

Overcoming Atlas Heterogeneity in Federated Learning for Cross-site Connectome-based Predictive Modeling

Qinghao Liang¹(✉), Brendan D. Adkinson², Rongtao Jiang³(✉), and Dustin Scheinost^{1,2,3,4}

¹ Department of Biomedical Engineering, Yale University, New Haven, CT, USA

² Interdepartmental Neuroscience Program, Yale University, New Haven, CT, USA

³ Department of Statistics & Data Science, Yale University, New Haven, CT, USA

⁴ Department of Radiology and Biomedical Imaging, Yale School of Medicine, New Haven, CT, USA

⁵ Child Study Center, Yale School of Medicine, New Haven, CT, USA
{qinghao.liang, rongtao.jiang}@yale.edu

Abstract. Data-sharing in neuroimaging research alleviates the cost and time constraints of collecting large sample sizes at a single location, aiding the development of foundational models with deep learning. Yet, challenges to data sharing, such as data privacy, ownership, and regulatory compliance, exist. Federated learning enables collaborative training across sites while addressing many of these concerns. Connectomes are a promising data type for data sharing and creating foundational models. Yet, the field lacks a single, standardized atlas for constructing connectomes. Connectomes are incomparable between these atlases, limiting the utility of connectomes in federated learning. Further, fully reprocessing raw data in a single pipeline is not a solution when sample sizes range in the 10–100’s of thousands. Dedicated frameworks are needed to efficiently harmonize previously processed connectomes from various atlases for federated learning. We present Federate Learning for Existing Connectomes from Heterogeneous Atlases (FLECHA) to address these challenges. FLECHA learns a mapping between atlas spaces on an independent dataset, enabling the transformation of connectomes to a common target space before federated learning. We assess FLECHA using functional and structural connectomes processed with five atlases from the Human Connectome Project. Our results show improved prediction performance for FLECHA. They also demonstrate the potential of FLECHA to generalize connectome-based models across diverse silos, potentially enhancing the application of deep learning in neuroimaging.

Keywords: federated learning, optimal transport, connectomes

1 Introduction

Compiling large datasets at a single location in neuroimaging is challenging due to high costs and time constraints. Consequently, data-sharing will be crucial to

eventually train foundational models using deep learning [1, 2]. However, simply sharing raw data might not lead to these foundational models. Data privacy, ownership, and regulatory compliance concerns persist, preventing the sharing of raw data. Further, there is significant cost and complexity to preprocessing raw imaging data [3]. As sample size grows, it will not be possible to reprocess raw data in a harmonized manner for training models. Methods must be developed to efficiently harmonize previously processed data in a fraction of the time needed to reprocess the data fully.

Connectomes, which measure structural and functional connectivity between brain regions based on an atlas, are a promising data type for creating foundational models. However, the field lacks a standard atlas for constructing connectomes, leading to variability and incomparability between atlases [4, 5]. To share connectomes from multiple sites, either *a priori* harmonization or frameworks to account for different atlases are needed [6, 7]. These needs limit the use of federated learning—a decentralized approach for collaborative training of distributed data—due to connectomic heterogeneity [8, 9]. Federated learning alleviates many privacy and regulatory concerns by sharing only model parameters, not the data.

Bridging heterogeneous domains is challenging due to differences in feature spaces [10]. Techniques like the entropic Gromov-Wasserstein distance [11] and feature transformation [12] are limited by their need for data from both domains, hindering their application in federated learning. Connectomes from different atlases, derived from the same imaging data, can be transformed to a common target space. Inspired by [13, 14], we introduce Federated Learning for Existing Connectomes from Heterogeneous Atlases (FLECHA), which learns a mapping between atlas spaces on an independent dataset and transforms connectomes locally before training.

We evaluate FLECHA’s effectiveness by modeling phenotypes using connectomes processed with five different atlases from the Human Connectome Project, learning the atlas mapping on the Yale dataset. We validate FLECHA by comparing predictive performance of models trained on a single silo and federated models trained across all silos. Results show improved prediction performance in federated learning settings, regardless of imaging modality, prediction task, and machine learning algorithms. FLECHA demonstrates potential for generalizing connectome-based models across silos with varying atlases, advancing the application of deep learning in neuroimaging research.

2 Materials and Methods

2.1 Federate Learning for Existing Connectomes from Heterogeneous Atlases (FLECHA)

FLECHA consists of two steps. The first is to create a mapping between different connectomes in an independent dataset using optimal transport. The second is to use the mapping and federation learning to create a cross-silo predictive model.




Fig. 1: Transforming connectomes across atlases. The mappings between atlases are learned with optimal transport on an independent dataset. Node-wise time-series data yields transportation matrices at each timepoint, with the average providing the final mapping. The source atlas connectomes are decomposed into node factors, transported to the target space using the mapping downloaded online for estimating target connectomes.

Transforming connectomes across atlas space We transform connectomes of the source atlas $\mathbf{C}_s \in \mathbb{R}^{n_s \times n_s}$ to the target atlas, resulting in $\mathbf{C}_t \in \mathbb{R}^{n_t \times n_t}$, n_s and n_t are the number of regions in the source and targets atlas. As shown in Fig. 1, we first learn a mapping $\mathbf{T}$ between the source domain $\Omega_s \in \mathbb{R}^{n_s}$ and target domain $\Omega_t \in \mathbb{R}^{n_t}$ on an independent parcellated fMRI dataset. The mapping is learned by transporting the brain activity at each timepoint between the spaces of the two atlases. Given the a subject's node-wise fMRI timeseries data $\mathbf{V}_s = [v_{s,1}^T, \dots, v_{s,d}^T] \in \mathbb{R}^{n_s \times d}$ of atlas α and $\mathbf{V}_t = [v_{t,1}^T, \dots, v_{t,d}^T] \in \mathbb{R}^{n_t \times d}$ of atlas β , the optimal transport between nodes $\mathbf{T}_i$ at time i is learned by solving the following optimal transport problem:

$$\mathbf{T}_i = \arg \min_{\mathbf{T} \in \mathcal{B}_i} \langle \mathbf{T}, \mathbf{C} \rangle_F$$

, where $\langle \cdot, \cdot \rangle_F$ is the Frobenius dot product and $\mathbf{C} \geq 0$ is the cost matrix. $\mathcal{B}_i$ is the set of probabilistic couplings between the two empirical distributions defined as: $\mathcal{B}_i = \left\{ \mathbf{T} \in (\mathbb{R}^+)^{n_s \times n_t} \mid \mathbf{T} \mathbf{1}_{n_t} = v_{s,i}, \mathbf{T}^T \mathbf{1}_{n_s} = v_{t,i} \right\}$. We used a distance metric based on the similarity of timeseries from different atlases as the cost matrix:

$$C = 1 - \begin{pmatrix} \rho(\sigma_1^s, \sigma_1^t) & \cdots & \rho(\sigma_1^s, \sigma_{n_t}^t) \\ \vdots & \ddots & \vdots \\ \rho(\sigma_{n_s}^s, \sigma_1^t) & \cdots & \rho(\sigma_{n_s}^s, \sigma_{n_t}^t) \end{pmatrix} \in \mathbb{R}^{n_s \times n_t},$$

where σ_i^s and $\sigma_j^t \in \mathbb{R}^{1 \times d}$ are timeseries data of region i of the source atlas and region j of the target atlas, $\rho(\cdot, \cdot)$ is the Pearson correlation. The final mapping is determined by averaging the transport maps across all training subjects and timepoints. We then decompose the source connectome into node factors $\mathbf{U}_s \in$




Fig. 2: FLECHA and the experimental setup for silos with connectomes from different atlases. (a) Initially, connectomes are locally transformed into the target atlas using online mapping resources. The local models, trained on the transformed data, are subsequently aggregated on the server and distributed back to the local nodes for updates. (b) The participants are divided into five partitions (silos) randomly, each of equal size. Different atlases are utilized at each silo. One silo is designated as the testing set, while the remaining four serve as the training set. We conducted a comparative analysis between the federated learning model trained on data from all four silos and models trained independently on each silo. The whole process was repeated 100 times.

$\mathbb{R}^{n_s \times m}$ and map them into the target space $\mathbf{U}_t \in \mathbb{R}^{n_t \times m}$ using the mapping, where m is the latent dimension. The target connectome could be estimated using the multiplication of the node factors.

Federated learning We consider N silos, denoted as $\{P_1, P_2, \dots, P_N\}$. Each silo P_k possesses a private dataset $D_k = \{(x_i^k, y_i^k)\}_{i=1}^{N_k}$, where $|x^k| = N_k$ and the total data points across all silos is $N = \sum_{k=1}^K N_k$. There can be three components in a complete dataset: the feature space $\mathcal{X}$, the label space $\mathcal{Y}$ and the ID space $\mathcal{I}$. The parties may not be identical in the feature space $\mathcal{X}$ and ID space $\mathcal{I}$. In our case: $\mathcal{X}$ is the connectomes of an atlas and label $\mathcal{Y}$ can be the diagnosis or phenotype we want to predict. Since the atlases are different in different silos, $\mathcal{X}_i \neq \mathcal{X}_j$ but $\mathcal{Y}_i = \mathcal{Y}_j$. Typically, each silo P_k has either a pre-trained local network model or an initialized model, represented by $f(w_k)$. Thus, predicts the output for the private sample x^k using the local model $f(w_k)$. In contrast, traditional centralized machine learning frameworks rely on a large, centralized dataset $D = D_1 \cup D_2 \cup \dots \cup D_K$, which integrates the private datasets from each silo

to train a more effective centralized model $f(w)$. Yet, data privacy concerns and the limitations of data silos render traditional centralized learning impractical in privacy-sensitive scenarios. Federated learning offers a solution by allowing each silo to collaboratively train models while keeping their data private.

FLECHA consists of transforming the data to the target space locally and two steps in decentralized optimization: 1) local update, and 2) communication to a global server (Fig. 2a). In this paper, we chose two state-of-the-art federated learning methods: FedAvg [15] and FedProx [8]. FedAvg operates by aggregating the locally updated models from each silo to form a global model. The global model update in FedAvg can be expressed as: $w_{global} = \frac{1}{N} \sum_{k=1}^K N_k w_k$, where w_{global} is the global model, N_k is the number of data points in silo P_k 's dataset, and w_k is the local model of silo P_k . FedProx, on the other hand, introduces a proximal term to the local objective function to mitigate the heterogeneity in local data distributions. The local update in FedProx can be formulated as:

$$w_k^{t+1} = \arg \min_w \left(\mathcal{L}_k(w) + \frac{\mu}{2} \|w - w^t\|^2 \right)$$

, where $\mathcal{L}_k(w)$ is the local loss function for client P_k , μ is a non-negative proximal term coefficient, and w^t is the global model at iteration t . Since linear models are also widely used in connectome modeling, we implemented ridge regression as a comparison. For the federated learning setting, we simply averaged the coefficients of the ridge regression model across silos (Coef_avg).

2.2 Datasets

Yale dataset We used resting-state data from 100 individuals at the Yale School of Medicine. 48 minutes of functional data were collected per individual—50 females (mean age=33.3±12.3) and 50 males (mean age=34.9±10.1). Dataset and processing information can be found in [16]. Using five atlases, we parcellated the brain into regions per each atlas. We computed regional time series for each subject by averaging voxel-wise fMRI data within each region, resulting in an $n \times t$ matrix representing brain activity, with t being 1872 time points.

HCP dataset We used the HCP S900 dataset [17], which included 515 subjects (241 males, ages 22–37) meeting strict criteria: full participation in nine fMRI conditions, minimal motion (mean frame-to-frame displacement ≤ 0.1 mm, max ≤ 0.15 mm), and available working memory measures (WM_Task_0bk_Acc). We also parcellated the brain using the five atlases and computed regional time series. The pairwise Pearson correlation between all node time series was calculated and Fisher z-transformed, yielding a $n \times n$ matrix for each scan. We predicted working memory using functional connectomes in the HCP dataset.

HCP-D dataset The Human Connectome Project in Development (HCP-D) is a comprehensive neuroimaging dataset from ages 5 to 21. This dataset includes 635 subjects (292 males; ages 5-21). Diffusion MRI data were processed

using the SHARD pipeline and reconstructed with generalized q-sampling imaging, followed by whole-brain fiber tracking in DSI-studio, resulting in 1,000,000 streamlines per individual. Structural connectomes were built for each individual using five atlases with a connectivity threshold of 0.001, and pairwise connectivity strength was computed as the average QA value of each connecting fiber, resulting in an $n \times n$ matrix representing each participant’s brain structural connectome. We predicted age using structural connectomes in the HCP-D dataset.

2.3 Atlases

We used five different atlases. The Shen atlas [18] was created from functional connectivity data of 45 adults, resulting in a 268-node atlas covering the entire brain, using a group-wise spectral clustering algorithm. The Craddock atlas [19] was based on data from 41 adults and consists of a 200-node atlas covering the complete brain, constructed using the N-cut algorithm. The Schaefer atlas [20] utilized data from 744 adults, yielding a 400-node atlas focused on the cortex, created through a gradient-weighted Markov Random Field (gwMRF) model. The Brainnetome atlas [21] was derived from structural connectivity data of 40 adults from the HCP, resulting in a 246-node atlas covering both the cortex and sub-cortex constructed using tractography. The Dosenbach atlas [22] was developed through meta-analyses of task-related fMRI studies, comprising 160 nodes covering the cortex, cerebellum, and some sub-cortical regions.

2.4 Experiments

We mapped atlas spaces using parcellated fMRI data from the Yale dataset and conducted predictive modeling on the HCP dataset for functional connectomes and the HCP-D dataset for structural connectomes. Age and working memory were chosen as primary variables due to their ease of measurement, common use in brain-behavior models, and clinical relevance. Structural connectomes are suitable for predicting age, while functional connectomes correlate with higher-order cognitive functions like working memory. Both variables are treated as continuous.

Participants were divided into five equal-sized silos, each using a distinct atlas (Fig. 2b). A 3-layer Multilayer Perceptron (MLP) was trained using federated learning, with four silos for training and one for testing, and compared to models trained independently on each silo. The experiment was repeated 100 times with different partitions, and performance was evaluated using Pearson correlation for its standardization and comparability.

The MLP comprised three fully connected layers with $n(n-1)/2$, 64, and 8 nodes, and a dropout rate of 0.25. Batch normalization and ReLU activation followed the first two layers. Training parameters included mean squared error loss, SGD optimizer (learning rate: 0.001, momentum: 0.9, weight decay: 0.001), and local updates based on communication pace τ . We set 20 total steps per epoch, with the batch size of each silo determined by dividing the training data by 20. Models were aggregated every $\tau = 5$ steps.




Fig. 3: Prediction performance of the average ridge regression model (Coef_avg) and models trained independently on each silo, with colors indicating the atlas used in the training silos. The left plot presents the results of working memory prediction using functional connectomes (FC) on the HCP dataset, while the right plot depicts age prediction using structural connectomes (SC). The horizontal axis lists the atlas names associated with the testing silos. In the experiments, all participants were divided into five partitions (silos) 100 times. Stars above the boxplots for Coef_avg signify significantly higher prediction performance than all other models ($p < 0.001$).

3 Results

3.1 Aggregating linear models across silos improves prediction

We compared the prediction performance of ridge regression models trained on a single silo and the averaged model across silos. As shown in Fig. 3, for ridge regression, averaging coefficients of models across silos (Coef_avg) showed a significant improvement ($p < 0.001$) in prediction performance relative to models trained at a single silo using functional connectivity (FC) data. However, this enhancement was less marked when structural connectivity (SC) data were employed. When the Shen and Craddock atlases were the targets, the Coef_avg method demonstrated only a slightly better performance ($0.001 < p < 0.005$) than the second-best model from a single silo.

3.2 Federated learning across silos improves generalizability

We compared the prediction performance of DNNs trained on a single silo and trained across silos in federated approaches. Federated learning models (FedAvg and FedProx) demonstrated superior performance to models trained on individual silos across all data modalities, predictive tasks, and target atlases (Fig. 4). DNNs trained on a single silo overfit due to the limited sample size and noise generated by atlas transformation. Yet, FLECHA significantly improve generalizability. No significant differences exist between the performance of FedAvg and FedProx.




Fig. 4: Prediction performance of FLECHA (FedAvg and FedProx) and models trained independently on each silo, with colors indicating the atlas used in the training silo. There was no significant difference between the performance of FedAvg and FedProx. Stars above the boxplots for FedAvg and FedProx signify significantly higher prediction performance than comparison models ($p < 0.001$)

4 Conclusions

In this paper, we present FLECHA, a method to locally transform connectomes to a common atlas space, facilitating collaborative training across silos with connectomes from different atlases. FLECHA improves predictive performance, reduces overfitting, and enhances model generalizability by leveraging larger, more diverse datasets. Despite potential noise from atlas transformations, FLECHA ensures cross-silo robustness, allowing the model to focus on core connectome signals and enhance generalization.

In summary, FLECHA offers a novel approach for federated learning with connectomes from heterogeneous atlas spaces, addressing data sharing and heterogeneity challenges in connectomics.

Future work will include more tasks, encompassing classification and other labels. Ablation studies with a broader range of atlases, varying in node number and coverage, will help identify atlases too coarse for FLECHA. We will explore machine unlearning to mitigate the impact of individual silos on the centralized model. Additionally, FLECHA will be tested in real-world scenarios where data heterogeneity results from different scanners and preprocessing steps.

Acknowledgments. 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.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Markiewicz, C.J., Gorgolewski, K.J., Feingold, F., Blair, R., Halchenko, Y.O., Miller, E., Hardcastle, N., Wexler, J., Esteban, O., Goncavles, M., Jwa, A., Poldrack, R.: The OpenNeuro resource for sharing of neuroscience data **10**, e71774. <https://doi.org/10.7554/eLife.71774>, <https://doi.org/10.7554/eLife.71774>, publisher: eLife Sciences Publications, Ltd
2. Poldrack, R.A., Gorgolewski, K.J.: Making big data open: data sharing in neuroimaging **17**(11), 1510–1517. <https://doi.org/10.1038/nn.3818>, <https://www.nature.com/articles/nn.3818>, number: 11 Publisher: Nature Publishing Group
3. Horien, C., Noble, S., Greene, A.S., Lee, K., Barron, D.S., Gao, S., O'Connor, D., Salehi, M., Dadashkarimi, J., Shen, X., Lake, E.M.R., Constable, R.T., Scheinost, D.: A hitchhiker's guide to working with large, open-source neuroimaging datasets **5**(2), 185–193. <https://doi.org/10.1038/s41562-020-01005-4>, <https://www.nature.com/articles/s41562-020-01005-4>, number: 2 Publisher: Nature Publishing Group
4. Arslan, S., Ktena, S.I., Makropoulos, A., Robinson, E.C., Rueckert, D., Parisot, S.: Human brain mapping: A systematic comparison of parcellation methods for the human cerebral cortex **170**, 5–30. <https://doi.org/10.1016/j.neuroimage.2017.04.014>, <https://www.sciencedirect.com/science/article/pii/S1053811917303026>
5. Dadi, K., Rahim, M., Abraham, A., Chyzyk, D., Milham, M., Thirion, B., Varoquaux, G.: Benchmarking functional connectome-based predictive models for resting-state fMRI **192**, 115–134. <https://doi.org/10.1016/j.neuroimage.2019.02.062>, <https://www.sciencedirect.com/science/article/pii/S1053811919301594>
6. Ye, M., Fang, X., Du, B., Yuen, P.C., Tao, D.: Heterogeneous federated learning: State-of-the-art and research challenges, <http://arxiv.org/abs/2307.10616>
7. Huang, C., Huang, J., Liu, X.: Cross-silo federated learning: Challenges and opportunities, <http://arxiv.org/abs/2206.12949>
8. Li, T., Sahu, A.K., Zaheer, M., Sanjabi, M., Talwalkar, A., Smith, V.: Federated optimization in heterogeneous networks. <https://doi.org/10.48550/arXiv.1812.06127>, <http://arxiv.org/abs/1812.06127>
9. Li, X., Gu, Y., Dvornek, N., Staib, L.H., Ventola, P., Duncan, J.S.: Multi-site fMRI analysis using privacy-preserving federated learning and domain adaptation: ABIDE results **65**, 101765. <https://doi.org/10.1016/j.media.2020.101765>, <https://www.sciencedirect.com/science/article/pii/S1361841520301298>
10. Gao, D., Yao, X., Yang, Q.: A survey on heterogeneous federated learning. <https://doi.org/10.48550/arXiv.2210.04505>, <http://arxiv.org/abs/2210.04505>
11. Yan, Y., Li, W., Wu, H., Min, H., Tan, M., Wu, Q.: Semi-supervised optimal transport for heterogeneous domain adaptation. In: Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence. pp. 2969–2975. International Joint Conferences on Artificial Intelligence Organization. <https://doi.org/10.24963/ijcai.2018/412>, <https://www.ijcai.org/proceedings/2018/412>
12. Shen, C., Guo, Y.: Unsupervised heterogeneous domain adaptation with sparse feature transformation. In: Proceedings of The 10th Asian Conference on Machine Learning. pp. 375–390. PMLR, <https://proceedings.mlr.press/v95/shen18b.html>, ISSN: 2640-3498

13. Dadashkarimi, J., Karbasi, A., Liang, Q., Rosenblatt, M., Noble, S., Foster, M., Rodriguez, R., Adkinson, B., Ye, J., Sun, H., Camp, C., Farruggia, M., Tejavibulya, L., Dai, W., Jiang, R., Pollatou, A., Scheinost, D.: Cross atlas remapping via optimal transport (CAROT): Creating connectomes for different atlases when raw data is not available **88**, 102864. <https://doi.org/10.1016/j.media.2023.102864>, <https://www.sciencedirect.com/science/article/pii/S136184152300124X>
14. Liang, Q., Dadashkarimi, J., Dai, W., Karbasi, A., Chang, J., Zhou, H.H., Scheinost, D.: Transforming connectomes to “any” parcellation via graph matching. In: Manfredi, L., Ahmadi, S.A., Bronstein, M., Kazi, A., Lomanto, D., Mathew, A., Magerand, L., Mullakaeve, K., Papiez, B., Taylor, R.H., Trucco, E. (eds.) *Imaging Systems for GI Endoscopy, and Graphs in Biomedical Image Analysis*. pp. 118–127. Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-21083-9_12
15. McMahan, H.B., Moore, E., Ramage, D., Hampson, S., Arcas, B.A.y.: Communication-efficient learning of deep networks from decentralized data. <https://doi.org/10.48550/arXiv.1602.05629>, <http://arxiv.org/abs/1602.05629>
16. Scheinost, D., Finn, E.S., Tokoglu, F., Shen, X., Papademetris, X., Hampson, M., Constable, R.T.: Sex differences in normal age trajectories of functional brain networks **36**(4), 1524–1535. <https://doi.org/10.1002/hbm.22720>, <https://onlinelibrary.wiley.com/doi/abs/10.1002/hbm.22720>, <https://onlinelibrary.wiley.com/doi/pdf/10.1002/hbm.22720>, [eprint:](#)
17. Van Essen, D.C., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T.E.J., Bucholz, R., Chang, A., Chen, L., Corbetta, M., Curtiss, S.W., Della Penna, S., Feinberg, D., Glasser, M.F., Harel, N., Heath, A.C., Larson-Prior, L., Marcus, D., Michalareas, G., Moeller, S., Oostenveld, R., Petersen, S.E., Prior, F., Schlaggar, B.L., Smith, S.M., Snyder, A.Z., Xu, J., Yacoub, E., WU-Minn HCP Consortium: The human connectome project: a data acquisition perspective **62**(4), 2222–2231. <https://doi.org/10.1016/j.neuroimage.2012.02.018>
18. Shen, X., Tokoglu, F., Papademetris, X., Constable, R.T.: Groupwise whole-brain parcellation from resting-state fMRI data for network node identification **82**, 403–415. <https://doi.org/10.1016/j.neuroimage.2013.05.081>, <https://www.sciencedirect.com/science/article/pii/S1053811913005818>
19. Craddock, R.C., James, G., Holtzheimer III, P.E., Hu, X.P., Mayberg, H.S.: A whole brain fMRI atlas generated via spatially constrained spectral clustering **33**(8), 1914–1928. <https://doi.org/10.1002/hbm.21333>, <https://onlinelibrary.wiley.com/doi/abs/10.1002/hbm.21333>, <https://onlinelibrary.wiley.com/doi/pdf/10.1002/hbm.21333>, [eprint:](#)
20. Schaefer, A., Kong, R., Gordon, E.M., Laumann, T.O., Zuo, X.N., Holmes, A.J., Eickhoff, S.B., Yeo, B.T.T.: Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI **28**(9), 3095–3114. <https://doi.org/10.1093/cercor/bhx179>, <https://doi.org/10.1093/cercor/bhx179>
21. Fan, L., Li, H., Zhuo, J., Zhang, Y., Wang, J., Chen, L., Yang, Z., Chu, C., Xie, S., Laird, A.R., Fox, P.T., Eickhoff, S.B., Yu, C., Jiang, T.: The human brainnetome atlas: A new brain atlas based on connectional architecture **26**(8), 3508–3526. <https://doi.org/10.1093/cercor/bhw157>
22. Dosenbach, N.U.F., Fair, D.A., Miezin, F.M., Cohen, A.L., Wenger, K.K., Dosenbach, R.A.T., Fox, M.D., Snyder, A.Z., Vincent, J.L., Raichle, M.E., Schlaggar, B.L., Petersen, S.E.: Distinct brain networks for adaptive and stable task control in humans **104**(26), 11073–11078. <https://doi.org/10.1073/pnas.0704320104>, <https://www.pnas.org/doi/10.1073/pnas.0704320104>, publisher: Proceedings of the National Academy of Sciences