

# Cross Atlas Remapping via Optimal Transport (CAROT): Creating connectomes for any atlas when raw data is not available

Javid Dadashkarimi<sup>a,\*</sup>, Amin Karbasi<sup>a,b</sup>, Qinghao Liang<sup>c</sup>, Matthew Rosenblatt<sup>c</sup>, Stephanie Noble<sup>g</sup>, Maya Foster<sup>c</sup>, Raimundo Rodriguez<sup>e</sup>, Brendan Adkinson<sup>e</sup>, Jean Ye<sup>e</sup>, Huili Sun<sup>c</sup>, Chris Camp<sup>e</sup>, Michael Farruggia<sup>e</sup>, Link Tejavibulya<sup>e</sup>, Wei Dai<sup>c</sup>, Rongtao Jiang<sup>g</sup>, Angeliki Pollatou<sup>h</sup>, Dustin Scheinost<sup>c,d,f,g</sup>

<sup>a</sup>Computer Science Department, Yale University, New Haven, CT, USA

<sup>b</sup>Department of Electrical Engineering, Yale University, New Haven, CT, USA

<sup>c</sup>Department of Biomedical Engineering, Yale University, New Haven, CT, USA

<sup>d</sup>Department of Statistics & Data Science, Yale University, New Haven, CT, USA

<sup>e</sup>Interdepartmental Neuroscience Program, Yale University, New Haven, CT, USA

<sup>f</sup>Child Study Center, Yale School of Medicine, New Haven, CT, USA

<sup>g</sup>Department of Radiology and Biomedical Engineering, Yale University, New Haven, CT, USA

<sup>h</sup>Developing Brain Institute, Children's National Hospital, Washington, DC

## Abstract

Open-source, publicly available neuroimaging datasets—whether from large-scale data collection efforts or pooled from multiple smaller studies—offer unprecedented sample sizes and promote generalization efforts. Releasing data can democratize science, increase the replicability of findings, and lead to discoveries. Due to patient privacy and data storage concerns, researchers typically release preprocessed data with the voxelwise time series parcellated into a map of predefined regions, known as an atlas. However, releasing preprocessed data also limits the choices available to the end-user. This is especially true for connectomics, as connectomes created from different atlases are not directly comparable. Since there exist several atlases with no gold standards, it is unrealistic to have processed, open-source data available from all atlases. Together, these limitations directly inhibit the potential benefits of open-source neuroimaging data. To address these limitations, we introduce Cross Atlas Remapping via Optimal Transport (CAROT) to find a mapping between two atlases. This approach allows data processed from one atlas to be directly transformed into a connectome based on another atlas without the need for raw data access. To validate CAROT, we compare reconstructed connectomes against their original counterparts (i.e., connectomes generated directly from an atlas), demonstrate the utility of transformed connectomes in downstream analyses, and show how a connectome-based predictive model can generalize to publicly available data that was processed with different atlases. Overall, CAROT can reconstruct connectomes from an extensive set of atlases—without ever needing the raw data—allowing already processed connectomes to be easily reused in a wide range of analyses while eliminating redundant processing efforts. We share this tool as both source code and as a stand-alone web application (<http://carotproject.com/>).

**Keywords:** functional connectivity, brain-behavior associations, dataset harmonization, optimal transport

## 1. Introduction

A connectome—a matrix describing the connectivity between any pair of brain regions—is a popular approach used to model the brain as a graph-like structure (Sporns et al., 2004; Bassett and Bullmore, 2006; Bullmore and Sporns, 2009). They are created by parcellating the brain into distinct areas using an atlas (i.e., the nodes of a graph) and estimating the connections between these regions (i.e., the edges of a graph). A wide range of works demonstrates the value of connectomics in studying individual differences in brain function (Elliott et al., 2019; Dubois and Adolphs, 2016), associating brain-behavior associations (Sui et al., 2020; Jiang et al., 2019; Beaty et al., 2018), and understanding brain alterations in neuropsychiatric

disorders (Yan et al., 2019). Overall, connectomes have high potential as a biomarker of various phenotypic information.

Nevertheless, the need for an atlas to create a connectome hinders comparisons across studies and replication and generalization efforts. Different atlases divide the brain into different regions of varying size and topology. Thus, connectomes created from different atlases are not directly comparable. In other words, simply comparing the results from two independent studies that use different atlases is challenging. Further, several atlases exist with no gold standards (Arslan et al., 2018), and more are being developed yearly. Currently, no solutions exist to extend previous results and potential biomarkers to a connectome generated from a different atlas, limiting the broader use of potential connectome-based biomarkers.

Transforming an existing connectome into one generated from a different atlas would help these efforts and increase the utility of existing connectomes. For example, large-scale

\*Corresponding author: Tel.: +1-203-781-6115

projects—like the Human Connectome Project (HCP) (Van Essen et al., 2013), the Adolescent Brain Cognitive Development (ABCD) study (Casey et al., 2018), and the UK Biobank (Sudlow et al., 2015)—share fully processed connectomes. However, the released connectomes for each project are based on different atlases, preventing these datasets from being combined without reprocessing data from thousands of participants. Smaller labs might not have the resources to store and reprocess these data from scratch (Horien et al., 2021). Finally, due to privacy concerns of being able to identify a participant based on unprocessed data, some datasets are only released as fully processed connectomes (Yan et al., 2019). Critically, in this case, it is not possible to go to the data to create connectomes from another atlas. Thus, algorithms to map and transform connectomes have applications for preserving participant privacy and democratizing data access, as well as improving the generalizability of scientific findings.

To this aim, we propose Cross Atlas Remapping via Optimal Transport (CAROT), which uses optimal transport theory, or the mathematics of converting a probability distribution from one set to another, to find an optimal mapping between two atlases. CAROT is designed for functional connectomes based on functional magnetic imaging (fMRI) data. It allows a connectome constructed from one atlas to be directly transformed into a connectome based on a different atlas without needing to access or preprocess the raw data. We define raw data as data in any form other than fully preprocessed timeseries from an atlas, which is the final form of the data used to create a connectome. Fully preprocessed timeseries from an atlas have several benefits over other intermediate forms derived from a connectomic processing pipeline. As these data consist of only 200 – 500 timeseries, they require less storage than voxel-wise or vertex-wise preprocessed data in common space (1 – 3 MB compared to 500 – 1000 MB per individual). These data are also not identifiable if privacy concerns exist.

First, in a training sample with fMRI time series data from two different atlases, we find a mapping by solving the Monge–Kantorovich transportation problem (Kantorovich, 1942). Then, by employing this optimal mapping, time series data based on the first atlas (from individuals independent of the training data) can be reconstructed into connectomes based on the second atlas without ever needing to be preprocessed. To validate CAROT, we compare reconstructed connectomes against their original counterparts (i.e., connectomes generated directly from an atlas), demonstrate the utility of transformed connectomes in downstream analyses, and show how a connectome-based predictive model can be generalized to publicly available data preprocessed with different atlases. Overall, CAROT can reconstruct connectomes from an extensive set of atlases—without ever needing the raw data—enabling comparison across connectome-based results from different atlases and the reuse of already processed connectomes in a wide range of downstream analyses.

This work builds upon two conference papers presented at the 2021 and 2022 International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI) (Dadashkarimi et al., 2021, 2022). The conference pa-

pers present our initial results using optimal transport to map and transform connectomes from different atlases. We expand our previous results by presenting an extensive set of validation studies, increasing the number of atlases tested, and sharing this tool as source code and a stand-alone web application (<http://carotproject.com/>).

## 2. Theory and calculations

### 2.1. Optimal transport

The optimal transport problem solves how to transport resources from one location  $\alpha$  to another  $\beta$  while minimizing the cost  $C$  (Tolstoi, 1930; Hitchcock, 1941; Koopmans, 1949; Gangbo and McCann, 1996). It has been used for contrast equalization (Delon, 2004), image matching (Li et al., 2013), image watermarking (Mathon et al., 2014), text classification (Huang et al., 2016), and music transportation (Flamary et al., 2016). Optimal transport is one of the few methods that provides a well-defined distance metric when the supports of the distributions are different. Other mapping approaches, such as Kullback–Leibler divergence, do not make this guarantee.

The original formulation of the optimal transport problem is known as the Monge problem. Assuming we have some resources  $x_1, \dots, x_n$  in location  $\alpha$  and some other resources  $y_1, \dots, y_m$  in location  $\beta$ , we specify weight vectors  $a$  and  $b$  over these resources and define matrix  $C$  as a measure of pairwise distances between points  $x_i \in \alpha$  and comparable points  $\mathcal{T}(x_i)$ . The Monge problem aims to solve the following optimizing problem (Monge, 1781):

$$\min_{\mathcal{T}} \left\{ \sum_i C(x_i, \mathcal{T}(x_i)) : \mathcal{T}_{\sharp} \alpha = \beta \right\}, \quad (1)$$

where the push forward operator  $\sharp$  indicates that mass from  $\alpha$  moves towards  $\beta$  assuming that weights absorbed in  $b_j = \sum_{\mathcal{T}(x_i) = y_j} a_i$ . An assignment problem when the number of elements in the measures are not equal is a special case of this problem, where each point in  $\alpha$  can be assigned to several points in  $\beta$ .

As a generalization of the Monge problem, the Kantorovich relaxation solves the mass transportation problem using a probabilistic approach in which the amount of mass located at  $x_i$  potentially dispatches to several points in the target (Kantorovich, 1942). An admissible solution for Kantorovich relaxation is defined as  $\mathcal{T} \in \mathbb{R}_+^{n \times m}$  indicating the amount of mass being transferred from location  $x_i$  to  $y_j$  by  $\mathcal{T}_{i,j}$ :

$$U(a, b) = \{ \mathcal{T} \in \mathbb{R}_+^{n \times m} : \mathcal{T} \mathbb{1}_m = a, \mathcal{T}^T \mathbb{1}_n = b \}, \quad (2)$$

where  $\mathbb{1}$  represents a vector of all 1's. An optimum solution is obtained by solving the following problem for a given  $C \in \mathbb{R}^{n \times m}$  (Rubner et al., 2000):

$$L_c(a, b) = \min_{\mathcal{T} \in U(a, b)} \langle C, \mathcal{T} \rangle = \sum_{i,j} C_{i,j} \mathcal{T}_{i,j}. \quad (3)$$

While a unique solution is not guaranteed (Peyré et al., 2019), an optimal solution exists (see proof in Birkhoff (1946); Bertsimas and Tsitsiklis). Kantorovich and Monge problems are equivalent under certain conditions (see proof in Brenier (1991)).

## 2.2. Cross Atlas Remapping via Optimal Transport (CAROT)

CAROT operates by transforming timeseries data from one atlas (labeled the source atlas) into timeseries from an unavailable atlas (labeled the target atlas). This transformation is a spatial mapping between the two atlases. Next, the corresponding functional connectomes can be estimated using standard approaches (e.g., full or partial correlation). Transforming the timeseries data rather than connectomes themselves has two benefits. First, this results in a lower dimensional mapping, which is more robust to estimate. Second, connectomes can be constructed with standard methods (like correlation), guaranteeing properties like semi-positive definite. Directly mapping between connectomes may not guarantee this property.

Formally, let us assume we have training timeseries data consisting of  $T$  timepoints from the same individuals but from two different atlases (atlas  $\mathcal{P}_s$  with  $n$  regions and atlas  $\mathcal{P}_t$  with  $m$  regions). Additionally, let  $\mu_t \in \mathbb{R}^n$  and  $\nu_t \in \mathbb{R}^m$  to be the vectorized brain activity at single timepoint  $t$  based on atlases  $\mathcal{P}_s$  and  $\mathcal{P}_t$ , respectively. For a fixed cost matrix  $C \in \mathbb{R}^{n \times m}$ , which measures the pairwise distance between regions in  $\mathcal{P}_s$  and  $\mathcal{P}_t$ , we aim to find a mapping  $\mathcal{T} \in \mathbb{R}^{n \times m}$  that minimizes transportation cost between  $\mu_t$  and  $\nu_t$ :

$$L_c(\mu_t, \nu_t) = \min_{\mathcal{T}} C^T \mathcal{T} \text{ s.t. } A \underline{\mathcal{T}} = \begin{bmatrix} \mu_t \\ \nu_t \end{bmatrix}, \quad (4)$$

in which  $\underline{\mathcal{T}} \in \mathbb{R}^{nm}$  is vectorized version of  $\mathcal{T}$  such that the  $i + n(j-1)$ 's element of  $\mathcal{T}$  is equal to  $\mathcal{T}_{ij}$  and  $A$  is defined as:

$$A = \begin{matrix} & \begin{matrix} 1 & 2 & \dots & n \end{matrix} \\ \begin{matrix} m \\ n \end{matrix} & \begin{pmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \\ (1 & 1 & \dots & 1) \\ \vdots \\ (1 & 1 & \dots & 1) \end{pmatrix} & \begin{pmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \\ (1 & 1 & \dots & 1) \\ \vdots \\ (1 & 1 & \dots & 1) \end{pmatrix} & \dots & \begin{pmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \\ (1 & 1 & \dots & 1) \\ \vdots \\ (1 & 1 & \dots & 1) \end{pmatrix} \end{matrix} \quad (5)$$

$\mathcal{T}$  represents the optimal way of transforming the brain activity data from  $n$  regions into  $m$  regions. Thus, by applying  $\mathcal{T}$  to every timepoint from the timeseries data of the source atlas, we can estimate the timeseries data of the target atlas. As solving this large linear program is computationally hard (Dantzig, 1983), we use the entropy regularization, which gives an approximation solution with complexity of  $\mathcal{O}(n^2 \log(n) \eta^{-3})$  for  $\epsilon = \frac{4 \log(n)}{\eta}$  Peyré et al. (2019), and instead solve the following:

$$L_c(\mu_t, \nu_t) = \min_{\mathcal{T}} C^T \mathcal{T} - \epsilon H(\mathcal{T}) \text{ s.t. } A \underline{\mathcal{T}} = \begin{bmatrix} \mu_t \\ \nu_t \end{bmatrix}. \quad (6)$$

Specifically, we use the Sinkhorn algorithm—an iterative solution for Equation 6 (Altschuler et al., 2017)—to find  $\mathcal{T}$ . For training data with  $S$  participants and  $K$  time points per participant, first, we estimate the optimal mapping  $\mathcal{T}_{s,k}$ , independently, for time point  $k$  for a given participant  $s$  using Equation 6. Next, we average  $\mathcal{T}_{s,k}$  overall time points and participants to produce a single optimal mapping  $\mathcal{T}$  in the training data (e.g.,  $\mathcal{T} = \frac{1}{|S||K|} \sum_{s=1}^{|S|} \sum_{k=1}^{|K|} \mathcal{T}_{s,k}$ ).

For the cost matrix  $C$ , we used a distance metric (labeled functional distance) that is based on the similarity of pairs of timeseries from the different atlases:

$$C = 1 - \begin{pmatrix} \rho(U_{1,1}, N_{1,1}) & \dots & \rho(U_{1,1}, N_{n,1}) \\ \vdots & \ddots & \vdots \\ \rho(U_{m,1}, N_{1,1}) & \dots & \rho(U_{m,1}, N_{n,1}) \end{pmatrix} \in \mathbb{R}^{m \times n} \quad (7)$$

where  $U_x$  and  $N_x$  are timeseries from  $\mathcal{P}_s$  and  $\mathcal{P}_t$  and  $\rho(U_x, N_y)$  is Spearman correlation between them. To increase a reliable estimation of  $C$ , we calculate the timeseries correlation independently for each individual in the training data and average over these correlations. Functional distance was used over Euclidean distance between nodes for two mains reasons: (i) functional distance does not require having access to the atlas or node locations, which provides greater flexibility should a unknown and unavailable atlas be used, and (ii) spatial proximity in the brain does not guarantee similar function. For example, the medial prefrontal nodes of the default mode network are more correlated with nodes in the posterior cingulate cortex than other nodes in the frontal lobe. Nevertheless, we formally compare the performance of functional and Euclidean distances.

## 3. Material and methods

### 3.1. Evaluation datasets

We evaluated CAROT on six prominent functional atlases from the literature using three datasets, the Human Connectome Project (HCP), the REST-Meta-MDD Consortium, and the Yale Low-Resolution Controls Dataset.

#### 3.1.1. Atlases

The Shen atlas (Shen et al., 2013) was created using functional connectivity data from 45 adult participants. The 268-node atlas was constructed using a group-wise spectral clustering algorithm (derived from the N-cut algorithm) and covers the entire cortex, sub-cortex, and cerebellum. The Craddock atlas (Craddock et al., 2012) was created using functional connectivity data from 41 adult participants. The 200-node atlas was constructed using an N-cut algorithm and covers the entire cortex, sub-cortex, and cerebellum. The Schaefer atlas (Schaefer et al., 2018) was created using functional connectivity data from 744 adult participants from the Genomics Superstruct Project (Holmes et al., 2015). The 400-node atlas was constructed using a gradient-weighted Markov Random Field (gwMRF) model, covering only the cortex. The Brainnetome atlas (Fan et al., 2016) was created using structural connectivity data from 40 adult participants from the HCP. The 246-node atlas was constructed using a tractography-based approach and covers the cortex and sub-cortex. The Dosenbach atlas (Dosenbach et al., 2010) was created from meta-analyses of task-related fMRI studies and consists of 160 nodes that cover the cortex, cerebellum, and a few sub-cortical nodes. The Power atlas (Power et al., 2011) was created by combining the meta-analytical approach of the Dosenbach atlas with areal boundary detection based on functional connectivity data. The 264-node atlas covers the cortex, sub-cortex, and cerebellum.

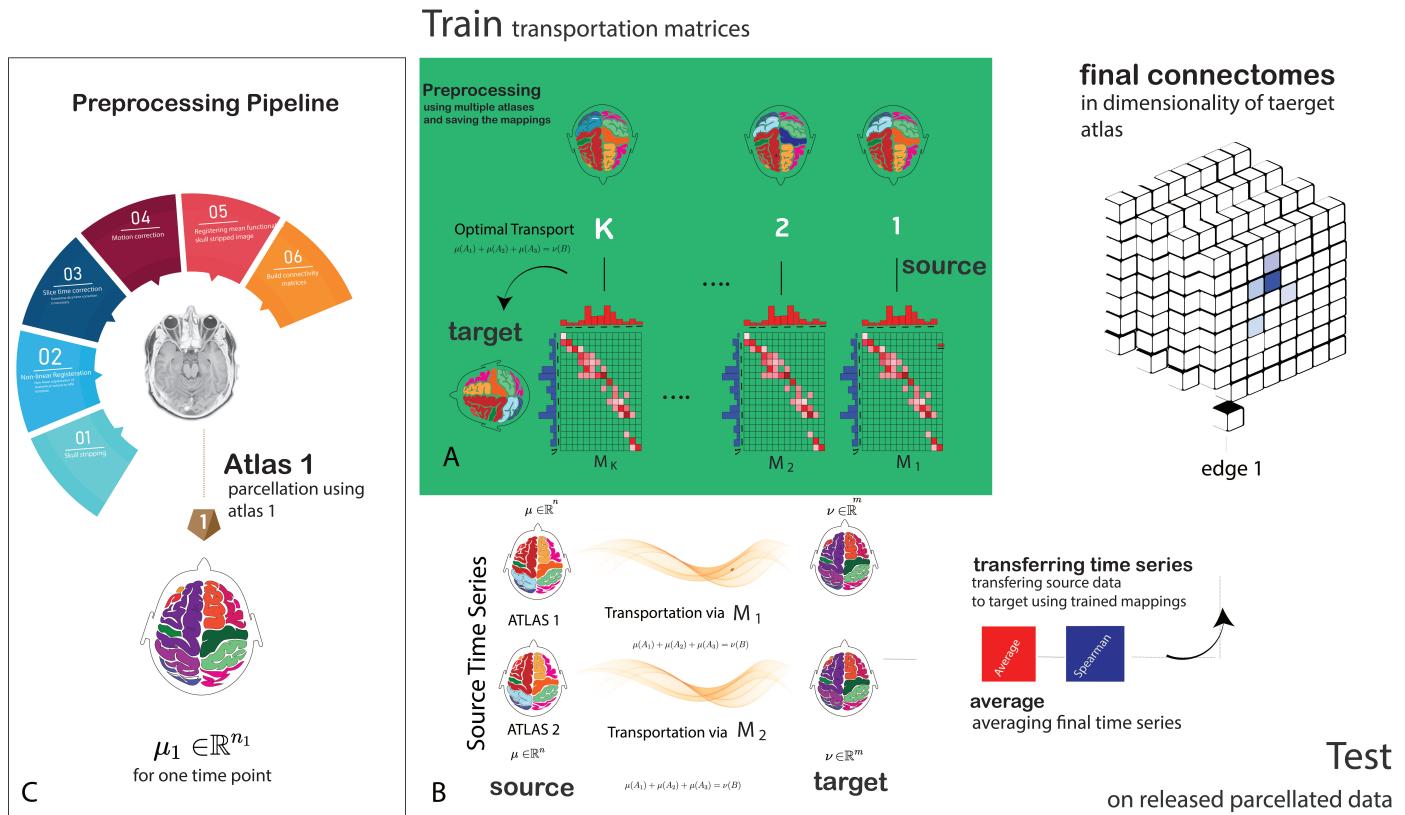


Figure 1: Schematic of CAROT: A) During training, CAROT transforms timeseries fMRI data from multiple source atlases to a target atlas to obtain transportation mappings. Mappings between the source and target atlases are found by employing optimal transport and solving Monge–Kantorovich transportation problem using the Sinkhorn approximation. The solution provides a transformation that maps the brain activity parcellated using the source atlas to brain activity parcellated based on the target atlas. B) During testing, for each pair of source and target atlases and a single time point in the timeseries data, the offline solutions are used, and time series and functional connectomes accordingly will be reconstructed in the desired target atlas. Results from several pairs of source and target atlases can be combined to improve the quality of the final reconstructed connectome. C) A standard image preprocessing pipeline to create functional connectomes.

### 3.1.2. HCP participants

We used behavioral and functional imaging data from this data set as previously described (Gao et al., 2019). We restricted our analyses to those subjects who participated in all nine fMRI conditions (seven tasks, two rest), whose mean frame-to-frame displacement was less than 0.1mm and whose maximum frame-to-frame displacement was less than 0.15mm, and for whom IQ measures were available ( $n=515$ ; 241 males; ages 22–36+). The HCP minimal preprocessing pipeline was used on these data, which includes artifact removal, motion correction, and registration to common space (Glasser et al., 2013). All subsequent preprocessing was performed in BioImage Suite (Joshi et al., 2011) and included standard preprocessing procedures, including removal of motion-related components of the signal; regression of mean time courses in white matter, cerebrospinal fluid, and gray matter; removal of the linear trend; and low-pass filtering.

### 3.1.3. REST-meta-MDD

Fully processed data was downloaded from <http://tfmri.org/REST-meta-MDD>. Full details about the dataset have been previously published elsewhere (Yan et al., 2019). We used data from 21 of the 24 sites. Two sites were removed due to large imbalance between male and female

participants (i.e., < 30% male or female; sites 2 and 12). One site was removed as self-reported sex was not provided (site 4). Briefly, the data was processed as follows. First, the initial 10 volumes were discarded, and slice-timing correction was performed. Then, the time series of images for each subject were realigned using a six-parameter linear transformation. After realignment, individual T1-weighted images were co-registered to the mean functional image using a 6 degrees-of-freedom linear transformation without re-sampling and then segmented into gray matter, white matter, and cerebrospinal fluid. Finally, transformations from individual native space to MNI space were computed with the Diffeomorphic Anatomical Registration Through Exponentiated Lie algebra (DARTEL) tool. To minimize head motion confounds, the Friston 24-parameter model was regressed from the data. Scrubbing (removing time points with  $FD > 0.2\text{mm}$ ) was also utilized to verify results using an aggressive head motion control strategy. Other sources of spurious variance (global, white matter, and CSF signals) were also removed from the data through linear regression. Additionally, linear trend were included as a regressor to account for drifts in the blood oxygen level dependent (BOLD) signal. Temporal bandpass filtering (0.01–0.1Hz) was performed on all time series.

### 3.1.4. Yale participants

In addition, we used resting-state data collected from 100 participants at the Yale School of Medicine. This dataset included 50 females (age=33.3±12.3) and 50 males (age=34.9±10.1) with eight functional scans (48 minutes total). The dataset and processing details can be found in (Scheinost et al., 2014). Briefly, standard preprocessing procedures were applied to these data. Structural scans were skull stripped using an optimized version of the FMRIB’s Software Library (FSL) pipeline. Slice time and motion correction were performed in SPM8. The remainder of image preprocessing was performed in BioImage Suite. The data was cleaned by regressing nuisance variables (motion parameters, drift terms, and the mean time courses of the white matter, cerebrospinal fluid, and gray matter signals) and band-pass filtering) and was nonlinearly registered to the MNI template.

### 3.1.5. Generating connectomes

After processing, the Shen, Schaefer, Craddock, Brainnetome, Power, and Dosenbach atlases were applied to the pre-processed fMRI data to create mean timeseries for each node. For each atlas and dataset, connectomes were generated by calculating the Pearson’s correlation between each pair of these mean timeseries and then tasking the fisher transform of these correlations. Connectomes reconstructed by CAROT were also generated using Pearson’s Correlation.

## 3.2. Evaluation overview

We performed several evaluations of CAROT. First, we performed a baseline evaluation of CAROT, investigating the similarity of the original and reconstructed connectomes, the impact of free parameters (e.g., the number of participants used to train CAROT), and the number of available source atlases. Second, we investigate how reconstructed connectomes perform in standard downstream analyses (i.e., do reconstructed connectomes give similar neuroscience results as the original connectomes?). Finally, we present a real-world evaluation of how CAROT can generalize a preexisting connectome-based predictive model when data from the required atlas is unavailable.

## 3.3. Baseline evaluation of CAROT

### 3.3.1. Similarity between the original and reconstructed connectomes

We compared reconstructed connectomes to their original counterpart using HCP data. We partitioned our data into a 25/75 split, where 25% of the individuals is used to estimate the optimal mapping  $\mathcal{T}$  and 75% is used to evaluate the reconstructed connectomes. Reconstructed connectomes were created using single and multiple source atlases. To evaluate the similarity between CAROT reconstructed and original connectomes, the upper triangles of the connectomes were vectorized and correlated with Spearman’s rank correlation.

### 3.3.2. Evaluation of free parameters

We investigated the sensitivity of CAROT to the number of time points and number of participants used to find the mappings and the value of  $\epsilon$  in the Sinkhorn approximation. Using

the same 25/75 split of the HCP participants for training and testing as above, we varied the number of time points used from 100 to 1100 in increments of 100, varied the number of participants from 100 to 515 in increments of 100, and varied  $\epsilon$  from 0.01 to 10 in increments of 1.

### 3.3.3. Extending CAROT for multiple atlases

A vital drawback of the single-source optimal transport is that it relies on a single pair of source and target atlases (i.e., one source atlas and one target atlas), which ignores additional information when multiple source atlases exist. As preprocessed data is often released with timeseries data from multiple atlases (Yan et al., 2016), we investigated using these additional data to better reconstruct connectomes from an unavailable atlases. To incorporate multiple atlases, we applied CAROT to each pair of atlases, transforming timeseries data for each source atlas into the timeseries data for the target atlas. Next, the transformed timeseries data are averaged across all source atlases, and a single connectome for the target atlas is created (Fig. 1). Further, we investigated the impact of using a smaller number of source atlases by only including  $k$  random source atlases when creating a connectome for the target atlas. This process was repeated with 100 iterations over  $k = 2-6$ .

### 3.3.4. Generalizing mappings across datasets

We investigated if CAROT mappings trained in one dataset generalize to other datasets. In other words, we tested if CAROT can be trained once (for example, using the HCP) and then be applied to any new datasets without the need to re-run CAROT (for example, the Yale dataset). First, we trained CAROT using only the HCP dataset. Then, we reconstructed connectomes using the Yale dataset using these  $\mathcal{T}$ ’s. Spearman’s rank correlation between the upper triangles of the connectomes was used to assess the similarity between the reconstructed and original connectomes.

## 3.4. Evaluation of downstream analyses

### 3.4.1. Consistency of aging results

We tested that the reconstructed connectomes produced consistent neuroscience results compared to the original connectomes. First, we used 25% of the HCP participants to train CAROT. Next, using the original and reconstructed connectomes for the remaining 75% of participants, we calculated the association between connectomes and age using mass univariate, edge-wise correlations. Results were threshold at  $P < 0.05$ , corrected for multiple comparisons using the Network-based Statistic (NBS) (Zalesky et al., 2010). To assess whether the overlap of the significant edges found using the original and reconstructed connectomes was statistically significant (i.e., edge-level), we calculated the probability of the overlap being due to chance using the hypergeometric cumulative distribution:

$$F = \sum_{i=0}^K \frac{\binom{K}{i} \binom{M-K}{N-i}}{\binom{M}{N}},$$

where  $F$  is the probability of drawing up to  $i$  of a possible  $K$  items in  $N$  drawings without replacement from a group of  $M$

objects. The p-value for the significance of overlap is then calculated as  $1 - F$ . We also assess the similarity of results at the node-level by summing over all significant edges for a node (i.e., the network theory measure—degree) and correlating these maps for results from the original and reconstructed connectomes.

### 3.4.2. IQ prediction

To show that meaningful brain-phenotype associations are retained in reconstructed connectomes, we used reconstructed connectomes to predict fluid intelligence using connectome-based predictive modeling (CPM) (Shen et al., 2017). We partitioned the HCP dataset into three groupings:  $g_1$ , consisting of 25% of the participants;  $g_2$ , consisting of 50% of the participants; and,  $g_3$ , consisting of the final 25% of the participants. In  $g_1$ , we trained CAROT for each source and target atlases pair. We then applied the learned  $\mathcal{T}$  on  $g_2$  and  $g_3$  to estimate connectomes for each target atlas, resulting in nine connectomes for each atlas (seven reconstructed connectomes based on a single source atlas, one reconstructed connectome based on all source atlases, and the original connectome). Finally, we trained a CPM model of fluid intelligence for each set of connectomes using  $g_2$  and tested this model in  $g_3$ . Fluid intelligence was quantified using a 24-item version of the Penn Progressive Matrices test. Spearman correlation between observed and predicted values was used to evaluate prediction performance. This procedure was repeated with 100 random splits of the data into three groups.

### 3.4.3. Identification rate

We investigated if the individual uniqueness of connectomes is retained in reconstructed connectomes by identifying individuals scanned on repeated days (Finn et al., 2015). As mentioned above, we used the HCP data and a 25/75 split to create reconstructed connectomes based on all available source atlases. In an iterative process, one individual's connectome was selected from the target set and compared against each of the connectivity matrices in the database to find the matrix that was maximally similar. Spearman correlation between the target connectome and each in the database was used to assess similarity. A score of 1 was assigned if the predicted identity matched the true identity, or 0, if it did not. Each target connectome was tested against the database in an independent trial. Connectomes generated from the day 1 resting-state data were used as the target set, and connectomes generated from the day 2 resting-state data were used as the database. We performed this identification procedure for the original and reconstructed connectomes independently. We used permutation testing to generate a null distribution to determine if identification rates were achieved at above-chance levels. Specifically, participants' identities were randomly shuffled and identification was performed with these shuffled labels. Identification rates obtained using the correct labels were then compared to this null distribution to determine significance.

### 3.5. Real-world evaluation

In this evaluation, we generalize a sex classification model (using 100 adults collected at the Yale School of Medicine and

created with the Shen atlas) to the REST-Meta-MDD dataset (Yan et al., 2016), which only provides preprocessed timeseries data from the Dosenbach, Power, and Craddock atlases. First, we trained the sex classification model using the Yale dataset's resting-state data from 100 individuals (50 males). We trained a  $\ell_2$ -penalized logistic regression model with 10-fold cross-validation to classify self-reported sex. Then, we used  $\mathcal{T}$  estimated from the HCP to transform the publicly available preprocessed data (i.e., timeseries data from the Dosenbach, Power, and Craddock atlases) from the REST-Meta-MDD dataset into the Shen atlas. Data from each source atlas were combined to create a single connectome based on the Shen atlas for the 1005 (585 females) health controls. Finally, the sex classification model created in the Yale dataset was applied to these reconstructed Shen connectomes.

### 3.6. Data availability

All datasets used in this study are open-source: HCP (ConnectomeDB database, <https://db.humanconnectome.org>), REST-meta-MDD (<http://rfmri.org/REST-meta-MDD>), and Yale dataset ([http://fcon\\_1000.projects.nitrc.org/indi/retro/yale\\_lowres.html](http://fcon_1000.projects.nitrc.org/indi/retro/yale_lowres.html)). BioImage Suite tools used for processing can be accessed at (<https://bioimagesuiteweb.github.io/>). CAROT and associated canonical mappings are on GitHub (<https://github.com/dadashkarimi/carot>). The Python Optimal Transport (POT) toolbox is available at <https://pythonot.github.io/>.

## 4. Results

### 4.1. Baseline evaluation of CAROT

#### 4.1.1. Reconstructed connectomes are similar to original connectomes

As shown in Table 1, the correlation between the reconstructed connectomes and their original counterparts depends on the atlas pairing, with more similar atlases appearing to have higher correlations. For instance, a strong correlation is seen while transforming data from the Craddock atlas to the Shen atlas ( $\rho = 0.48$ ,  $p < 0.001$ ). Both atlases are based on clustering timeseries fMRI data using variants of the N-cut algorithm. In contrast, a weaker correlation exists for transforming data from the Dosenbach atlas (which was constructed based on a meta-analysis of task activations) to the Shen atlas ( $\rho = 0.24$ ,  $p < 0.001$ ). Finally, similarity between the reconstructed and original connectomes was much lower when using Euclidean distance (Fig. S1).

#### 4.1.2. Using multiple atlases improves CAROT

Overall, we observed a considerable improvement when including data from multiple source atlases. In every case, using all available data produced more similar connectomes to their original counterparts (all  $\rho$ 's  $> 0.50$ ; Fig. 2). For most atlases, explained variance is more than tripled using CAROT with multiple source atlases compared to using a single source atlas. As shown in Fig. 2, while the similarity between reconstructed and original connectomes increases as the number of source atlases

	Shen	Schaefer	Craddock	Brainnetome	Power	Dosenbach
Shen		0.46 $\pm 0.010$	0.55 $\pm 0.001$	0.42 $\pm 0.003$	0.36 $\pm 0.030$	0.39 $\pm 0.004$
Schaefer	0.31 $\pm 0.001$		0.38 $\pm 0.010$	0.38 $\pm 0.001$	0.33 $\pm 0.010$	0.34 $\pm 0.002$
Craddock	0.48 $\pm 0.010$	0.54 $\pm 0.010$		0.51 $\pm 0.003$	0.43 $\pm 0.002$	0.43 $\pm 0.001$
Brainnetome	0.17 $\pm 0.003$	0.23 $\pm 0.003$	0.23 $\pm 0.002$		0.19 $\pm 0.004$	0.18 $\pm 0.002$
Power	0.24 $\pm 0.002$	0.35 $\pm 0.003$	0.32 $\pm 0.001$	0.29 $\pm 0.070$		0.32 $\pm 0.002$
Dosenbach	0.24 $\pm 0.001$	0.32 $\pm 0.003$	0.28 $\pm 0.020$	0.28 $\pm 0.001$	0.28 $\pm 0.010$	

Table 1: Spearman correlation between reconstructed connectomes and original connectomes for each source-target pair. Presented results show as mean  $\pm$  standard deviation over 100 random splitting the data into training and testing sets.

increases, strong correlations (e.g.,  $\rho > 0.6$ ) can be observed with as little as two or three source atlases, suggesting that a small number of atlases may be sufficient for most applications.

#### 4.1.3. CAROT is insensitive to parameter choices

No clear pattern of performance change was observed across the tested parameter range, suggesting that CAROT is not affected the number of frames and participants, and the range of  $\epsilon$ 's (Fig. S2). However, using only 100 participants and 100 time points significant ( $p < 0.05$ ) reduced the processing time from 2,975 s to 467s.

#### 4.1.4. Mappings generalize across datasets

When applying the mapping trained in the HCP dataset to the Yale dataset, we observed a strong correspondence between the reconstructed connectomes and their original counterparts with  $\rho$ 's  $> 0.50$  (Shen:  $\rho = 0.59$ ; Schaefer:  $\rho = 0.66$ ; Craddock:  $\rho = 0.71$ ; Brainnetome:  $\rho = 0.54$ ; Power:  $\rho = 0.50$ ; Dosenbach:  $\rho = 0.54$ ). Notably, these correlations are in the same range as those observed when applied these mappings to the HCP data (i.e., the same dataset used for training the mappings). Together, these results exhibit that mappings can be trained in one dataset and applied to another.

### 4.2. Evaluation of reconstructed connectomes in downstream analyses

#### 4.2.1. Similar patterns of aging are found with reconstructed connectomes

At the edge-level, the reconstructed connectomes for all atlases produced aging results that significantly overlapped with the results from using the original connectomes ( $p < 0.00001$ ). Similarly, node-level correlations were all significant ( $r$ 's  $> 0.60$ ,  $p$ 's  $< 0.001$ ). Fig ?? shows a representative example of node-level results between the reconstructed and original connectomes for the Shen atlas.

#### 4.2.2. Reconstructed connectomes predict IQ

In all cases, connectomes reconstructed using all source atlases performed as well in prediction as the original connectomes (Fig 4A). The reconstructed connectomes using all source atlases performed better than the original connectomes for the Schaefer and Power atlases. Similar to other analyses, connectomes reconstructed from a single atlas varied in prediction performance, depending on the combination of source and target atlases.

#### 4.2.3. Reconstructed connectomes are unique to an individual

For all analyses, identification of individuals demonstrated a high success rate that were significantly greater than chance (5%;  $p < 0.001$ ; based on permutation testing). (Fig. 4B). Reconstructed connectomes performed slightly better than the originals (original connectomes: mean rate=79%; reconstructed connectomes: mean rate=90%). Overall, these results suggest that the reconstructed connectomes retain similar levels of individual differences as their original connectome counterparts.

### 4.3. CAROT facilitates external validation of connectome-based predictive models

Overall, the sex classification model demonstrated significant classification accuracy in the Yale dataset (Accuracy=60.5%  $\pm$  6%; Naive model accuracy=50%;  $\chi^2 = 5.8$ ;  $p = 0.03$ ). Next, the sex classification model preformed significantly better than chance in the REST-Meta-MDD dataset when using the reconstructed connectomes (Accuracy=66.5%; Naive model accuracy=52.3%;  $\chi^2 = 13.9$ ;  $p = 0.0002$ ). To better this result into context, we created connectomes for the Dosenbach, Power, and Craddock atlases in the Yale dataset, created a sex classification model for connectome type, and generalized these models to the REST-Meta-MDD dataset. The generalization accuracy of reconstructed connectomes (Shen: 66.5%) was numerically superior to the generalization accuracies based on original connectomes (Dosenbach: 59.6%, Power: 59.0%, and Craddock: 64.5%), suggesting that using CAROT and reconstructed connectomes perform as well as original connectomes in generalizing a preexisting predictive model.

### 4.4. Software availability and implementation

To facilitate open science and the broader adoption of CAROT, we have created <http://carotproject.com/>. This web application allows end-users to convert timeseries data from the Shen, Schaefer, Craddock, Brainnetome, Power, and Dosenbach atlases to connectomes for any of the other atlases. As a web application, it works without software installation and across multiple platforms (e.g., Windows, Linux, MacOS, Android). The only requirement is a modern web browser, such as Google Chrome. Please note, that any data used on <http://carotproject.com/> remains on the local computer and is never uploaded or stored on a remoted server. In addition, we provide the CAROT software and associated canonical mapping as opensource at <https://github.com/dadashkarimi/carot/>.

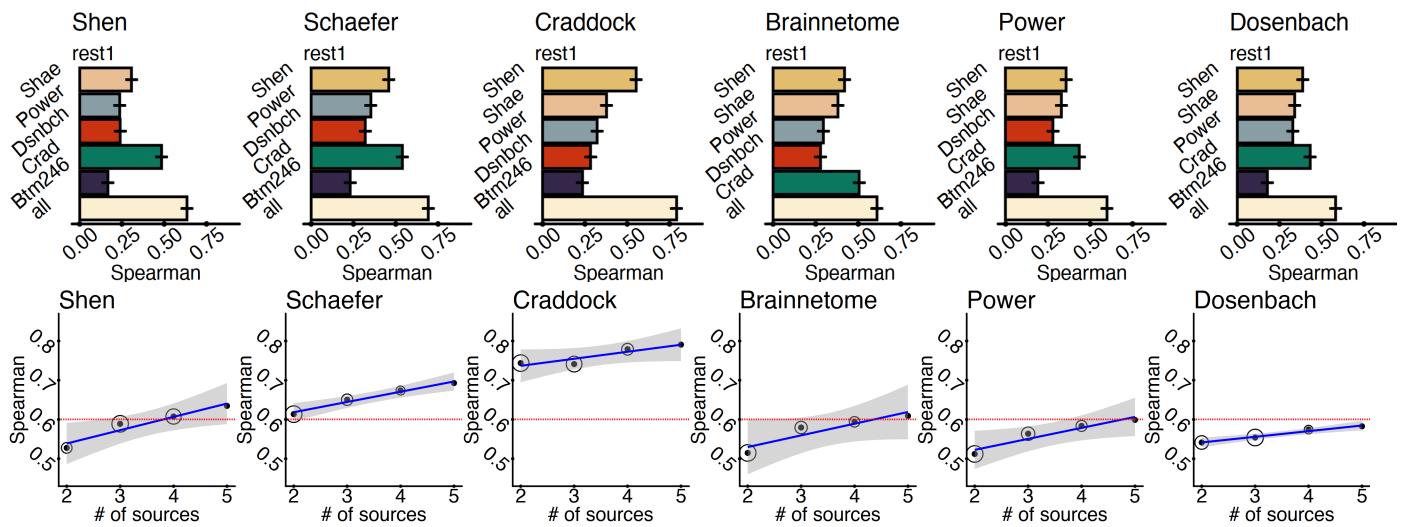


Figure 2: Using multiple source atlases improves the similarity of reconstructed connectomes. A) The Spearman's rank correlation between the reconstructed connectomes and connectomes generated directly with the target atlases are shown for each pair of source and target atlas as well as reconstructed connectomes using all of the source atlases. Using all source atlases produces higher quality reconstructed connectomes for each target atlases. Error bars are generated from 100 iterations of randomly splitting the data into 25% for training and 75% for testing. B) For each target atlas, increasing the source atlases increases the similarity of reconstructed and original connectomes. For most atlases, a Spearman's correlation of  $\rho > 0.60$  (red line) can be achieved by using fewer than five source atlases (i.e., all available source atlases). Circle size represents the variability of the correlation over 100 iterations of splitting the data into training and testing sets.

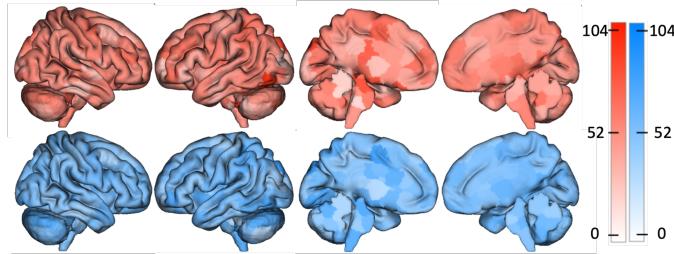


Figure 3: Reconstructed connectomes give similar aging results as the original connectomes. The top row shows the nodes with the largest number of edges significantly associated with age for original connectomes from the HCP created with the Shen atlas. The bottom row shows the same but using reconstructed Shen connectomes. These spatial maps correlate at  $r = 0.61$ , suggesting that analyses with the reconstructed connectomes produce the same neuroscientific insights as analyses with the original connectomes.

Specifically, we provide functionality: (i) to generate the cost matrix based on functional distance for timeseries data from two different atlases; (ii) to generate the mapping  $\mathcal{T}$  between two atlases based on the cost matrix defined above; and (iii) to convert timeseries data from one or more source atlases to connectomes based on a target atlas. In addition, we provide canonical mappings based on the HCP data to map between every pair of the Shen, Schaefer, Craddock, Brainnetome, Power, and Dosenbach atlases. Based on the results present here, these mappings should work in other datasets, saving researchers the need to regenerate these mappings for themselves. We will look to provide mappings between additional atlases as they become available. CAROT is implemented in Python 3, building on the Python Optimal Transport (POT) toolbox (Flamary and Courty, 2017).

## 5. Discussion and conclusions

Neuroimaging is at a crossroads, facing a need to increase replication efforts and use larger-than-ever samples (Yarkoni, 2009; Szucs and Ioannidis, 2020; Marek et al., 2022). These

are tough challenges for functional connectomics, where connectomes created from different atlases are incomparable. As such, processed connectomes or connectomic results from different atlases must be reprocessed from raw data. Here, we introduced and validated CAROT, a method that will allow us to overcome the limitation of not being able to combine connectomes and results from different atlases. CAROT allows functional connectomes from different atlases to be transformed to a common atlas and combined in downstream analyses. CAROT relies on optimal transport to find a frame-to-frame mapping of fMRI time series data used to create functional connectomes for a missing atlas. We show that these reconstructed connectomes are highly similar to the original ones and perform similarly in downstream analyses. Specially, reconstructed connectomes retain sufficient individual differences to predict IQ and uniqueness to identify individuals. Finally, we provide a real-world example of how a connectome-based predictive model (based on the Shen atlas) can be generalized to open source, preprocessed data that was not processed with the Shen atlas.

Critically, the mappings between connectomes are general to the dataset used to create the mappings. As such, a single set

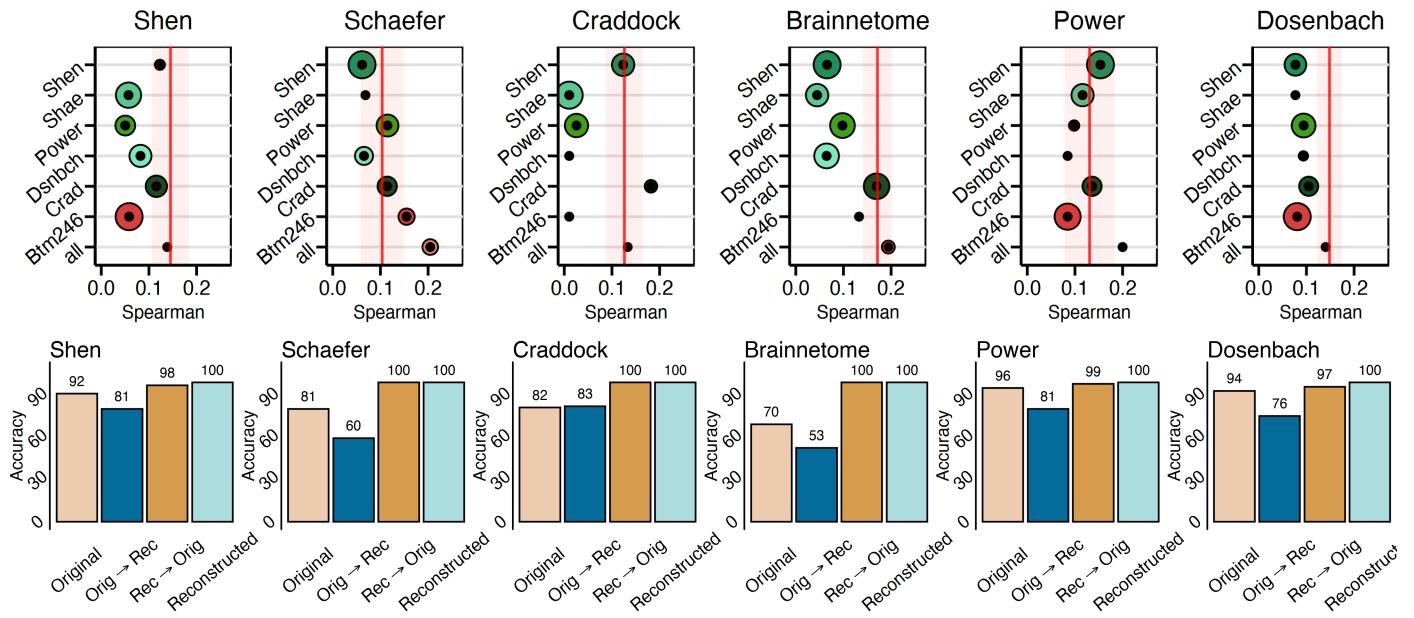


Figure 4: Reconstructed connectomes behave the same as original connectomes in downstream analyses. A) The reconstructed connectomes retain sufficient individual differences to predict IQ using connectome-based predictive modeling. In all cases, reconstructed connectomes based on all available source atlases (bottom circle) predicted IQ as well or better than the original connectome (red line). Size of circle represents the variance of prediction performance of 100 iteration of 10-fold cross-validation. B) The reconstructed connectomes retain sufficient individual uniqueness to identify individuals using the reconstructed connectomes.

of canonical or gold-standard mappings can be trained with one dataset and be distributed to work in new datasets without re-training the mappings. Accordingly, we have released initial mappings based on the HCP data to map between every pair of the Shen, Schaefer, Craddock, Brainnetome, Power, and Dosenbach atlases as part of our software. We hope that CAROT and <http://carotproject.com/> will save researchers time and effort by eliminating data reprocessing and increase the ease of performing mega-analysis and external validation efforts.

Across analyses, we show that CAROT produces reconstructed connectomes that achieve similar results in IQ prediction and fingerprinting as connectomes created directly from the data. This observation holds across a range of atlases that differ in their construction and constituent brain regions. While atlas pairs that are more similar in terms of their construction and coverage produced better pair-wise mappings (e.g., the Craddock and Shen atlases were created with N-cut algorithms and cover the cortex, sub-cortex, and cerebellum (Shen et al., 2013; Craddock et al., 2012)), using multiple source atlases is even better. Likely, combining transformed time series averages out the minor idiosyncrasies in the individual mappings between atlas pairs, producing more stable results. Overall, when using multiple source atlases, CAROT is robust to differences between the source and target atlases.

While including all available data generated the most similar connectomes, strong correspondence between reconstructed and original connectomes was observed when as little as 2 or 3 different source atlases were used. Together, this suggests that an exhaustive list of every possible atlas does not need to be released, but that only including a few different atlases could vastly increase the utility of any released preprocessed data. Balancing the utility of released data and the effort to release

it is a delicate task. If the data is not in a convenient form for end-users, it will not be used and, if the effort is too high to share data, data will not be shared. We believe that CAROT can help balance these, by increasing the utility of the shared data with only a slight increase in effort for sharing the data.

Given that using multiple source atlases produces more robust results, it encourages future studies to release preprocessed data from a few atlases. Not only does this increase the chances that the needed atlas is available for an end-user, but also better facilitates the use of the data when the needed atlas is unavailable. Some open source datasets already release data from multiple atlases (e.g., REST-Meta-MDD). That CAROT performs better with multiple atlases may be relevant for large-scale projects, like ABCD and UK Biobank, that share raw data and curated releases. Given that these datasets range in the several thousands of participants, curated data from multiple atlases further facilitates the use of this data by smaller labs and research groups with a more expansive range of atlases and tools. Additionally, CAROT may help with connectome-based meta analyses, for which there are few, by allowing results to be pooled across studies. Coordinate-based meta analyses are popular for task activation and brain morphometry studies (Laird et al., 2005; Yarkoni et al., 2011; Eickhoff et al., 2009) and are possible as most neuroimaging studies rely on a common template (*i.e.*, the MNI template). This common template allows for spatial comparisons and pooling of results across different studies.

There are a few notable strengths and limitations of CAROT. First, CAROT appears to be robust to the choices of algorithmic parameters such as the number of fMRI frames, sample size used for training, the choice of cost matrix, and the equation used to solve the optimal transport problem. We showed that

the method is not sensitive to parameter search in part due to the large amount of spatial and temporal autocorrelation in fMRI data (Shinn et al., 2021), which allow something as complex as a connectome to be compactly parameterized.

Future work includes generalizing CAROT to other functional time series data—such as electroencephalography (EEG), functional near infrared spectroscopy (fNIRS), or even wide-field CA2+ imaging data in mice (Lake et al., 2020)—where spatial and temporal autocorrelation patterns will be different. One limitation is that since CAROT is based on time series data, it is only appropriate for functional connectomes. Nevertheless, the “missing atlas” problem also exists for structural connectomes, for which no solution exists. Hence, the problem still exists for studies looking to uncover structure-function relationships at the connectome level. However, perhaps CAROT based on Euclidean distance rather than functional distance may be a reasonable approach to map between atlases used to create structural connectomes as well as map between different atlases used in morphometric analyses (such as the Desikan-Killiany and Destrieux atlases used in FreeSurfer). While we tested CAROT with an extensive range of atlases, we could not test CAROT in every functional atlas, as there are many. Nevertheless, given the range in atlas size (200-500 nodes) and atlas coverage (whole-brain and cortical only), we expect CAROT to work well for modern atlases not tested here and look to update CAROT when a new generation of brain atlases emerges.

In sum, CAROT allows a connectome generated based on one atlas to be directly transformed into a connectome based on another without needing raw data. These reconstructed connectomes are similar to and, in downstream analyses, behave like the original connectomes created from the raw data. Using CAROT on preprocessed open source data will increase its utility, accelerate the use of big data, and help make generalization and replication efforts easier.

## 6. Acknowledgments

This work was supported by NIMH R01 MH121095, NSF (IIS-1845032), ONR (N00014-19-1-2406), and Tata. Data were provided in part by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; U54 MH091657) and funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University. Data was also provided by the REST-meta-MDD Consortium Data Sharing, which was supported by the National Key R&D Program of China (2017YFC1309902), the National Natural Science Foundation of China (81671774, 81630031, 81471740 and 81371488), the 13th Five-year Informatization Plan (XXH13505) of Chinese Academy of Sciences, Beijing Municipal Science & Technology Commission (Z161100000216152, Z171100000117016, Z161100002616023 and Z171100000117012), Department of Science and Technology, Zhejiang Province (2015C03037) and the National Basic Research (973) Program (2015CB351702). The remainder of the data used in this study were pro-

vided by the Philadelphia Neurodevelopmental Cohort (Principal Investigators: Hakon Hakonarson and Raquel Gur; phs000607.v1.p1). Support for the collection of the data sets was provided by grant RC2MH089983 awarded to Raquel Gur and RC2MH089924 awarded to Hakon Hakonarson.

## References

Altschuler, J., Weed, J., Rigollet, P., 2017. Near-linear time approximation algorithms for optimal transport via sinkhorn iteration. arXiv preprint arXiv:1705.09634 .

Arslan, S., Ktena, S.I., Makropoulos, A., Robinson, E.C., Rueckert, D., Parisot, S., 2018. Human brain mapping: A systematic comparison of parcellation methods for the human cerebral cortex. *NeuroImage* 170, 5–30. doi:<https://doi.org/10.1016/j.neuroimage.2017.04.014>. segmenting the Brain.

Bassett, D.S., Bullmore, E., 2006. Small-world brain networks. *The Neuroscientist* 12, 512–523.

Beaty, R.E., Kenett, Y.N., Christensen, A.P., Rosenberg, M.D., Benedek, M., Chen, Q., Fink, A., Qiu, J., Kwapił, T.R., Kane, M.J., et al., 2018. Robust prediction of individual creative ability from brain functional connectivity. *Proceedings of the National Academy of Sciences* 115, 1087–1092.

Bertsimas, D., Tsitsiklis, J., . Introduction to linear optimization, athena scientific, 1997 .

Birkhoff, G., 1946. Tres observaciones sobre el algebra lineal. *Univ. Nac. Tucuman, Ser. A* 5, 147–154.

Brenier, Y., 1991. Polar factorization and monotone rearrangement of vector-valued functions. *Communications on pure and applied mathematics* 44, 375–417.

Bullmore, E., Sporns, O., 2009. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature reviews neuroscience* 10, 186–198.

Casey, B., Cannonier, T., Conley, M.I., Cohen, A.O., Barch, D.M., Heitzeg, M.M., Soules, M.E., Teslovich, T., Dellarco, D.V., Garavan, H., et al., 2018. The adolescent brain cognitive development (abcd) study: imaging acquisition across 21 sites. *Developmental cognitive neuroscience* 32, 43–54.

Craddock, R.C., James, G.A., Holtzheimer III, P.E., Hu, X.P., Mayberg, H.S., 2012. A whole brain fmri atlas generated via spatially constrained spectral clustering. *Human brain mapping* 33, 1914–1928.

Dadashkarimi, J., Karbasi, A., Scheinost, D., 2021. Data-driven mapping between functional connectomes using optimal transport, in: de Bruijne, M., Cattin, P.C., Cotin, S., Padov, N., Speidel, S., Zheng, Y., Essert, C. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*, Springer International Publishing, Cham. pp. 293–302.

Dadashkarimi, J., Karbasi, A., Scheinost, D., 2022. Combining multiple atlases to estimate data-driven mappings between functional connectomes using optimal transport, in: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, Springer Nature Switzerland, Cham. pp. 386–395.

Dantzig, G.B., 1983. Reminiscences about the origins of linear programming, in: *Mathematical Programming The State of the Art*. Springer. pp. 78–86.

Delon, J., 2004. Midway image equalization. *Journal of Mathematical Imaging and Vision* 21, 119–134.

Dosenbach, N.U.F., Nardos, B., Cohen, A.L., Fair, D.A., Power, J.D., Church, J.A., Nelson, S.M., Wig, G.S., Vogel, A.C., Lessov-Schlaggar, C.N., Barnes, K.A., Dubis, J.W., Feczkó, E., Coalson, R.S., Pruett, Jr, J.R., Barch, D.M., Petersen, S.E., Schlaggar, B.L., 2010. Prediction of individual brain maturity using fMRI. *Science* 329, 1358–1361.

Dubois, J., Adolphs, R., 2016. Building a science of individual differences from fmri. *Trends in Cognitive Sciences* 20, 425 – 443. doi:<https://doi.org/10.1016/j.tics.2016.03.014>.

Eickhoff, S.B., Laird, A.R., Grefkes, C., Wang, L.E., Zilles, K., Fox, P.T., 2009. Coordinate-based activation likelihood estimation meta-analysis of neuroimaging data: A random-effects approach based on empirical estimates of spatial uncertainty. *Human Brain Mapping* 30, 2907–2926. doi:[10.1002/hbm.20718](https://doi.org/10.1002/hbm.20718).

Elliott, M.L., Knott, A.R., Cooke, M., Kim, M.J., Melzer, T.R., Keenan, R., Ireland, D., Ramrakha, S., Poulton, R., Caspi, A., Moffitt, T.E., Hariri, A.R., 2019. General functional connectivity: Shared features of resting-state and task fmri drive reliable and heritable individual differences in functional brain networks. *NeuroImage* 189, 516 – 532.

Fan, L., Li, H., Zhuo, J., Zhang, Y., Wang, J., Chen, L., Yang, Z., Chu, C., Xie, S., Laird, A.R., et al., 2016. The human brainnetome atlas: a new brain atlas based on connectional architecture. *Cerebral cortex* 26, 3508–3526.

Finn, E.S., Shen, X., Scheinost, D., Rosenberg, M.D., Huang, J., Chun, M.M., Papademetris, X., Constable, R.T., 2015. Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity. *Nature neuroscience* 18, 1664.

Flamary, R., Courty, N., 2017. Pot python optimal transport library.

Flamary, R., Févotte, C., Courty, N., Emiya, V., 2016. Optimal spectral transportation with application to music transcription. arXiv preprint arXiv:1609.09799 .

Gangbo, W., McCann, R.J., 1996. The geometry of optimal transportation. *Acta Mathematica* 177, 113–161.

Gao, S., Greene, A.S., Constable, R.T., Scheinost, D., 2019. Combining multiple connectomes improves predictive modeling of phenotypic measures. *NeuroImage* 201, 116038. doi:<https://doi.org/10.1016/j.neuroimage.2019.116038>.

Glasser, M.F., Sotiroopoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., Van Essen, D.C., Jenkinson, M., 2013. The minimal preprocessing pipelines for the human connectome project. *NeuroImage* 80, 105–124. doi:<https://doi.org/10.1016/j.neuroimage.2013.04.127>. mapping the Connectome.

Hitchcock, F.L., 1941. The distribution of a product from several sources to numerous localities. *Journal of mathematics and physics* 20, 224–230.

Holmes, A.J., Hollinshead, M.O., O’Keefe, T.M., Petrov, V.I., Fariello, G.R., Wald, L.L., Fischl, B., Rosen, B.R., Mair, R.W., Roffman, J.L., Smoller, J.W., Buckner, R.L., 2015. Brain genomics superstruct project initial data release with structural, functional, and behavioral measures. *Scientific Data* 2. doi:[10.1038/sdata.2015.31](https://doi.org/10.1038/sdata.2015.31).

Horien, C., Noble, S., Greene, A.S., Lee, K., Barron, D.S., Gao, S., O’Connor, D., Salehi, M., Dadashkarimi, J., Shen, X., et al., 2021. A hitchhiker’s guide to working with large, open-source neuroimaging datasets. *Nature human behaviour* 5, 185–193.

Huang, G., Quo, C., Kusner, M.J., Sun, Y., Weinberger, K.Q., Sha, F., 2016. Supervised word mover’s distance, in: *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pp. 4869–4877.

Jiang, R., Calhoun, V.D., Fan, L., Zuo, N., Jung, R., Qi, S., Lin, D., Li, J., Zhuo, C., Song, M., Fu, Z., Jiang, T., Sui, J., 2019. Gender Differences in Connectome-based Predictions of Individualized Intelligence Quotient and Sub-domain Scores. *Cerebral Cortex* 30, 888–900. URL: <https://doi.org/10.1093/cercor/bhz134>, arXiv:<https://academic.oup.com/cercor/article-pdf/30/3/888/33011437>.

Joshi, A., Scheinost, D., Okuda, H., Belhachemi, D., Murphy, I., Staib, L.H., Papademetris, X., 2011. Unified framework for development, deployment and robust testing of neuroimaging algorithms. *Neuroinformatics* 9, 69–84.

Kantorovich, L., 1942. On the transfer of masses (in russian), in: *Doklady Akademii Nauk*, pp. 227–229.

Koopmans, T.C., 1949. Optimum utilization of the transportation system. *Econometrica: Journal of the Econometric Society* , 136–146.

Laird, A.R., Lancaster, J.L., Fox, P.T., 2005. BrainMap: The social evolution of a human brain mapping database. *Neuroinformatics* 3, 065–078. doi:[10.1385/ni:3:1:065](https://doi.org/10.1385/ni:3:1:065).

Lake, E.M.R., Ge, X., Shen, X., Herman, P., Hyder, F., Cardin, J.A., Higley, M.J., Scheinost, D., Papademetris, X., Crair, M.C., Constable, R.T., 2020. Simultaneous cortex-wide fluorescence ca2 imaging and whole-brain fmri. *Nature Methods* 17, 1262–1271.

Li, P., Wang, Q., Zhang, L., 2013. A novel earth mover’s distance methodology for image matching with gaussian mixture models, in: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1689–1696.

Marek, S., Tervo-Clemmens, B., Calabro, F.J., Montez, D.F., Kay, B.P., Hattoum, A.S., Donohue, M.R., Foran, W., Miller, R.L., Hendrickson, T.J., Malone, S.M., Kandala, S., Feczkó, E., Miranda-Domínguez, O., Graham, A.M., Earl, E.A., Perrone, A.J., Cordova, M., Doyle, O., Moore, L.A., Conan, G.M., Uriarte, J., Snider, K., Lynch, B.J., Wilgenbusch, J.C., Pengo, T., Tam, A., Chen, J., Newbold, D.J., Zheng, A., Seider, N.A., Van, A.N., Metoki, A., Chauvin, R.J., Laumann, T.O., Greene, D.J., Petersen, S.E., Garavan, H., Thompson, W.K., Nichols, T.E., Yeo, B.T.T., Barch, D.M., Luna, B., Fair, D.A., Dosenbach, N.U.F., 2022. Reproducible brain-wide association studies require thousands of individuals. *Nature* doi:[10.1038/s41586-022-04492-9](https://doi.org/10.1038/s41586-022-04492-9).

Mathon, B., Cayre, F., Bas, P., Macq, B., 2014. Optimal transport for secure spread-spectrum watermarking of still images. *IEEE Transactions on Image Processing* 23, 1694–1705.

Monge, G., 1781. Mémoire sur la théorie des déblais et des remblais. *Histoire de l’Académie Royale des Sciences de Paris* .

Peyré, G., Cuturi, M., et al., 2019. Computational optimal transport: With applications to data science. *Foundations and Trends in Machine Learning* 11, 355–607.

Power, J.D., Cohen, A.L., Nelson, S.M., Wig, G.S., Barnes, K.A., Church, J.A.,

Vogel, A.C., Laumann, T.O., Miezin, F.M., Schlaggar, B.L., et al., 2011. Functional network organization of the human brain. *Neuron* 72, 665–678.

Rubner, Y., Tomasi, C., Guibas, L.J., 2000. The earth mover's distance as a metric for image retrieval. *International journal of computer vision* 40, 99–121.

Schaefer, A., Kong, R., Gordon, E.M., Laumann, T.O., Zuo, X.N., Holmes, A.J., Eickhoff, S.B., Yeo, B.T., 2018. Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity mri. *Cerebral cortex* 28, 3095–3114.

Scheinost, D., Finn, E.S., Tokoglu, F., Shen, X., Papademetris, X., Hampson, M., Constable, R.T., 2014. Sex differences in normal age trajectories of functional brain networks. *Human Brain Mapping* 36, 1524–1535. doi:10.1002/hbm.22720.

Shen, X., Finn, E.S., Scheinost, D., Rosenberg, M.D., Chun, M.M., Papademetris, X., Constable, R.T., 2017. Using connectome-based predictive modeling to predict individual behavior from brain connectivity. *nature protocols* 12, 506.

Shen, X., Tokoglu, F., Papademetris, X., Constable, R.T., 2013. Groupwise whole-brain parcellation from resting-state fmri data for network node identification. *Neuroimage* 82, 403–415.

Shinn, M., Hu, A., Turner, L., Noble, S., Achard, S., Anticevic, A., Scheinost, D., Constable, R.T., Lee, D., Bullmore, E.T., Murray, J.D., 2021. Spatial and temporal autocorrelation weave human brain networks. *bioRxiv* doi:10.1101/2021.06.01.446561.

Sporns, O., Chialvo, D.R., Kaiser, M., Hilgetag, C.C., 2004. Organization, development and function of complex brain networks. *Trends in cognitive sciences* 8, 418–425.

Sudlow, C., Gallacher, J., Allen, N., Beral, V., Burton, P., Danesh, J., Downey, P., Elliott, P., Green, J., Landray, M., et al., 2015. Uk biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS medicine* 12, e1001779.

Sui, J., Jiang, R., Bustillo, J., Calhoun, V., 2020. Neuroimaging-based individualized prediction of cognition and behavior for mental disorders and health: Methods and promises. *Biological Psychiatry* 88, 818–828. URL: <https://www.sciencedirect.com/science/article/pii/S0006322320301116>, doi:<https://doi.org/10.1016/j.biopsych.2020.02.016>. neuroimaging Biomarkers of Psychological Trauma.

Szucs, D., Ioannidis, J.P., 2020. Sample size evolution in neuroimaging research: An evaluation of highly-cited studies (1990–2012) and of latest practices (2017–2018) in high-impact journals. *NeuroImage* 221, 117164. doi:10.1016/j.neuroimage.2020.117164.

Tolstoi, A., 1930. Methods of finding the minimal total kilometrage in cargo transportation planning in space. *TransPress of the National Commissariat of Transportation* 1, 23–55.

Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K., Consortium, W.M.H., et al., 2013. The wu-minn human connectome project: an overview. *Neuroimage* 80, 62–79.

Yan, C.G., Chen, X., Li, L., Castellanos, F.X., Bai, T.J., Bo, Q.J., Cao, J., Chen, G.M., Chen, N.X., Chen, W., Cheng, C., Cheng, Y.Q., Cui, X.L., Duan, J., Fang, Y.R., Gong, Q.Y., Guo, W.B., Hou, Z.H., Hu, L., Kuang, L., Li, F., Li, K.M., Li, T., Liu, Y.S., Liu, Z.N., Long, Y.C., Luo, Q.H., Meng, H.Q., Peng, D.H., Qiu, H.T., Qiu, J., Shen, Y.D., Shi, Y.S., Wang, C.Y., Wang, F., Wang, K., Wang, L., Wang, X., Wang, Y., Wu, X.P., Wu, X.R., Xie, C.M., Xie, G.R., Xie, H.Y., Xie, P., Xu, X.F., Yang, H., Yang, J., Yao, J.S., Yao, S.Q., Yin, Y.Y., Yuan, Y.G., Zhang, A.X., Zhang, H., Zhang, K.R., Zhang, L., Zhang, Z.J., Zhou, R.B., Zhou, Y.T., Zhu, J.J., Zou, C.J., Si, T.M., Zuo, X.N., Zhao, J.P., Zang, Y.F., 2019. Reduced default mode network functional connectivity in patients with recurrent major depressive disorder. *Proceedings of the National Academy of Sciences* 116, 9078–9083. doi:10.1073/pnas.1900390116, arXiv:<https://www.pnas.org/content/116/18/9078.full.pdf>.

Yan, C.G., Wang, X.D., Zuo, X.N., Zang, Y.F., 2016. Dpabi: data processing & analysis for (resting-state) brain imaging. *Neuroinformatics* 14, 339–351.

Yarkoni, T., 2009. Big correlations in little studies: Inflated fMRI correlations reflect low statistical power—commentary on vul et al. (2009). *Perspectives on Psychological Science* 4, 294–298. doi:10.1111/j.1745-6924.2009.01127.x.

Yarkoni, T., Poldrack, R.A., Nichols, T.E., Essen, D.C.V., Wager, T.D., 2011. Large-scale automated synthesis of human functional neuroimaging data. *Nature Methods* 8, 665–670. doi:10.1038/nmeth.1635.

Zalesky, A., Fornito, A., Bullmore, E.T., 2010. Network-based statistic: Identifying differences in brain networks. *NeuroImage* 53, 1197 – 1207.

## Supplemental Materials

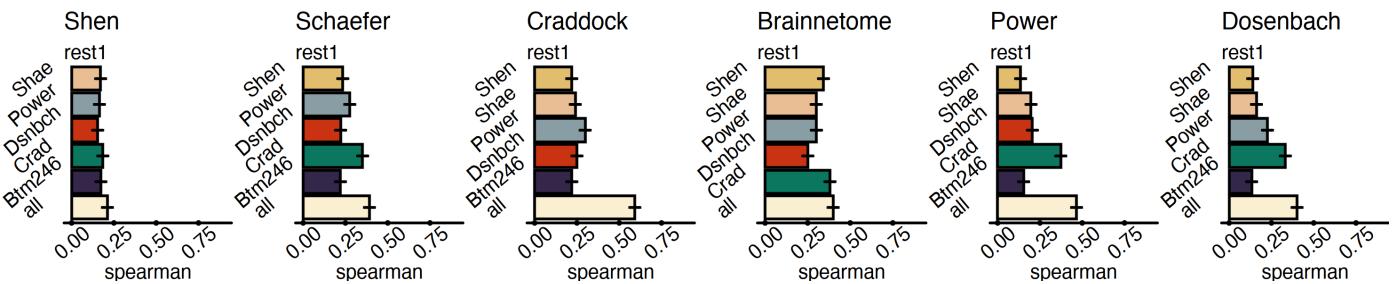


Figure S1: CAROT performance (rest) using euclidean distance between the center of gravity for each ROI as the cost measure. The results exhibit significantly lower performance compared to functional distance.

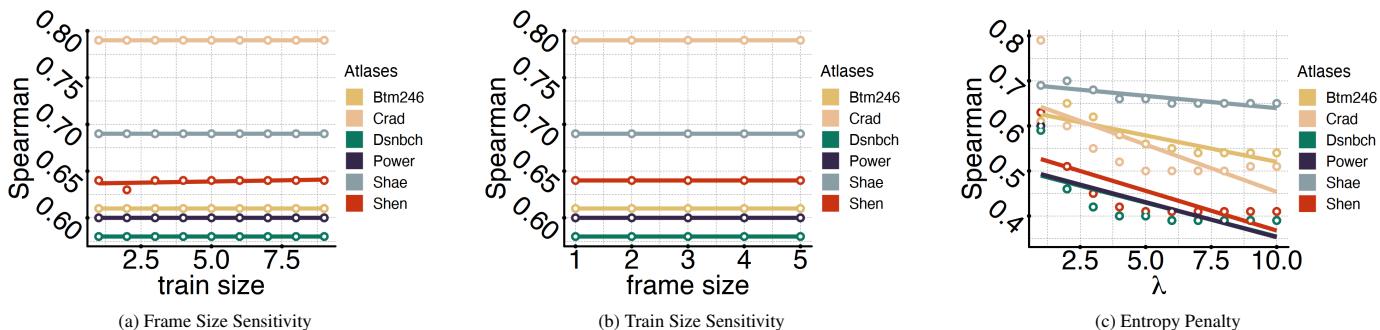


Figure S2: Parameter sensitivity of frame size, training data, and entropy regularization  $\epsilon$  for different target atlases.

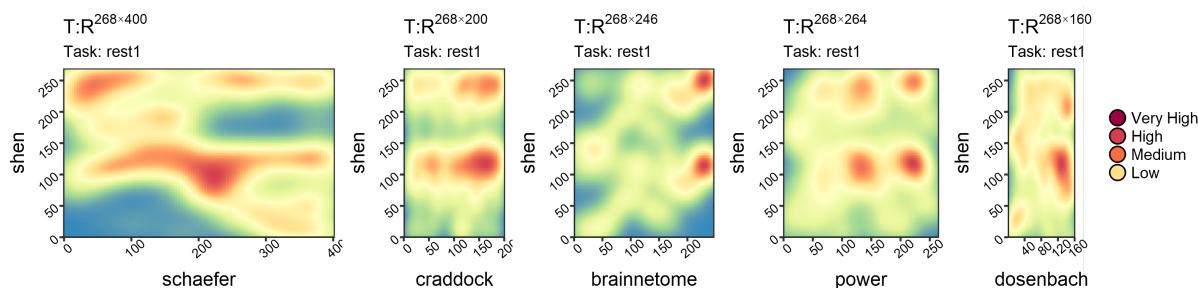


Figure S3: Optimal transport mappings derived from resting-resting and task data from the Shen atlas (source atlas) to each other atlas. Warmer color indicates regions that contribute the most towards mapping between atlases. Horizontal blue areas may indicate locations that are missing in the source atlas. For example, the Schaefer atlas does not include regions in the cerebellum, while the Shen does.