j.tics.2004.07.008. URL <https://www.sciencedirect.com/science/article/pii/S1364661304001901>.
- 936 Lianglong Sun, Tengda Zhao, Xinyuan Liang, Mingrui Xia, Qionglin Li, Xuhong Liao, Gaolang
937 Gong, Qian Wang, Chenxuan Pang, Qian Yu, Yanchao Bi, Pindong Chen, Rui Chen, Yuan Chen,
938 Taolin Chen, Jingliang Cheng, Yuqi Cheng, Zaixu Cui, Zhengjia Dai, Yao Deng, Yuyin Ding,
939 Qi Dong, Dingna Duan, Jia-Hong Gao, Qiyong Gong, Ying Han, Zaizhu Han, Chu-Chung Huang,
940 Ruiwang Huang, Ran Huo, Lingjiang Li, Ching-Po Lin, Qixiang Lin, Bangshan Liu, Chao Liu,
941 Ningyu Liu, Ying Liu, Yong Liu, Jing Lu, Leilei Ma, Weiwei Men, Shaozheng Qin, Jiang Qiu,
942 Shijun Qiu, Tianmei Si, Shuping Tan, Yanqing Tang, Sha Tao, Dawei Wang, Fei Wang, Jiali
943 Wang, Pan Wang, Xiaoqin Wang, Yanpei Wang, Dongtao Wei, Yankun Wu, Peng Xie, Xiufeng
944 Xu, Yuehua Xu, Zhilei Xu, Liyuan Yang, Huishu Yuan, Zilong Zeng, Haibo Zhang, Xi Zhang,
945 Gai Zhao, Yanting Zheng, Suyu Zhong, Cam-CAN Alzheimer’s Disease Neuroimaging Initiative,
946 and Yong He. Functional connectome through the human life span, September 2023. URL
947 <https://www.biorxiv.org/content/10.1101/2023.09.12.557193v2>.
- 948 Valerie J. Sydnor, Bart Larsen, Danielle S. Bassett, Aaron Alexander-Bloch, Damien A. Fair, Conor
949 Liston, Allyson P. Mackey, Michael P. Milham, Adam Pines, David R. Roalf, Jakob Seidlitz,
950 Ting Xu, Armin Raznahan, and Theodore D. Satterthwaite. Neurodevelopment of the association
951 cortices: Patterns, mechanisms, and implications for psychopathology. *Neuron*, 109(18):2820–
952 2846, September 2021. ISSN 0896-6273. doi: 10.1016/j.neuron.2021.06.016. URL <https://www.sciencedirect.com/science/article/pii/S0896627321004578>.
- 953 Enzo Tagliazucchi, Frederic von Wegner, Astrid Morzelewski, Verena Brodbeck, Kolja Jahnke, and
954 Helmut Laufs. Breakdown of long-range temporal dependence in default mode and attention
955 networks during deep sleep. *Proceedings of the National Academy of Sciences*, 110(38):15419–
956 15424, September 2013. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1312848110. URL
957 <https://pnas.org/doi/full/10.1073/pnas.1312848110>.
- 958 PT Jean Talairach. Co-planar stereotaxic atlas of the human brain. (*No Title*), 1988.
- 960 Jiaxin Cindy Tu, Michael Myers, Wei Li, Jiaqi Li, Xintian Wang, Donna Dierker, Trevor Day,
961 Abraham Z Snyder, Aidan Latham, Jeanette K Kenley, et al. Early life neuroimaging: The gen-
962 eralizability of cortical area parcellations across development. *bioRxiv : the preprint server for*
963 *biology*, pp. 2024–09, 2024. URL <https://doi.org/10.1101/2024.09.09.612056>.
- 964 Sebastian Urchs, Jonathan Armoza, Clara Moreau, Yassine Benhajali, Jolène St-Aubin, Pierre Or-
965 ban, and Pierre Bellec. Mist: A multi-resolution parcellation of functional brain networks. *MNI*
966 *Open Research*, 1:3, 2019. ISSN 2515-5059.
- 967 D. C. Van Essen, M. F. Glasser, D. L. Dierker, J. Harwell, and T. Coalson. Parcellations and
968 hemispheric asymmetries of human cerebral cortex analyzed on surface-based atlases. *Cerebral*
969 *Cortex*, 22(10):2241–2262, October 2012a. ISSN 1047-3211, 1460-2199. doi: 10.1093/cercor/

- 972 bhr291. URL <https://academic.oup.com/cercor/article-lookup/doi/10.1093/cercor/bhr291>.
973
974
- 975 D. C. Van Essen, K. Ugurbil, E. Auerbach, D. Barch, T. E. J. Behrens, R. Bucholz, A. Chang,
976 L. Chen, M. Corbetta, S. W. Curtiss, S. Della Penna, D. Feinberg, M. F. Glasser, N. Harel,
977 A. C. Heath, L. Larson-Prior, D. Marcus, G. Michalareas, S. Moeller, R. Oostenveld, S. E.
978 Petersen, F. Prior, B. L. Schlaggar, S. M. Smith, A. Z. Snyder, J. Xu, and E. Yacoub. The
979 human connectome project: A data acquisition perspective. *NeuroImage*, 62(4):2222–2231,
980 October 2012b. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2012.02.018. URL <https://www.sciencedirect.com/science/article/pii/S1053811912001954>.
981
- 982 David C. Van Essen, Heather A. Drury, James Dickson, John Harwell, Donna Hanlon, and
983 Charles H. Anderson. An integrated software suite for surface-based analyses of cerebral cortex. *Journal of the American Medical Informatics Association*, 8(5):443–459, September 2001. ISSN 1067-5027. doi: 10.1136/jamia.2001.0080443. URL <https://doi.org/10.1136/jamia.2001.0080443>.
984
985
986
- 987 Reinder Vos de Wael, Oualid Benkarim, Casey Paquola, Sara Lariviere, Jessica Royer, Shahin
988 Tavakol, Ting Xu, Seok-Jun Hong, Georg Langs, Sofie Valk, Bratislav Mistic, Michael Milham,
989 Daniel Margulies, Jonathan Smallwood, and Boris C. Bernhardt. Brainspace: A toolbox for the
990 analysis of macroscale gradients in neuroimaging and connectomics datasets. *Communications Biology*, 3(1):1–10, March 2020. ISSN 2399-3642. doi: 10.1038/s42003-020-0794-7. URL <https://www.nature.com/articles/s42003-020-0794-7>.
991
992
993
- 994 Gagan S. Wig. Segregated systems of human brain networks. *Trends in Cognitive Sciences*, 21(12):
995 981–996, December 2017. ISSN 1364-6613. doi: 10.1016/j.tics.2017.09.006. URL <https://www.sciencedirect.com/science/article/pii/S1364661317301948>.
996
997
- 998 Yunman Xia, Mingrui Xia, Jin Liu, Xuhong Liao, Tianyuan Lei, Xinyu Liang, Tengda Zhao, Ziyi
999 Shi, Lianglong Sun, Xiaodan Chen, Weiwei Men, Yanpei Wang, Zhiying Pan, Jie Luo, Siya
1000 Peng, Menglu Chen, Lei Hao, Shuping Tan, Jia-Hong Gao, Shaozheng Qin, Gaolang Gong,
1001 Sha Tao, Qi Dong, and Yong He. Development of functional connectome gradients during
1002 childhood and adolescence. *Science Bulletin*, 67(10):1049–1061, May 2022. ISSN 2095-9273.
1003 doi: 10.1016/j.scib.2022.01.002. URL <https://www.sciencedirect.com/science/article/pii/S2095927322000020>.
1004
- 1005 B T Thomas Yeo, Fenna M. Krienen, Jorge Sepulcre, Mert R. Sabuncu, Danial Lashkari, Marisa
1006 Hollinshead, Joshua L. Roffman, Jordan W. Smoller, Lilla Zöllei, Jonathan R. Polimeni, Bruce
1007 Fischl, Hesheng Liu, and Randy L. Buckner. The organization of the human cerebral cortex
1008 estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, 106(3):1125–1165,
1009 September 2011. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.00338.2011. URL <https://www.physiology.org/doi/10.1152/jn.00338.2011>.
1010
- 1011 Alibek Zhakubayev and Greg Hamerly. Clustering faster and better with projected data. In *Proceedings of the 6th International Conference on Information System and Data Mining*, pp. 1–6, 2022.
1012
1013
1014
- 1015 Lu Zhang, Li Wang, and Dajiang Zhu. Recovering brain structural connectivity from functional connectivity via multi-gcn based generative adversarial network. In Anne L. Martel, Purang Abolmaesumi, Danail Stoyanov, Diana Mateus, Maria A. Zuluaga, S. Kevin Zhou, Daniel Racoceanu, and Leo Joskowicz (eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, pp. 53–61, Cham, 2020. Springer International Publishing. ISBN 978-3-030-59728-3. doi: 10.1007/978-3-030-59728-3_6.
1016
1017
1018
1019
1020
- 1021 Jianlong Zhao, Jinjie Huang, Dongmei Zhi, Weizheng Yan, Xiaohong Ma, Xiao Yang, Xianbin
1022 Li, Qing Ke, Tianzi Jiang, Vince D. Calhoun, and Jing Sui. Functional network connectivity (fnc)-based generative adversarial network (gan) and its applications in classification of mental disorders. *Journal of Neuroscience Methods*, 341:108756, July 2020. ISSN 0165-0270. doi: 10.1016/j.jneumeth.2020.108756. URL <https://www.sciencedirect.com/science/article/pii/S0165027020301795>.
1023
1024
1025

1026 Qiankun Zuo, Junhua Hu, Yudong Zhang, Junren Pan, Changhong Jing, Xuhang Chen, Xiaobo
 1027 Meng, and Jin Hong. Brain functional network generation using distribution-regularized adver-
 1028 sarial graph autoencoder with transformer for dementia diagnosis. *Computer modeling in engi-
 1029 neering & sciences : CMES*, 137(3):2129, August 2023. doi: 10.32604/cmcs.2023.028732. URL
 1030 <https://pmc.ncbi.nlm.nih.gov/articles/PMC7615791/>.
 1031

1032 A APPENDIX

1033 A.1 COMMUNITY DETECTION IN A SHARED LATENT SPACE HELPS ESTABLISH 1034 CORRESPONDENCE

1037 One key pursuit of system neuroscience is to group the nodes in functional connectomes into func-
 1038 tional networks or communities across states, individuals, and the lifespan (Betzel et al., 2014;
 1039 Grayson & Fair, 2017; Mitra et al., 2017; Puxeddu et al., 2020; Tagliazucchi et al., 2013; Wig, 2017).
 1040 A key hurdle for examining those communities is the establishment of correspondence across the
 1041 different connectomes. Prior work has attempted this by detecting communities in individual con-
 1042 nectomes and then matching the degree of overlap in their topography either through visualization
 1043 (Gordon et al., 2017b) or using a Hungarian matching algorithm to minimize the Hamming distance
 1044 (Langs et al., 2016). Alternatively, a multilayer network can be constructed by linking multiple
 1045 functional connectomes as layers (Bassett et al., 2011; Betzel et al., 2019; Puxeddu et al., 2020),
 1046 where community detection methods were then applied. Neither approach could be applied to func-
 1047 tional connectomes with different nodal definitions optimized for individuals (Gordon et al., 2017b;
 1048 Laumann et al., 2015) or populations (Han et al., 2018; Myers et al., 2024).

1049 A.2 THE NEED FOR IMPROVING COMPUTATIONAL EFFICIENCY

1051 The idea of improving computational efficiency on clustering by first performing dimensional-
 1052 ity reduction rather than using the original high-dimension data has proven effective in other do-
 1053 mains(Zhakubayev & Hamerly, 2022). While the number of nodes (100-1000) and edges (4950-
 1054 499500) might not seem big for an individual functional connectivity matrix, the total data space can
 1055 become large when considering all individuals (e.g., for UK Biobank it would be 40000+ (Horien
 1056 et al., 2021)), longitudinal sessions, different data modality, and time windows within a session (de
 1057 Domenico, 2017; Muldoon & Bassett, 2016; Betzel et al., 2019). In addition, recently people com-
 1058 bine across multiple datasets for life span studies (Sun et al., 2023) which can increase the sample
 1059 size further.

1060 The theoretical time complexity of the Louvain algorithm in single-layer modularity maximization is
 1061 $O(n(\log(n)))$ where n is the number of nodes. We can approximately calculate the time complexity
 1062 of the Louvain algorithm in multi-layer modularity maximization to be $O(N \times L(\log(N \times L)))$
 1063 where N is the number of nodes per layer and L is the number of layers. Here, we monitored the
 1064 computational time for community detection with a multilayer network using the MATLAB code
 1065 provided by (Betzel et al., 2019) with different combinations of N and L . For each repetition,
 1066 we randomly drew hyperparameters ω and γ which control the strength of interlayer coupling and
 1067 resolution scale, respectively. The experiments were conducted on a server equipped with two AMD
 1068 EPYC 7713 64-core Processors, providing 128 cores and 256 threads, with a base clock speed of
 1069 2.56 GHz. Our empirical results matched the theoretical expectations (Table 1).

N (number of nodes)	L (number of layers)	Time in seconds to run 100 repetitions
200	5	8.3
200	50	120
200	500	3265
333	5	17
333	50	380
333	500	6694

1070 Table 1: Computational time monitoring for multilayer modularity maximization

1078 On the other hand, other popular clustering algorithms like Gaussian Mixture models have a
 1079 quadratic dependence on the number of dimensions due to the complexity of manipulating covari-

1080 ance matrices. Therefore, even though there is no strong need to go from the parcel space (hundreds)
1081 to a low-dimensional latent space, it should still provide noticeable improvements.

1083 A.3 NEUROIMAGING DATA ACQUISITION AND PROCESSING DETAILS

1085 A.3.1 WASHINGTON UNIVERSITY 120 (WU120) DATA

1087 In summary, data was obtained from 120 healthy young adults (60 females, average age = 25 years,
1088 age range = 19–32 years). Participants were right-handed, native English speakers recruited from
1089 the Washington University community. Screening via self-report questionnaire ensured no history
1090 of neurological or psychiatric diagnosis, nor head injuries resulting in more than 5 minutes of uncon-
1091 sciousness. All participants provided informed consent, and the study was approved by the Wash-
1092 ington University School of Medicine Human Studies Committee and Institutional Review Board.
1093 Data is available at <https://openneuro.org/datasets/ds000243/versions/00001/file-display/00001>

1094 Structural and functional MRI data were obtained with a Siemens MAGNETOM Trio Tim 3.0-T
1095 Scanner and a Siemens 12-channel Head Matrix Coil. Structural imaging included a T1-weighted
1096 sagittal magnetization-prepared rapid acquisition gradient-echo (MP-RAGE) structural image was
1097 obtained [time echo (TE) = 3.08 ms, time repetition, TR (partition) = 2.4 s, time to inversion (TI) =
1098 1000 ms, flip angle = 8°, 176 slices with $1 \times 1 \times 1$ mm voxels]. Functional scan slices were aligned
1099 parallel to the anterior commissure–posterior commissure plane of the MP-RAGE and centered on
1100 the brain using an auto-align pulse sequence protocol available in Siemens software. This alignment
1101 corresponds to the Talairach atlas (Talairach, 1988).

1102 For the functional MRI data acquisition, subjects were instructed to relax while maintaining fixa-
1103 tion on a black crosshair against a white background. Functional imaging used a BOLD contrast-
1104 sensitive gradient-echo echo-planar imaging (EPI) sequence (TE = 27 ms, flip angle = 90°, in-plane
1105 resolution = 4×4 mm). Full-brain EPI volumes (MR frames) of 32 contiguous, 4-mm-thick axial
1106 slices were obtained every 2.5 seconds. Additionally, a T2-weighted turbo spin-echo structural im-
1107 age (TE = 84 ms, TR = 6.8 s, 32 slices with $1 \times 1 \times 4$ mm voxels) in the same anatomical planes
1108 as the BOLD images was also captured to augment atlas alignment. The fMRI acquisition used An-
1109 terior→Posterior (AP) phase encoding. The number of volumes collected per subject ranged from
1110 184 to 724, with an average of 336 frames (14.0 min).

1111 Functional images were first processed to reduce artifacts including (1) correction of odd versus even
1112 slice intensity differences due to interleaved acquisition without gaps, (2) head movement correction
1113 within and across runs, and (3) across-run intensity normalization to a whole-brain mode value of
1114 1000. Each individual’s functional data was transformed into an atlas space using the MP-RAGE
1115 scan and resampled to an isotropic 3-mm atlas space (Talairach, 1988), using a single cubic spline
1116 interpolation (Lancaster et al., 1995).

1117 Additional preprocessing mitigated high-motion frame effects in two iterations. The first iteration
1118 included: (1) demeaning and detrending, (2) multiple regression including whole-brain, ventricu-
1119 lar cerebrospinal fluid (CSF), and white matter signals, and motion regressors derived by Volterra
1120 expansion and (3) a band-pass filter ($0.009\text{Hz} < f < 0.08\text{Hz}$). Temporal masks were created
1121 in this iteration to flag motion-contaminated frames, identified using framewise displacement (FD),
1122 calculated as the squared sum of the motion vectors (Power et al., 2012). Volumes exceeding $\text{FD} > 0.2$ mm and segments of fewer than 5 contiguous volumes were flagged for removal.

1123 The data were then reprocessed in a second iteration, incorporating the temporal masks described
1124 above. This reprocessing was identical to the initial processing stream but ignored censored data.
1125 Data were interpolated across censored frames using least squares spectral estimation (Power et al.,
1126 2014) of the values at censored frames so that continuous data could be passed through the band-pass
1127 filter ($0.009\text{Hz} < f < 0.08\text{Hz}$) without contaminating frames near high motion frames. Censored
1128 frames were ultimately ignored during functional connectivity matrix generation.

1129 Individual surfaces were generated from the structural images and functional data were sampled to
1130 surface space (Glasser et al., 2013). Following volumetric registration, left and right hemisphere
1131 anatomical surfaces were created from each subject’s MP-RAGE image using FreeSurfer’s recon-
1132 all processing pipeline (v5.0) (Fischl, 2012). This involved brain extraction, segmentation, white
1133 matter and pial surface generation, surface inflation to a sphere, and spherical registration of the
subject’s “native” surface to the fsaverage surface. The fsaverage-registered surfaces were then

1134 aligned and resampled to a resolution of 164000 vertices using Caret tools (Van Essen et al., 2001)
1135 and down-sampled to a 32492 vertex surface (32k-fsLR). Functional BOLD volumes were sampled
1136 to each subject’s individual “native” midthickness surface (generated as the average of the white and
1137 pial surfaces) using the ribbon-constrained sampling procedure available in Connectome Workbench
1138 (v0.84) and then deformed and resampled from the individual’s “native” surface to the 32k-fsLR sur-
1139 face. The final time series were smoothed along the 32k-fsLR surface using a Gaussian smoothing
1140 kernel ($\sigma = 2.55$ mm).

1142 A.3.2 HUMAN CONNECTOME PROJECT (HCP) DATA

1144 Data for the Human Connectome Project (Young Adult) was collected at Washington University
1145 in St. Louis and the University of Minnesota. The participants, healthy adults aged between
1146 22 to 35 years, underwent high-resolution T1-weighted (MP-RAGE, TR = 2.4s, voxel size =
1147 $0.7 \times 0.7 \times 0.7$ mm) and BOLD contrast-sensitive imaging (gradient echo EPI, multiband factor 8, TR
1148 = 0.72s, voxels = $2 \times 2 \times 2$ mm) using a custom Siemens SKYRA 3.0T MRI scanner equipped with a
1149 custom 32-channel Head Matrix Coil. Sequences with both left-to-right (LR) and right-to-left (RL)
1150 phase encoding were employed, with each participant completing a single run in each direction over
1151 two consecutive days, resulting in four runs in total, two for Rest 1 and another two for Rest 2. Data
1152 is available at [https://www.humanconnectome.org/study/hcp-young-adult/document/1200-subjects-](https://www.humanconnectome.org/study/hcp-young-adult/document/1200-subjects-data-release)
1153 [data-release](https://www.humanconnectome.org/study/hcp-young-adult/document/1200-subjects-data-release)

1154 Previous research has indicated that minimally pre-processed data does not sufficiently control for
1155 confounds such as subject head motion (Burgess et al., 2016). In addition, reliable associations
1156 between functional connectivity and behavior require sufficient low-motion functional connectivity
1157 data for each subject Gordon et al. (2017b); Laumann et al. (2015). A similar preprocessing method
1158 as before was implemented (Seitzman et al., 2020). First, to account for magnetization equilibrium
1159 and any responses evoked by the scan start (Laumann et al., 2015), the first 29.52 seconds or – 41
1160 frames – of each resting-state run were discarded. Then, the functional data were aligned to the
1161 first frame of the first run using rigid body transforms, motion corrected (3D-cross realigned), and
1162 whole-brain mode 1000 normalized (Miezin et al., 2000). The data, with $2 \times 2 \times 2$ mm voxels, was then
1163 registered to the T1-weighted image and a WashU MNI atlas using affine and FSL transformations
1164 (Smith et al., 2004).

1165 Further preprocessing of the resting-state BOLD data was applied to remove artifacts (Ciric et al.,
1166 2017; Power et al., 2014). This involved calculating frame-wise displacement (FD) to quantify
1167 motion between consecutive frames in fMRI data (Power et al., 2012) and artifact removal using a
1168 low-pass filter at 0.1 Hz to address respiration artifacts affecting the FD estimates (Fair, 2020; Siegel
1169 et al., 2017), along with a threshold for removing frames with FD greater than 0.04 mm after the
1170 low-pass filter (N.B. for section A.6, no frame removal was applied because this leaves different
1171 numbers of valid frames for each subject). For functional connectivity (FC) analysis preparation,
1172 regression of nuisance variables was performed, including 36 regressors: (1) 3 time series (whole-
1173 brain mean, mean ventricular CSF, mean white matter) with temporal derivatives from the Volterra
1174 expansion (12 total parameters) (Friston et al., 1996), and (2) 6 head motion parameters with tempo-
1175 ral derivatives from the Volterra expansion (24 total parameters). Spatial masks of the gray matter,
1176 white matter, and ventricles were created from the T1-weighted images for each of the individual-
1177 specific regressors using Freesurfer 5.3 automatic segmentation (Fischl et al., 2002). Data segments
1178 shorter than 5 contiguous frames were excluded, and least squares spectral estimation was used for
1179 interpolation over the censored frames (Hocke & Kämpfer, 2009; Power et al., 2014). The final
1180 data were then bandpass filtered from 0.009 to 0.08 Hz, and censored frames were excluded from
1181 the time series (Seitzman et al., 2020). It is crucial to perform censoring and interpolation before
1182 filtering to prevent high-motion noise artifacts from smearing into adjacent frames.

1183 Following that, the preprocessed BOLD time series data underwent surface processing, which in-
1184 volved using the ribbon-constrained sampling procedure in Connectome Workbench to sample the
1185 BOLD volumes to each subject’s native surface and exclude voxels with a time series coefficient
1186 with a variation of 0.5 SDs above that of the mean of nearby voxels (Glasser et al., 2013). After be-
1187 ing sampled to the surface, time courses were deformed, resampled, and smoothed with a Gaussian
smoothing kernel (FWHM = 4mm, $\sigma = 1.7$). Connectome Workbench was then used to combine
these surfaces with volumetric subcortical and cerebellar data into the CIFTI format, generating
full-brain time courses while excluding non-gray matter tissue (Glasser et al., 2013).

1188 A.4 GEOMETRIC REFORMATTING INTO 2-D IMAGES

1189
1190 The geometric reformatting procedure was done in four steps. First, the seed maps were mapped to
1191 the cortical surface using their coordinates in the 32k-fsLR mesh of the left and right hemispheres
1192 (32492 vertices per hemisphere with some of them empty due to the presence of the medial wall).
1193 Then, the surfaces in each hemisphere were inflated to a sphere using FreeSurfer (Fischl, 2012).
1194 After that, we used cart2sph.m in MATLAB to convert its Cartesian coordinates (x,y,z) to spherical
1195 coordinates (a,e), which reported the azimuth and elevation angles in a range from $-\pi$ to $+\pi$ and
1196 from $-\pi/2$ to $+\pi/2$, respectively. Lastly, we defined a 192×192 grid to resample the spherical
1197 surface with respect to azimuth and $\sin(\text{elevation})$ such that the resampled locations were uniformly
1198 distributed at approximation.

1199 A.5 AUTOENCODER NETWORK ARCHITECTURE

1200
1201 In the encoder network, the size of the output image of each convolutional layer (from left to right)
1202 is $96 \times 96 \times 64$ (32 channels per hemisphere), $48 \times 48 \times 128$, $24 \times 24 \times 128$, $12 \times 12 \times 256$, and $6 \times 6 \times 256$; for
1203 the decoder network, $6 \times 6 \times 256$, $12 \times 12 \times 256$, $24 \times 24 \times 128$, $48 \times 48 \times 128$, and $96 \times 96 \times 64$ (32 channels
1204 per image), from left to right. The convolution operations are defined as 1: convolution (kernel
1205 size=8, stride=2, padding=3) with rectified nonlinearity, 2-5: convolution (kernel size=4, stride=2,
1206 padding=1) with rectified nonlinearity, 6: fully-connected layer with re-parametrization, 7: fully-
1207 connected layer with rectified nonlinearity, 8-11: transposed convolution (kernel size=4, stride=2,
1208 padding=1) with rectified nonlinearity, 12: transposed convolution (kernel size=8, stride=2,
1209 padding=3). The code implementation was adapted from <https://github.com/libilab/rsfMRI-VAE/>.

1210 The adversarial network had a separate discriminator component and the adversarial loss
1211 was combined with reconstruction loss for regularization of the latent space (Makhzani et al.,
1212 2016). The code implementation was adapted from [https://github.com/eriklindernoren/PyTorch-
1213 GAN/blob/master/implementations/aae/aae.py](https://github.com/eriklindernoren/PyTorch-GAN/blob/master/implementations/aae/aae.py).

1214 A.6 DETERMINATION OF HYPERPARAMETERS

1215 A.6.1 NUMBER OF LATENT DIMENSIONS

1216
1217 While it is desirable to have the lowest dimension possible for best computational efficiency, a
1218 dimension that is too small might not capture all crucial structures in the original data. Based
1219 on prior literature, we believed that a dimension greater than 40 would provide little additional
1220 information (Gotts et al., 2020; Margulies et al., 2016). We validated this in our data by evaluating
1221 the reconstruction loss and KL divergence in models with a range of dimensions (2, 4, 8, 16, 32, 64,
1222 96, 128, 256) at $\beta = 1$ on the validation data (Figure 8A). We found that 32 dimensions were close
1223 to the elbow of the tradeoff between reconstruction performance and KL divergence, so we only
1224 tested models with latent dimensions 2, 4, and 32, with the reasoning that a latent dimensionality
1225 beyond 32 is likely to provide little improvement in reducing the reconstruction loss but would keep
1226 increasing the KL divergence. The same metric for dimensions = 2,4,8,16,32 in $\beta = 20$ was plotted
1227 as a comparison. We applied the same latent dimensions to the PCA and ICA approaches.
1228

1229 A.6.2 VAE BETA

1230
1231 As above, we explored a range of beta values from X to Y and found that $\beta = 20$ provided a good
1232 balance between reconstruction performance and KL divergence on the validation data (Figure 8B).
1233

1234 A.7 MODEL VARIABILITY

1235
1236 One limitation is that the performance of the models could be biased by the samples selected to
1237 train the model. Here, we demonstrate the model variability with an example VAE at $\beta = 20$ with
1238 5-fold cross-validation in the training data (Figure 9). The 100 subjects were divided into 20-subject
1239 groups. For each fold, we picked 4 groups to train the model and test on the same 10 subjects
1240 outside the 100 subjects as described in section A.3.1. We can see that the models achieved very
1241 similar performance on the training data, but had a larger variability in how well it generalize to an
unseen validation set.

1242
 1243
 1244
 1245
 1246
 1247
 1248
 1249
 1250
 1251
 1252
 1253
 1254
 1255
 1256
 1257
 1258
 1259
 1260
 1261
 1262
 1263
 1264
 1265
 1266
 1267
 1268
 1269
 1270
 1271
 1272
 1273
 1274
 1275
 1276
 1277
 1278
 1279
 1280
 1281
 1282
 1283
 1284
 1285
 1286
 1287
 1288
 1289
 1290
 1291
 1292
 1293
 1294
 1295

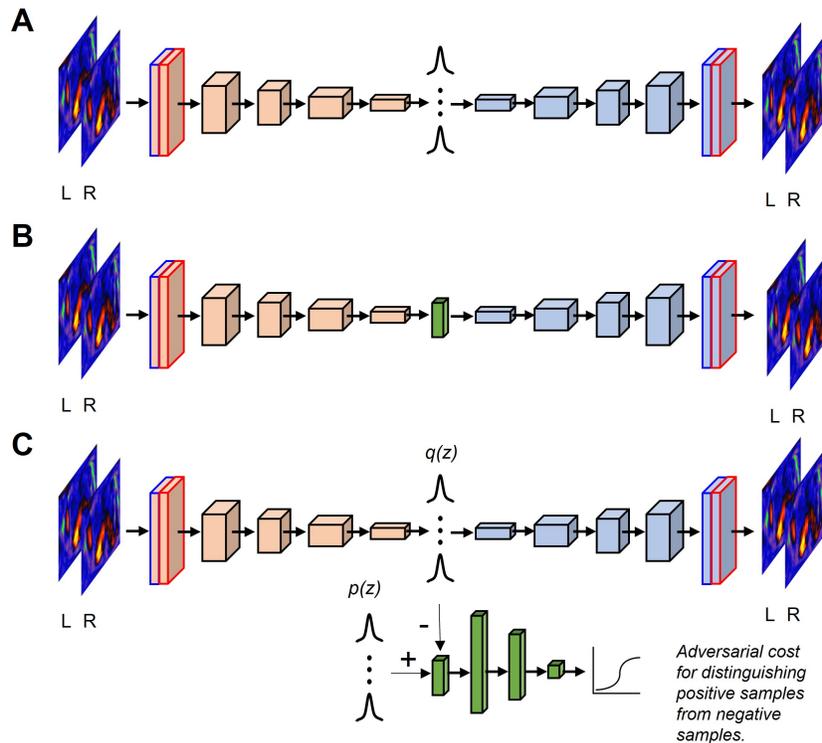


Figure 7: Autoencoder variants. A) Variational autoencoder. B) Autoencoder. C) Adversarial autoencoder. Blue and red boxes stand for the input images from left and right hemispheres, respectively. A-C: Adapted with permission from Kim, J., Zhang, Y., Han, K., Wen, Z., Choi, M., & Liu, Z. (2021). Representation learning of resting state fMRI with variational autoencoder. *NeuroImage*, 241, 118423. Copyright 2021 by Elsevier Inc. C: Adapted from Makhzani, A., Shlens, J., Jaitly, N., Goodfellow, I., Frey, B., ICLR 2016. Adversarial Autoencoders.

1296
 1297
 1298
 1299
 1300
 1301
 1302
 1303
 1304
 1305
 1306
 1307
 1308
 1309
 1310
 1311
 1312
 1313
 1314
 1315
 1316
 1317
 1318
 1319
 1320
 1321
 1322
 1323
 1324
 1325
 1326
 1327
 1328
 1329
 1330
 1331
 1332
 1333
 1334
 1335
 1336
 1337
 1338
 1339
 1340
 1341
 1342
 1343
 1344
 1345
 1346
 1347
 1348
 1349

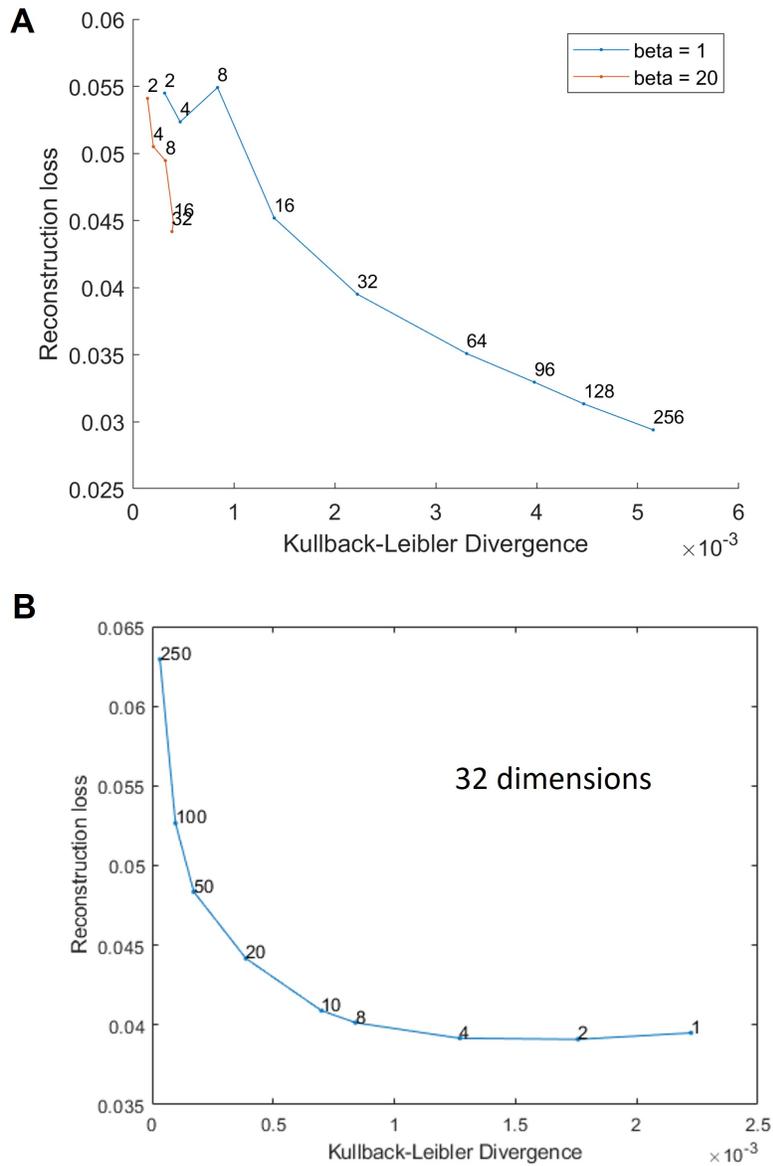


Figure 8: **Hyperparameter tuning.** A) Reconstruction loss and Kullback-Leibler divergence for different numbers of latent dimensions in the validation data. B) Reconstruction loss and Kullback-Leibler divergence for different β values in the validation data.

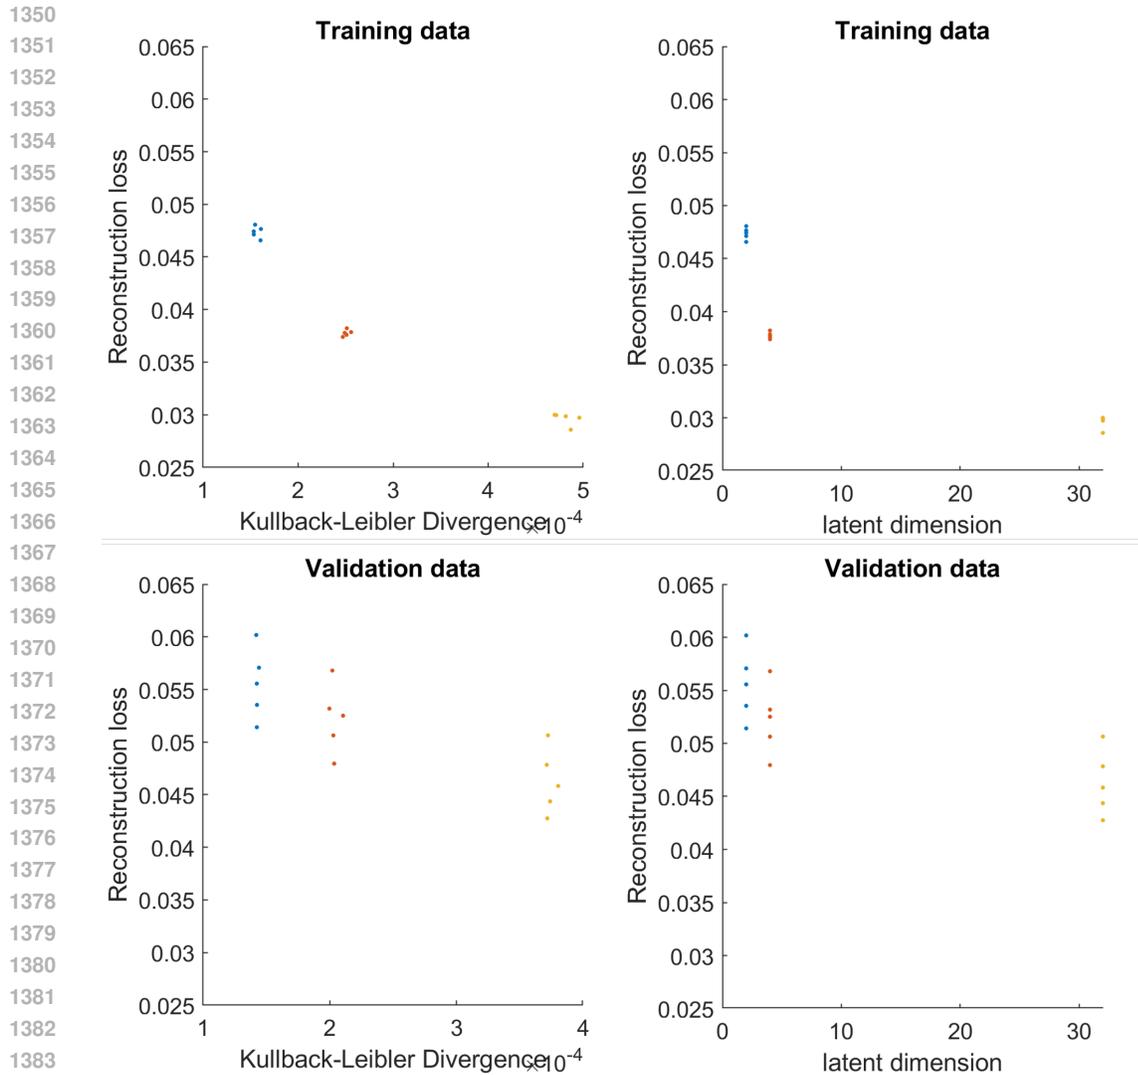


Figure 9: **5-fold cross-validation.** A) Reconstruction loss against KL-divergence for training data. B) Reconstruction loss against KL-divergence for test data. C-D) Same as A-B for validation data.

A.8 PSEUDO-CODE FOR ANALYSIS PIPELINE USING THE CONNECTOME LATENT REPRESENTATIONS

Below is the pseudo-code for generating the seed maps from each parcel and then embed them in a low-dimensional latent space.

```

1404 FOR each data session IN all_subjects:
1405   FOR each parcel i:
1406     // Compute the seed map for the parcel
1407     ParcelSeedMap = correlation(ParcelTimeSeries(:, i),
1408   VertexTimeSeries(:, i))
1409     // Encode the seed map to obtain latent representation
1410     ParcelSeedMapLatents = Encoder(ParcelSeedMap)
1411     // Concatenate the latent representations across all parcels
1412     Append ParcelSeedMapLatents TO ParcelSeedMapAllLatents
1413   END FOR
1414   // Save the concatenated latent representations for
1415   clustering/prediction
1416   SAVE ParcelSeedMapAllLatents
1417 END FOR

```

A.9 RECONSTRUCTION PERFORMANCE NORMALIZED TO THE NOISE CEILING

Since there is some variability in noise ceilings (a.k.a. the η^2 between the two resting scan sessions for the same subject), we additionally plot the normalized reconstruction performance calculated by dividing the actual reconstruction performance (η^2) by the noise ceiling for each subject in the HCP dataset.

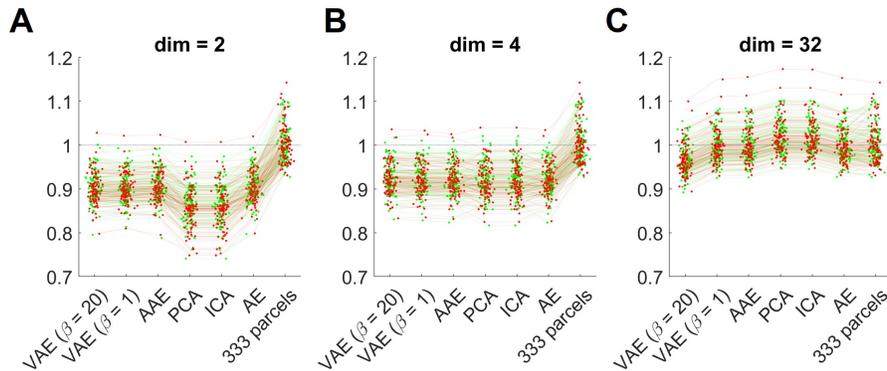


Figure 10: **Reconstruction performance in the HCP data.** A-C) The reconstruction performance for all subjects. Reconstruction Green line: HCP Rest1, Red line: HCP Rest2. Gray shaded area: mean and standard deviation of the noise ceiling and the null baseline. N.B.: 333 parcels always have 333 dimensions and were repeatedly displayed in all three panels as a reference. VAE = variational autoencoder. AAE = adversarial autoencoder. PCA = principal component analysis. ICA = independent component analysis. AE = autoencoder.

A.10 SEPARATION BY CANONICAL FUNCTIONAL NETWORKS

We found that the average SI of the 333 parcel connectome is 0.158 (95% CI: [0.145, 0.163], Figure 11). Also, the average SI was higher in dim = 4 and dim = 32 than in dim = 2 (Table 2).

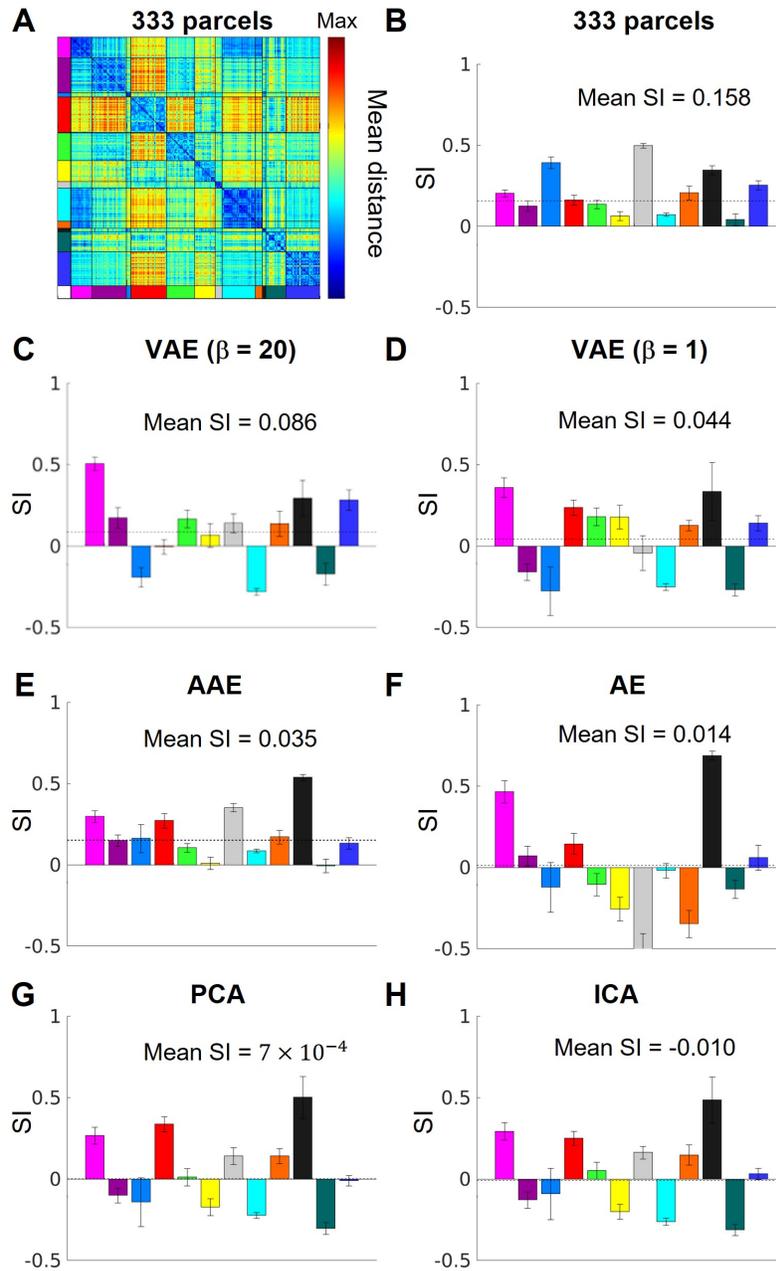


Figure 11: Separation of networks A) Mean distance between area parcels in the latent embeddings with the 333 parcels (dimension = 333). B) Mean \pm standard error SI for each functional network with the 333 parcels. C-H) Same as B but for the dimension = 2 across dimensionality reduction methods as shown in Figure 4.

Mean SI	VAE ($\beta = 20$)	VAE ($\beta = 1$)	AAE	AAE	PCA	ICA
dim = 2	0.086 [0.066, 0.099]	0.044 [0.015, 0.063]	0.035 [0.006, 0.054]	0.014 [-0.018, 0.030]	7×10^{-4} [-0.016, 0.009]	-0.010 [-0.028, 0.001]
dim = 4	0.213 [0.197, 0.221]	0.193 [0.176, 0.200]	0.181 [0.161, 0.188]	0.177 [0.158, 0.184]	0.173 [0.155, 0.181]	0.173 [0.156, 0.182]
dim = 32	0.171 [0.158, 0.175]	0.191 [0.177, 0.194]	0.152 [0.136, 0.152]	0.117 [0.094, 0.120]	0.178 [0.164, 0.183]	0.114 [0.104, 0.118]

Table 2: **Silhouette index for different dimensionality reduction methods for the Gordon network clusters.** "dim" stands for dimensions.

A.11 PROPERTIES OF DEFAULT SUB-CLUSTERS IN VAE LATENT SPACE (DIM = 2)

We applied a k-means clustering algorithm on the connectome latent representations from all default network parcels from the Rest1 session in all HCP 94 subjects in the 2-D VAE latent space (Figure 12A-B). The resultant cluster contains data from similar parcels across the subjects (Figure 12C) with some variability across subjects (Figure 10D). Cluster 1 seems to contain mostly parcels in the medial prefrontal cortex, the temporal cortex and the dorsolateral prefrontal cortex, while cluster 2 seems to contain mostly parcels from the inferior parietal cortex and the posterior cingulate cortex (Figure 12E). Subtle differences can be seen in the reconstructed seed maps from the centroids of the two clusters (Figure 12F). While the distribution of cluster 1 resembles the ventromedial and pregenual components of the default network, and the distribution of cluster 2 resembles the dorsolateral and parietal components of the default network as identified in an earlier study (Gordon et al., 2020), the biological and functional relevance of the two clusters identified here requires further investigation. However, the key takeaway of this analysis is that this eluded the possibility of identifying different subcomponents or individual/developmental-specific variants of existing functional connectivity networks using clustering of parcel connectome latent representations across subjects.

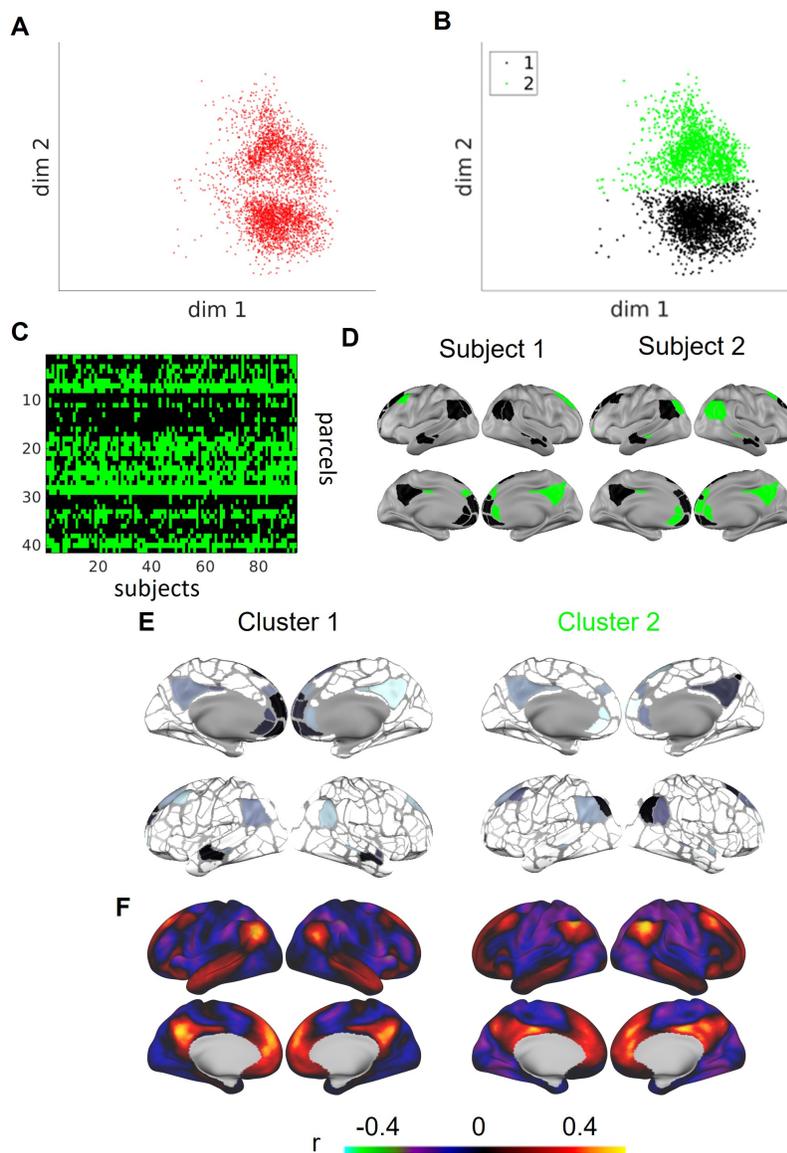


Figure 12: **Separation of DMN into subclusters.** A) DMN parcels across all subjects. B) Clustering the data in A into two sub-clusters with the k-means algorithm. C) The sub-cluster membership across parcels and subjects. D) The sub-cluster membership in two example subjects. E) The relative frequency of cluster membership across subjects (darker = more subjects) for the two sub-clusters. F) The reconstructed seed maps from the centroids of the two sub-clusters.

A.12 K-MEANS CLUSTERING OF ALL PARCEL SEED MAPS IN THE LATENT SPACE

We illustrated the initial and final results of the optimization of cluster separation of the 12 functional networks using k-means clustering in Figure 13. Here we only show the Rest1 sessions of the 94 HCP subjects for visualization purposes, but the k-means algorithm was run on data from both Rest1 and Rest2 sessions. The k-means algorithm minimizes the squared distance between all the points to their closest cluster center by changing the cluster membership of the data points. More details can be seen in Table 3 and Figure 14-21.

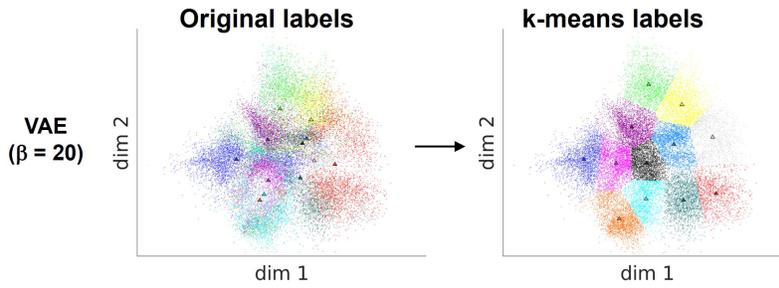
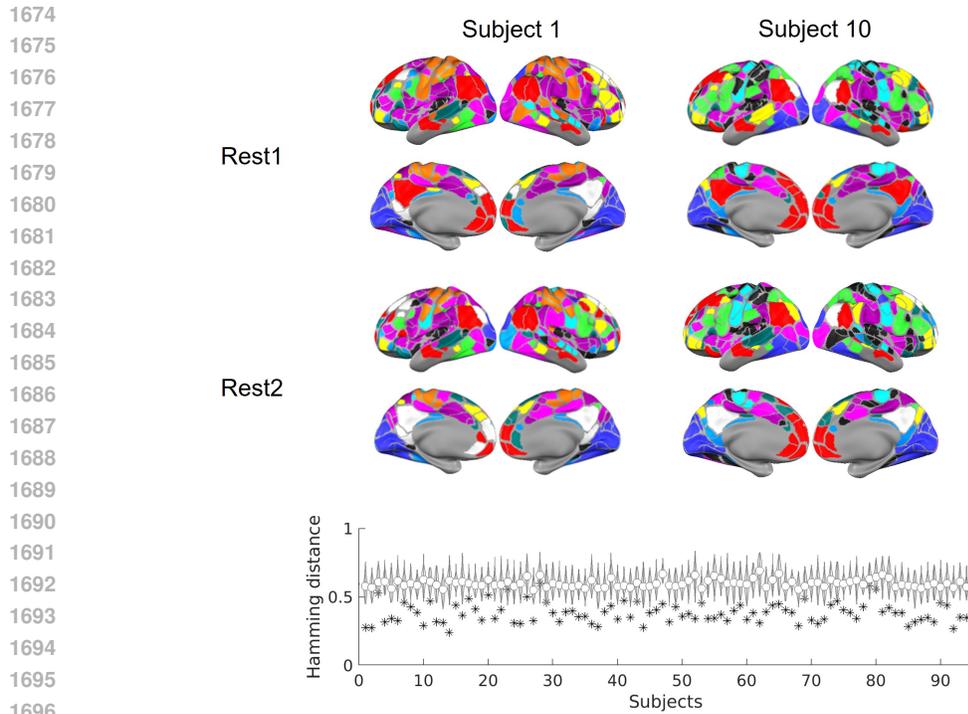


Figure 13: **k-means clustering in the latent space (2 dimensions)**. (Left) Data from Rest1 in HCP 94 subjects with the original labels (dots) and centroid (triangle). (Right) Data from Rest1 in HCP 94 subjects with the optimized label after k-means clustering with 12 clusters.

Intrasubject variability	dim = 2	dim = 4	dim = 32
VAE ($\beta = 20$)	0.39 ± 0.07	0.31 ± 0.07	0.29 ± 0.08
VAE ($\beta = 1$)	0.41 ± 0.08	0.40 ± 0.08	0.35 ± 0.08
AAE	0.40 ± 0.08	0.38 ± 0.08	0.45 ± 0.10
AE	0.39 ± 0.08	0.40 ± 0.09	0.54 ± 0.10
PCA	0.47 ± 0.11	0.38 ± 0.10	0.34 ± 0.09
ICA	0.48 ± 0.10	0.36 ± 0.09	0.27 ± 0.05
333 parcels	0.33 ± 0.08		
Intersubject variability	dim = 2	dim = 4	dim = 32
VAE ($\beta = 20$)	0.60 ± 0.03	0.51 ± 0.03	0.48 ± 0.03
VAE ($\beta = 1$)	0.62 ± 0.03	0.60 ± 0.03	0.58 ± 0.03
AAE	0.60 ± 0.03	0.56 ± 0.04	0.69 ± 0.03
AE	0.57 ± 0.04	0.60 ± 0.04	0.74 ± 0.02
PCA	0.70 ± 0.02	0.61 ± 0.03	0.57 ± 0.03
ICA	0.70 ± 0.02	0.58 ± 0.03	0.47 ± 0.05
333 parcels	0.55 ± 0.03		
vSNR	dim = 2	dim = 4	dim = 32
VAE ($\beta = 20$)	0.61 ± 0.29	0.74 ± 0.38	0.81 ± 0.50
VAE ($\beta = 1$)	0.58 ± 0.31	0.59 ± 0.34	0.73 ± 0.36
AAE	0.57 ± 0.31	0.56 ± 0.30	0.57 ± 0.33
AE	0.51 ± 0.32	0.56 ± 0.35	0.43 ± 0.28
PCA	0.57 ± 0.38	0.71 ± 0.42	0.76 ± 0.45
ICA	0.53 ± 0.31	0.71 ± 0.41	0.78 ± 0.38
333 parcels	0.74 ± 0.40		

Table 3: **Variability of parcel assignment within and between subjects using k-means clustering ($k = 12$ clusters)**.



1697 **Figure 14: Individual-specific network using k-means clustering with 12 networks (VAE ($\beta =$**
1698 **20), 2 dimensions).**

1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727

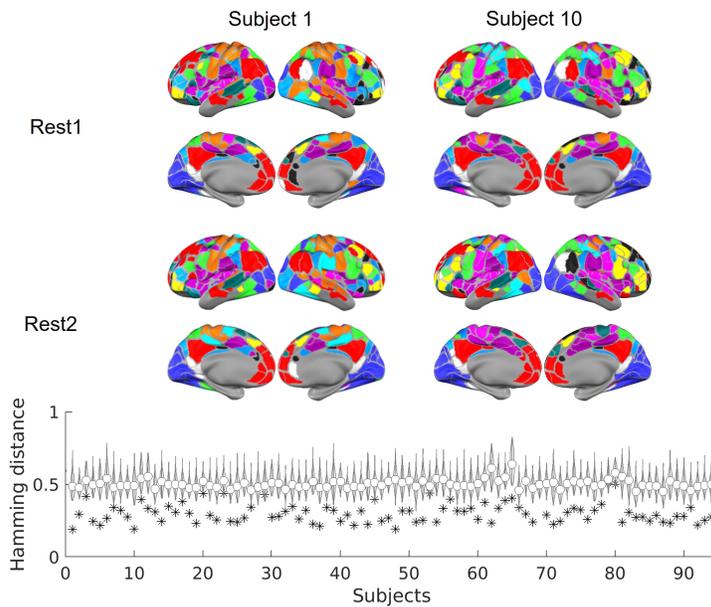
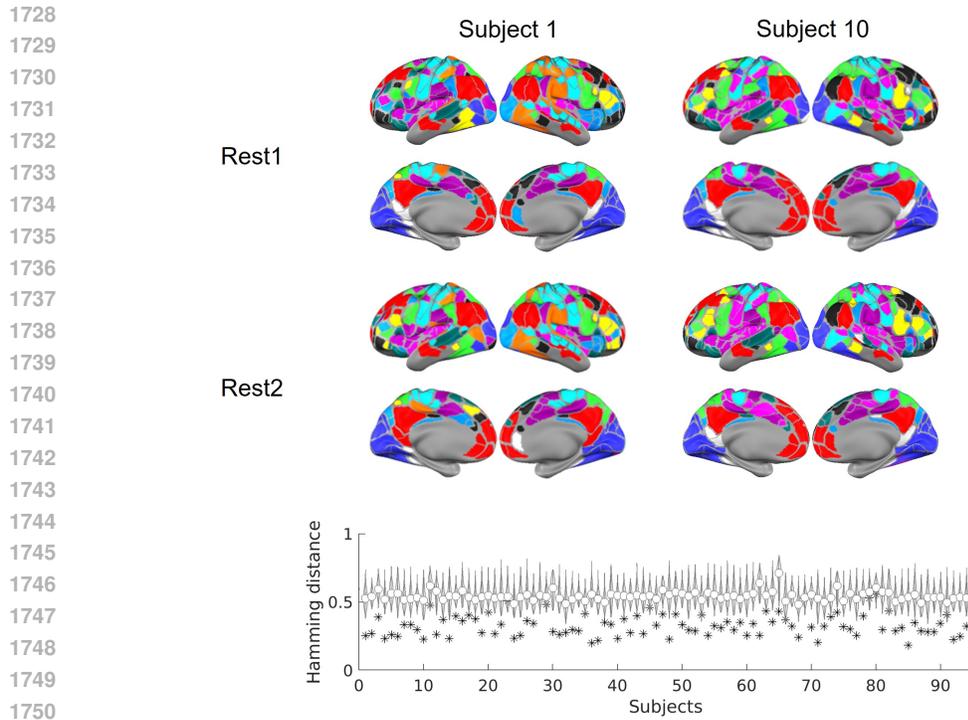


Figure 15: Individual-specific network using k-means clustering with 12 networks (VAE ($\beta =$
20), 4 dimensions).



1751 **Figure 16: Individual-specific network using k-means clustering with 12 networks (333**
1752 **parcels).**

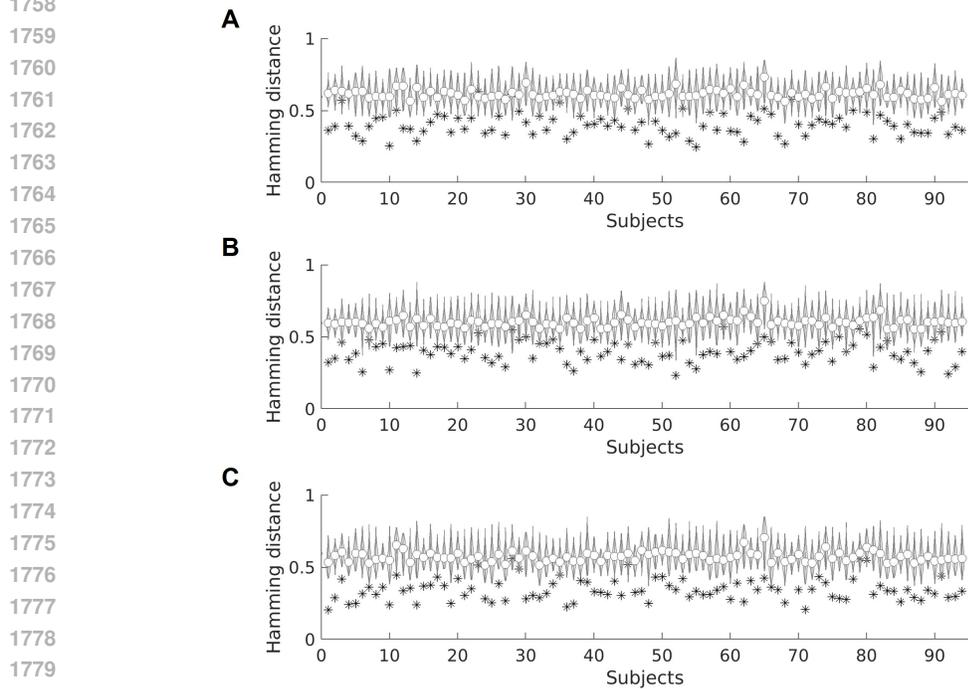


Figure 17: Individual specificity of network assignment using k-means clustering with 12 net-
works (VAE ($\beta = 1$)). A) dimension = 2, B) dimension = 4, C) dimension = 32

1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804

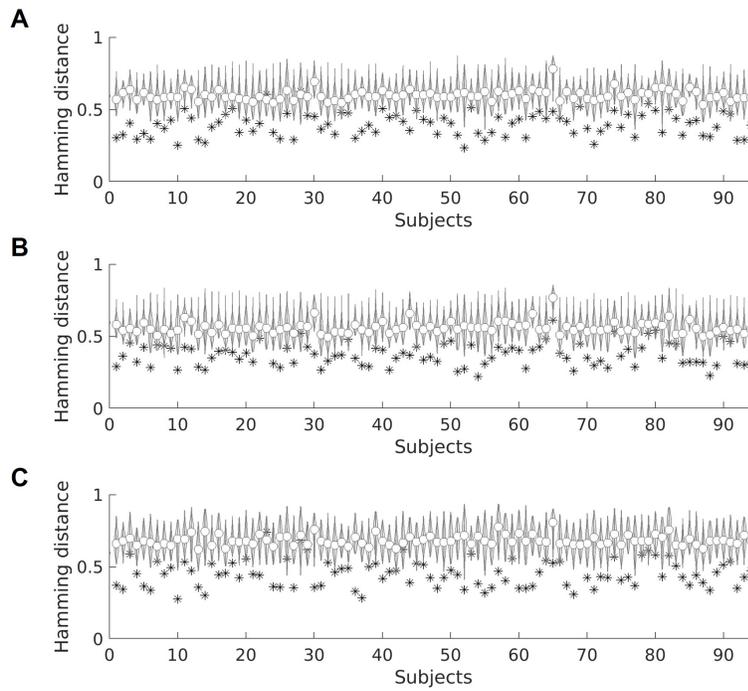


Figure 18: **Individual specificity of network assignment using k-means clustering with 12 networks (AAE).** A) dimension = 2, B) dimension = 4, C) dimension = 32

1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835

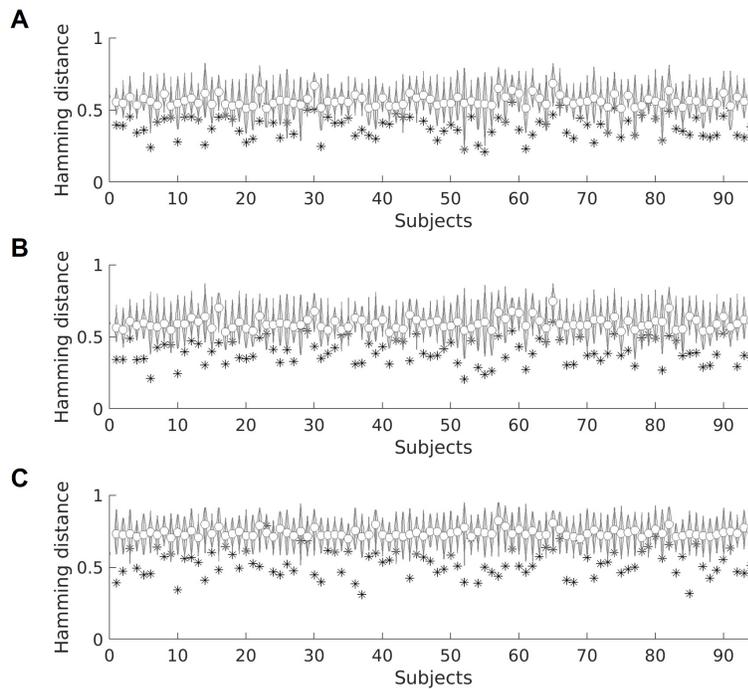


Figure 19: **Individual specificity of network assignment using k-means clustering with 12 networks (AE).** A) dimension = 2, B) dimension = 4, C) dimension = 32

1836
 1837
 1838
 1839
 1840
 1841
 1842
 1843
 1844
 1845
 1846
 1847
 1848
 1849
 1850
 1851
 1852
 1853
 1854
 1855
 1856
 1857
 1858
 1859
 1860
 1861
 1862
 1863
 1864
 1865
 1866
 1867
 1868
 1869
 1870
 1871
 1872
 1873
 1874
 1875
 1876
 1877
 1878
 1879
 1880
 1881
 1882
 1883
 1884
 1885
 1886
 1887
 1888
 1889

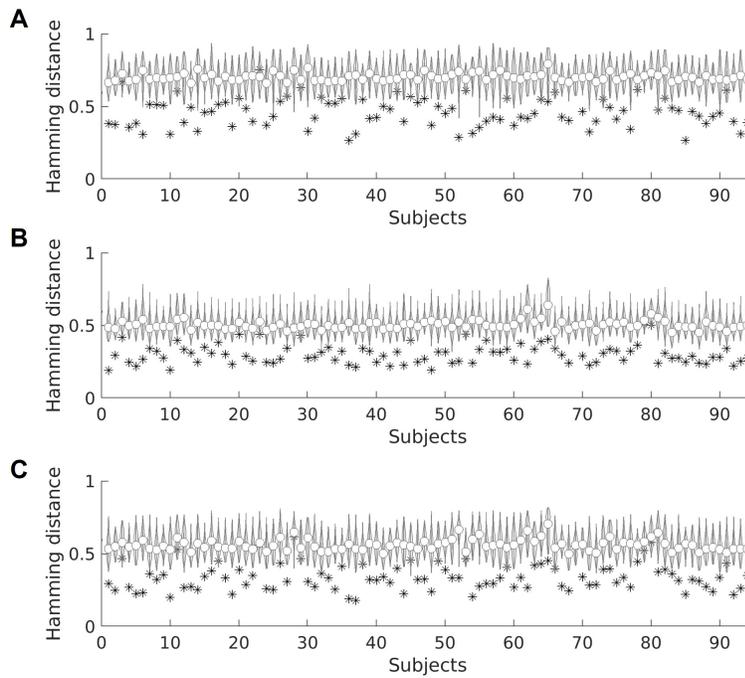


Figure 20: **Individual specificity of network assignment using k-means clustering with 12 networks (PCA)**. A) dimension = 2, B) dimension = 4, C) dimension = 32

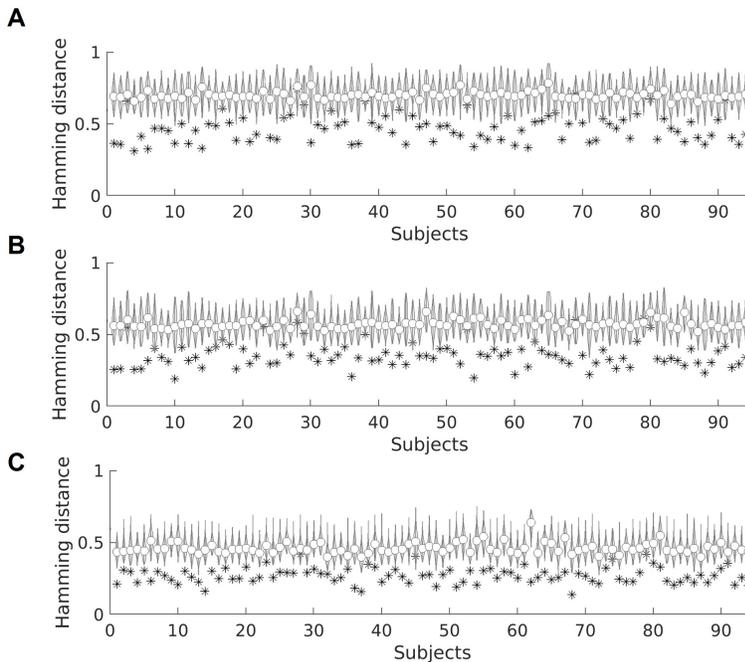


Figure 21: **Individual specificity of network assignment using k-means clustering with 12 networks (ICA)**. A) dimension = 2, B) dimension = 4, C) dimension = 32

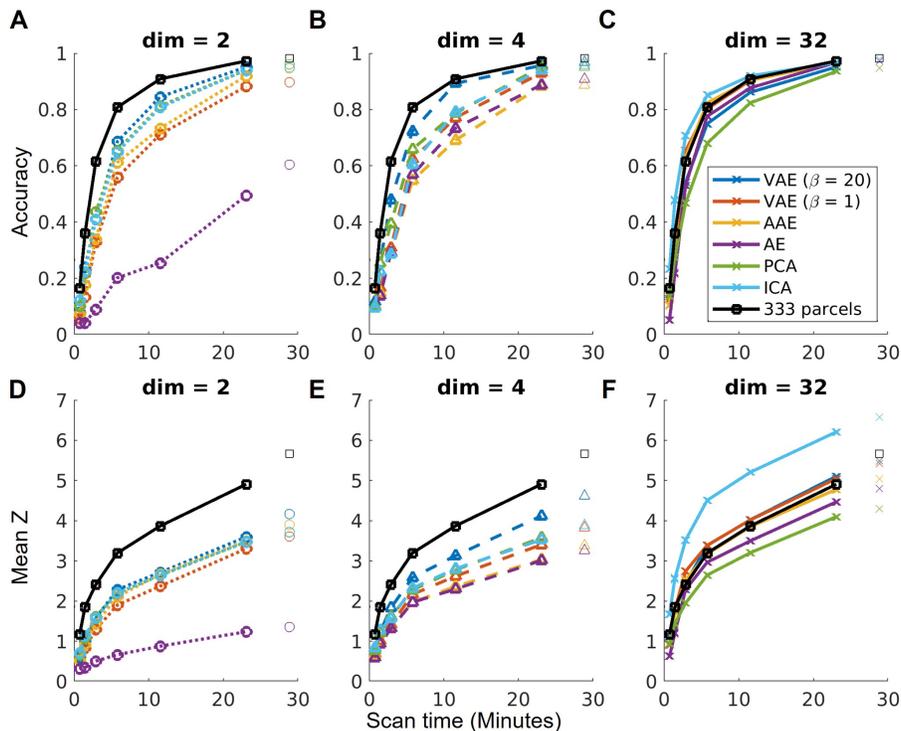
A.13 IDENTIFICATION OF INDIVIDUALS USING THE FINGERPRINTING ANALYSIS

An alternative measure of individual specificity is the “fingerprinting” analysis (Finn et al., 2015). Since the parcel-wise connectome matrix (333 x 333) is symmetric, we concatenated the lower trian-

1890 gle into a vector as a “barcode” for the first and second sessions, respectively. The similarity across
 1891 sessions for the same and different subjects was calculated using Pearson’s correlation between those
 1892 “barcodes”. For the latent representations with dimensions 2,4, and 32, we concatenated the values
 1893 from all dimensions for each session as the “barcode” for the session and repeated the fingerprinting
 1894 analysis. The average accuracy of identification was calculated as the average of two quantities:
 1895 the proportion of correct identification of the subject from the session 1 data, and the proportion of
 1896 correct identification from the session 2 data. We additionally calculated a Z-score which describes
 1897 the relative magnitude of correlation between the sessions from the same subject and the sessions
 1898 from different subjects. A Z-score was calculated for each subject as:

$$1899 \quad Z = \frac{r_{within} - Mean(r_{between})}{SD(r_{between})} \quad (5)$$

1900
 1901 We found that both the accuracy (Figure 22A-C) and within-subject to between-subject contrast
 1902 (Figure 22D-F) increased with the amount of data. Across all dimensions of 2,4 and 32 using VAE,
 1903 PCA, and ICA, the fingerprinting accuracy in latent spaces was similar to that in the parcel space,
 1904 especially for a long scan time (> 20 min). Overall, the difference in accuracy is larger with a shorter
 1905 scan time, with the accuracy from the VAE ranked the highest among the dimensionality reduction
 1906 methods with a marginal difference. Since incorrect identification would occur when even if only
 1907 one other subject had a similar connectome (or connectome latents) to the subject in query, we
 1908 also quantify the within-subject to between-subject contrast which is defined as the Z-score of the
 1909 within-subject correlation normalized by the mean and standard deviation of the between-subject
 1910 correlations. This contrast was higher with increasing latent dimensions and is higher for VAE
 1911 compared to PCA and ICA, especially at higher latent dimensions (Figure 22D-F).

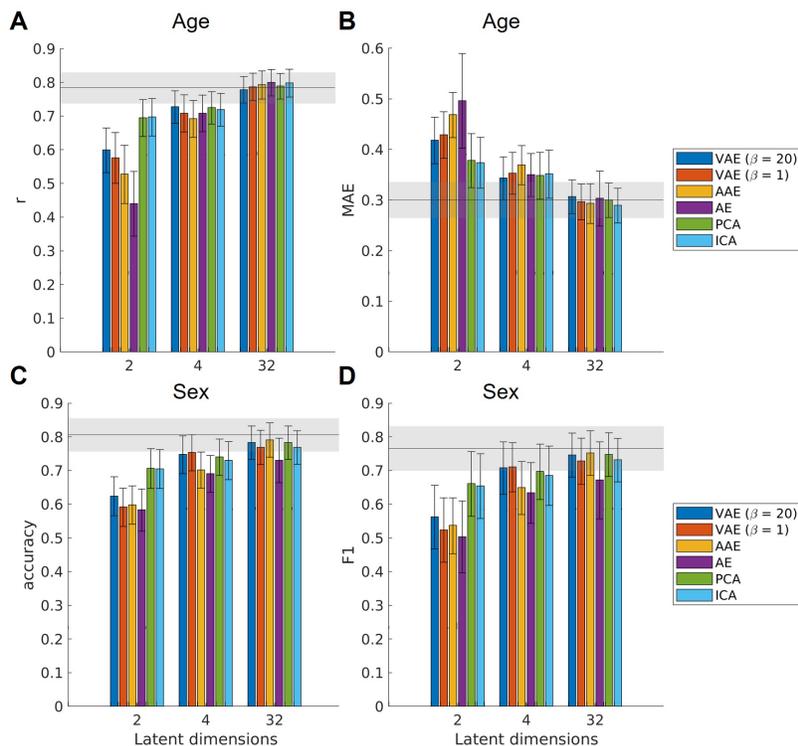


1912
 1913
 1914
 1915
 1916
 1917
 1918
 1919
 1920
 1921
 1922
 1923
 1924
 1925
 1926
 1927
 1928
 1929
 1930
 1931
 1932
 1933
 1934
 1935
 1936
 1937 **Figure 22: Fingerprinting performance for 94 HCP subjects.** A-C) Fingerprinting accuracy. D-F)
 1938 Mean Z-score of within-subject to between-subject contrast.

1939 A.14 PREDICTION OF BEHAVIORAL PHENOTYPE IN INDIVIDUALS

1940
 1941 To demonstrate the ability of the embedded data to retain information about behavioral phenotypes
 1942 in individual functional connectomes, we applied the learned dimensionality reduction mappings to
 1943 a separate infant dataset (Baby Connectome Project(Howell et al., 2019)) with 301 sessions at 8-60

1944 months. We predicted the sex and chronological age of each data sample in the parcel space (326
 1945 parcels with $326 \times (326-1)/2 = 52975$ unique features)(Tu et al., 2024) and the latent space ($326 \times$
 1946 $zdim$ features, where $zdim = 2, 4, 32$) using AE-based or linear dimensionality reduction methods.
 1947 We use a support vector machine regression for age and support vector machine classification for
 1948 sex with 80% of the data used for training and the remaining 20% used in testing. We applied a ridge
 1949 regularization where the hyper-parameter λ was optimized with a 5-fold cross-validation approach
 1950 from a range of 15 λ values from 10^{-5} to 10^{-1} evenly distributed on a log scale. This process
 1951 was repeated 1000 times by randomly splitting training and test data to generate an error bar. The
 1952 age prediction performance was measured with correlation (r) and mean absolute error (MAE). The
 1953 sex prediction performance was measured with accuracy and F1-score. We found that prediction
 1954 performance generally increased with the number of dimensions, but the difference across methods
 1955 was small except for when dimension = 2 (Figure 23).



1981 **Figure 23: Prediction of behavioral phenotype in individuals.** A) Age prediction correlation (r).
 1982 B) Age prediction Mean absolute error (MAE). C) Sex prediction (accuracy). D) Sex prediction (F1-
 1983 score). Error bars show mean and standard deviation across the 1000 samples. The horizontal line
 1984 and shaded area show the performance (mean and standard deviation) using features in the parcel
 1985 space.

1987 A.15 DIMENSIONALITY REDUCTION ON TIME SERIES DATA

1988 In theory, one can run dimensionality reduction directly on spatiotemporal fMRI time series data,
 1989 and then generate the seed map embeddings by correlating the original time series with this low-
 1990 dimensional time series. However, the seed map correlations are likely less noisy than the raw time
 1991 series because it is a summary metric. Latent embeddings obtained from dimensionality reduction
 1992 directly on the seed maps could be easily projected back to original seed map images, which can
 1993 then be overlaid with anatomy to provide intuitive visualizations and neuroscience insights (Figure
 1994 3). Also, doing dimensionality reduction directly on seed maps weighs each subject’s data equally
 1995 even when they have very different data acquisition lengths, and it is less likely that a few subjects
 1996 with long acquisition in the training data would bias the embedding axes. Furthermore, the seed
 1997 map correlations have constrained values ranging between -1 and 1 but the raw time series data can
 have different scales depending on the type of normalization. Direct seed map embedding can also

be applied to higher-level summary data such as the network-average functional connectivity across a group of subjects (Moore et al., 2024) so that individual seed maps can be juxtaposed with those network-average profiles to illustrate their positions relative to those network priors.

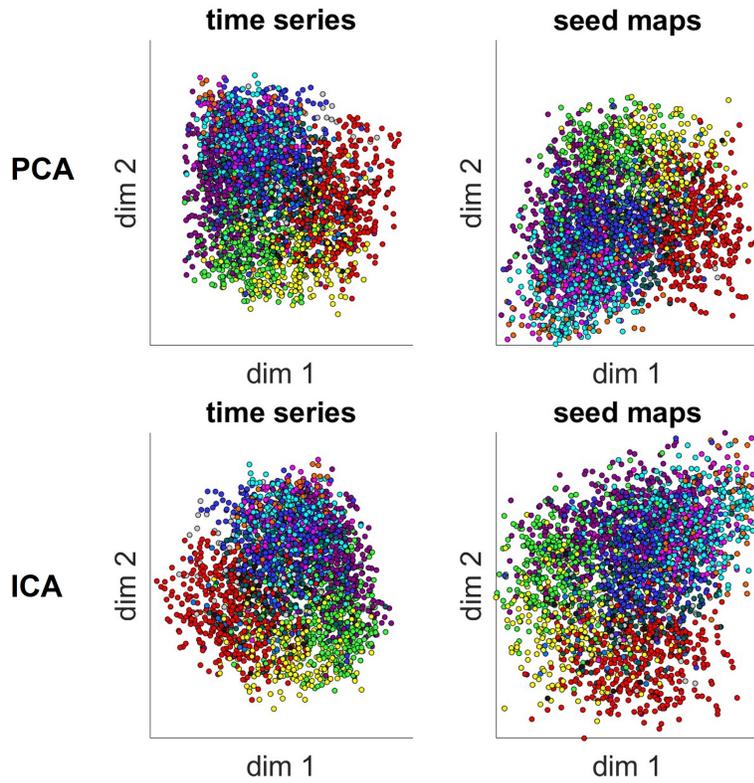
To empirically compare the dimensionality reduction on time series data and seed maps, we also conducted additional experiments using the time series data as the input to the model instead. We now have 27255 samples instead of the 594200 samples (5942 samples per subject x 100 subjects) in the training data. Similarly, the validation data samples decreased from 59420 to 3758. Overall, the reconstruction performance on the validation data is marginally lower when dimensionality reduction was performed on time series than on seed maps (Table 4). This difference in performance may be driven by the difference in the number of samples used to train the model.

	Dimensionality reduction on time series	Dimensionality reduction on functional connectivity
VAE ($\beta = 20$)		
dim = 2	0.57 ± 0.05	0.71 ± 0.10
dim = 4	0.54 ± 0.04	0.74 ± 0.08
dim = 32	0.51 ± 0.01	0.79 ± 0.07
VAE ($\beta = 1$)		
dim = 2	0.57 ± 0.05	0.71 ± 0.09
dim = 4	0.54 ± 0.04	0.73 ± 0.08
dim = 32	0.52 ± 0.02	0.83 ± 0.06
AAE		
dim = 2	0.57 ± 0.06	0.72 ± 0.09
dim = 4	0.61 ± 0.06	0.73 ± 0.08
dim = 32	0.73 ± 0.05	0.83 ± 0.06
AE		
dim = 2	0.57 ± 0.05	0.72 ± 0.09
dim = 4	0.55 ± 0.04	0.73 ± 0.07
dim = 32	0.52 ± 0.02	0.82 ± 0.06
PCA		
dim = 2	0.58 ± 0.07	0.67 ± 0.10
dim = 4	0.63 ± 0.07	0.74 ± 0.09
dim = 32	0.76 ± 0.05	0.84 ± 0.06
ICA		
dim = 2	0.58 ± 0.07	0.67 ± 0.10
dim = 4	0.63 ± 0.07	0.74 ± 0.09
dim = 32	0.76 ± 0.05	0.84 ± 0.06

Table 4: Reconstruction performance (η^2) on the validation data (10 subjects in WU120)

Next, we look at the embedded seed maps for 10 adult subjects from the test dataset of WU120 generated from 1) indirect embedding by embedding the time series ($2 \times T$) and then correlating with the parcel time series ($333 \times T$), and 2) direct embedding. We found that embedding the time series with linear methods produces seed map distributions that are similar to those obtained from directly embedding the seed maps (Figure 24). In addition, we repeated the behavioral phenotype prediction with the new seed map embeddings using the indirect method (a.k.a. embedding the time series first). The performance in predicting age and sex is similar between the direct and indirect embeddings (Figure 25).

2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077



2078 **Figure 24: Seed map latent embeddings using direct or indirect dimensionality reduction.** A) PCA on the time series data. B) PCA on the seed maps. C) ICA on the time series data. D) ICA on the seed maps. The colors correspond to the functional network labels in Figure 4A. The colors follow the network colors in Figure 4.

2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105

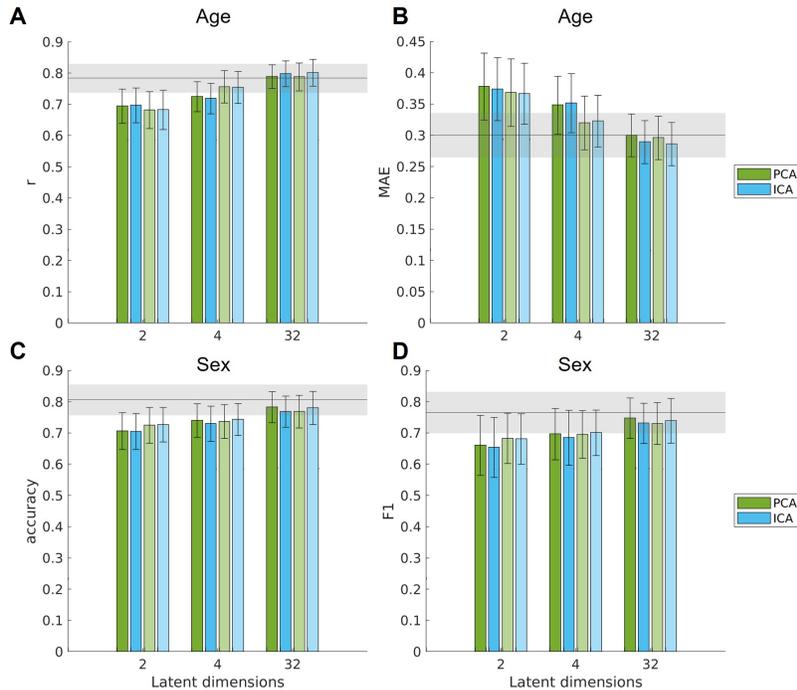


Figure 25: Prediction of behavioral phenotypes using seed map embeddings generated directly or indirectly through dimensionality reduction on the time series. The two darker bars on the left used direct seed map embedding and the two lighter bars on the right used indirect embedding through dimensionality reduction on the time series.