

A Matrix Autoencoder Framework to Align the Functional and Structural Connectivity Manifolds as Guided by Behavioral Phenotypes

Niharika Shimona D’Souza¹ ^{*}, Mary Beth Nebel²³, Deana Crocetti², Joshua Robinson², Stewart Mostofsky²³⁴, and Archana Venkataraman¹

¹ Dept. of Electrical and Computer Eng., Johns Hopkins University, Baltimore, USA

^{*}Shimona.Niharika.Dsouza@jhu.edu

² Center for Neurodevelopmental & Imaging Research, Kennedy Krieger Institute

³ Dept. of Neurology, Johns Hopkins School of Medicine, Baltimore, USA

⁴ Dept. of Psychiatry & Behavioral Science, Johns Hopkins School of Medicine, USA

Abstract. We propose a novel matrix autoencoder to map functional connectomes from resting state fMRI (rs-fMRI) to structural connectomes from Diffusion Tensor Imaging (DTI), as guided by subject-level phenotypic measures. Our specialized autoencoder infers a low dimensional manifold embedding for the rs-fMRI correlation matrices that mimics a canonical outer-product decomposition. The embedding is simultaneously used to reconstruct DTI tractography matrices via a second manifold alignment decoder and to predict inter-subject phenotypic variability via an artificial neural network. We validate our framework on a dataset of 275 healthy individuals from the Human Connectome Project database and on a second clinical dataset consisting of 57 subjects with Autism Spectrum Disorder. We demonstrate that the model reliably recovers structural connectivity patterns across individuals, while robustly extracting predictive and interpretable brain biomarkers in a cross-validated setting. Finally, our framework outperforms several baselines at predicting behavioral phenotypes in both real-world datasets.

Keywords: Matrix Autoencoder, Manifold Alignment, Functional Connectivity, Structural Connectivity, Phenotypic Prediction

1 Introduction

The brain is increasingly viewed as an interconnected network. Two key elements of this network are the structural pathways between brain regions and the functional signaling that rides on top. This structural connectivity information can be measured via Diffusion Tensor Imaging (DTI) tractography [123]. Likewise, resting-state fMRI (rs-fMRI) captures inter-regional co-activation, which can be used to infer functional connectivity [2513]. Several studies have found both direct and indirect correspondences between structural and functional connectivity [1614]. Going a step further, structural and functional connectivity have been shown to be predictive of each other at varying scales [31839]. Hence,

multimodal integration of these viewpoints has become a key area of focus for characterizing neuropsychiatric disorders such as autism and ADHD [28][40].

In the clinical neuroscience realm, techniques for integrating structural and functional connectivity focus on group-wise discrimination. These works include statistical tests on edge-based features to identify significant differences in Alzheimer's disease [15], parallel ICA using structure and function to identify discriminative biomarkers of schizophrenia [36], and classical machine learning techniques to predict diagnosis [6]. While highly informative at a group level, these methods do not directly address inter-individual variability, for example by predicting finer grained patient characteristics. This divide has been partially bridged by end-to-end deep learning models. Examples include MLPs [26] for age prediction from functional connectomes and convolutional neural networks [20] for predicting cognitive and motor measures from structural connectomes. The work of [11] takes the alternate approach of combining a dictionary learning model on the functional connectomes coupled with an Artificial Neural Network (ANN) that predicts multiple clinical measures, while also preserving interpretability. Even so, these models focus exclusively on a single neuroimaging modality and do not exploit the interplay between function and structure.

Geometric learning frameworks have recently shown great promise in multimodal connectomics studies, both for conventional manifold learning [38] and in the context of Graph Convolutional Networks (GCN) [27][40]. Their primary advantage is the ability to directly incorporate and exploit the underlying data geometry. Beyond associative analyses, the work of [39][4] employ multi-GCNs combined with a Generative Adversarial Network (GAN) for the alignment problem. Particularly, [39] examines the problem of recovering structural connectomes from patient functional connectomes. While this paper marks a seminal contribution to multimodal integration, the representations learned by end-to-end GCNs can be hard to interpret. It can also be difficult to train GANs on modest-sized datasets [19].

We propose an end-to-end matrix autoencoder that maps rs-fMRI correlation matrices to structural connectomes obtained from DTI tractography. Inspired by recent work in Riemannian deep learning [17][9], our matrix autoencoder, estimates a low dimensional embedding from rs-fMRI correlation matrices while taking into account the geometry of the functional connectivity (FC) manifold. Our second matrix decoder uses this embedding to reconstruct patient structural connectivity (SC) matrices akin to a manifold alignment [37] between the FC and SC data spaces. For regularization, the FC embedding is also used to predict behavioral phenotypes. We demonstrate that our framework reliably traverses from function to structure and extracts meaningful brain biomarkers.

2 A Matrix Auto-Encoder for Connectome Manifolds

Fig. 1 illustrates our matrix autoencoder framework consisting of an encoder-decoder for functional connectivity (gray box), manifold alignment for estimating structural connectivity (blue box), and ANN for prediction of behavioral phe-

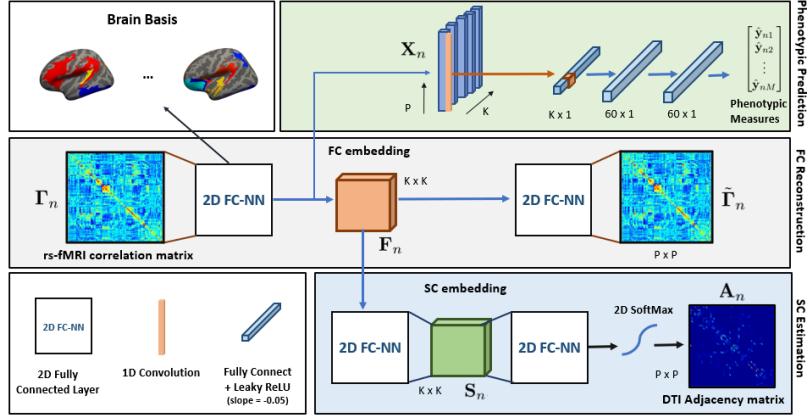


Fig. 1. A Matrix Autoencoder for aligning the FC and SC manifolds **Gray Box**: Matrix encoder-decoder for functional connectomes. **Blue Box**: Alignment Decoder for estimating DTI connectomes **Green Box**: ANN for predicting behavioral phenotypes

notypes (green box). Let N be the number of patients and P be the number of ROIs in our brain parcellation. We denote the rs-fMRI correlation matrix for patient n by $\Gamma_n \in \mathcal{R}^{P \times P}$. $\mathbf{A}_n \in \mathcal{R}^{P \times P}$ is the corresponding structural connectivity profile, and $\mathbf{y}_n \in \mathcal{R}^{M \times 1}$ is a vector of M concatenated phenotypic measures.

Functional Connectivity Reconstruction: By construction, the correlation matrices Γ_n belong to the manifold of symmetric positive semi-definite matrices \mathcal{P}_P^+ . Our matrix autoencoder estimates a latent functional embedding $\mathbf{F}_n \in \mathcal{P}_K^+$ using a 2D fully connected (2D FC-NN) layer [9][17]. Formally, this mapping $\Phi_{\text{ec}}(\cdot) : \mathcal{P}_P^+ \rightarrow \mathcal{P}_K^+$ is parametrized by weights $\mathbf{W} \in \mathcal{R}^{P \times K}$ and is computed as a cascade of two linear layers with tied weights: $\mathbf{F}_n = \Phi_{\text{ec}}(\Gamma_n) = \mathbf{W}^T \Gamma_n \mathbf{W}$. Our decoder is another 2D FC-NN that estimates $\tilde{\Gamma}_n$ from \mathbf{F}_n via a similar transformation $\Phi_{\text{dc}}(\cdot) : \mathcal{P}_K^+ \rightarrow \mathcal{P}_P^+$ that shares weights with the encoder. Mathematically, our FC reconstruction loss is represented as follows:

$$\mathcal{L}_{\text{FC}} = \frac{1}{N} \sum_n \|\Phi_{\text{dc}}(\Phi_{\text{ec}}(\Gamma_n)) - \Gamma_n\|_F^2 = \frac{1}{N} \sum_n \|\mathbf{W} \mathbf{W}^T \Gamma_n \mathbf{W} \mathbf{W}^T - \Gamma_n\|_F^2 \quad (1)$$

The second term of Eq. (1) encourages the columns of the brain basis \mathbf{W} to be orthonormal. Conceptually, this specialized loss helps us learn uncorrelated patterns that explain the rs-fMRI data well while acting as an implicit regularizer.

Structural Connectivity Estimation: The structural connectivity matrices \mathbf{A}_n are derived from DTI tractography and belong to the manifold of symmetric (non PSD) matrices \mathcal{S}_P . Our alignment decoder first generates an SC embedding $\mathbf{S}_n \in \mathcal{R}^{K \times K}$ from \mathbf{F}_n via a 2D FC-NN layer $\Phi_{\text{align}}(\cdot) : \mathcal{P}_K^+ \rightarrow \mathcal{P}_K^+$, followed by a second 2D FC-NN layer $\Phi_{\text{est}}(\cdot) : \mathcal{P}_K^+ \rightarrow \mathcal{P}_P^+$ which maps to the structural connectivity matrices. For stability our SC matrices do not have self-connections

and are normalized to $\|\mathbf{A}_n\|_1 = 1$. Accordingly, at the output layer, we apply a 2D softmax $\mathcal{SF}(\cdot)$, and then suppress the diagonal elements to generate the final output $\tilde{\mathbf{A}}_n \in \mathcal{R}^{P \times P}$. Our SC estimation objective is represented as follows:

$$\mathcal{L}_{\text{SC}} = \frac{1}{N} \sum_n \left\| \mathcal{SF} \left[\Phi_{\text{est}}(\Phi_{\text{align}}(\mathbf{F}_n)) \circ [\mathbf{1}\mathbf{1}^T - \mathcal{I}_P] \right] - \mathbf{A}_n \right\|_F^2 \quad (2)$$

where \circ is the element-wise Hadamard product. $\mathbf{1} \in \mathcal{R}^{P \times 1}$ is the vector of all ones, and \mathcal{I}_P is the identity matrix of dimension P . Conceptually, the loss in Eq. (2) is akin to manifold alignment [37] between the functional and structural embeddings based on a two sided Procrustes-like objective.

Phenotypic Prediction: We map the intermediate representation $\mathbf{X}_n = \Gamma_n \mathbf{W} \in \mathcal{R}^{P \times K}$ learned by the FC encoder to the phenotypes \mathbf{y}_n via a cascade of a 1D convolutional layer and an ANN. The convolutional layer $\mathcal{F}_{\text{conv}}(\cdot)$ collapses \mathbf{X}_n along its rows via a weighted sum to generate a K dimensional feature vector. This feature vector is input to a simple two layered ANN $\mathcal{G}(\cdot)$ to jointly estimate the elements in $\hat{\mathbf{y}}_n$. We use a Mean Squared Error (MSE) loss function:

$$\mathcal{L}_{\text{phen}} = \frac{1}{NM} \sum_n \|\hat{\mathbf{y}}_n - \mathbf{y}_n\|_F^2 = \frac{1}{NM} \sum_n \|\mathcal{G}(\mathcal{F}_{\text{conv}}(\mathbf{X}_n)) - \mathbf{y}_n\|_2^2 \quad (3)$$

This prediction task is a secondary regularizer that encourages our matrix autoencoder to learn representations predictive of inter-subject variability.

Implementation Details: We train our framework on a joint objective that combines Eqs. (1), (2) and (3) as follows:

$$\mathcal{L} = \mathcal{L}_{\text{FC}} + \gamma_1 \mathcal{L}_{\text{SC}} + \gamma_2 \mathcal{L}_{\text{phen}} \quad (4)$$

where γ_1 and γ_2 balance the tradeoff for the SC estimation and phenotypic prediction relative to the FC reconstruction objective. We employ an ADAM optimizer [22] with learning rate 0.005 and weight decay regularization [29] ($\delta = 0.0005$) run for a maximum of 400 epochs. Optimization parameters were fixed based on a validation set consisting of 30 additional patients from the HCP database. We use this strategy to set the the dimensionality of our autoencoder embedding at $K = 15$ and loss penalties to $\{\gamma_1, \gamma_2\} = 10^3, 3$. Finally, we utilize a spectral initialization for the encoder-decoder weights \mathbf{W} in Eq. (1) based on the top K eigenvectors of the average patient correlation matrix $\bar{\Gamma}_n$ for the training set. We use a similar initialization based on $\tilde{\mathbf{A}}_n$ for $\Phi_{\text{est}}(\cdot)$, and default initialization [24] for the remaining layers. Our model has a runtime of 10-12 minutes on an 8 core machine with 32GB RAM implemented in PyTorch (v1.5.1).

2.1 Baseline Methods for Phenotypic Prediction

Matrix AE without rs-fMRI decoder: We start with the architecture in Fig. 1 but omit the rs-fMRI decoder loss (\mathcal{L}_{FC}) in Eq. (4). This helps us evaluate the benefit of a tied encoder-decoder model for the rs-fMRI matrices.

Matrix AE without DTI decoder: We start with the architecture in Fig. 1 but remove the DTI decoder loss (\mathcal{L}_{SC}) in Eq. (4). This helps us evaluate the benefit of manifold alignment to constrain the functional embedding.

Decoupled Matrix AE + ANN: We start with the architecture in Fig. 1 but decouple the representation learning on the connectomics data from the prediction of phenotypic measures by training the models separately.

BrainNetCNN: This baseline integrates multimodal connectivity data via the BrainNetCNN [20]. We modify the original architecture, which is designed for a single modality, to have two branches, one for the rs-fMRI correlation matrices Γ_n , and another for the DTI connectomes \mathbf{A}_n . The ANN is modified to output M measures of clinical severity. We set the hyperparameters according to [20].

rs-fMRI Dictionary Learning + ANN: The framework in [11] uses rs-fMRI correlation matrices for the prediction of multiple clinical measures. The model combines a dictionary learning with a neural network predictor, with these two blocks optimized in an end-to-end fashion via a coupled optimization objective.

3 Experimental Evaluation and Results

HCP Dataset and Preprocessing: We download rs-fMRI and DTI scans for of 275 healthy individuals from the Human Connectome Project (HCP) S1200 database [12]. Rs-fMRI data was pre-processed according to the standard HCP pipeline [35], which accounts for motion and physiological confounds. We selected a 15-minute interval from the rs-fMRI scans to remain commensurate with standard rs-fMRI protocols. We used the Neurodata MR Graphs package [21] to pre-process the DTI scans and estimate fiber bundles via streamline tractography. Our phenotypic measure was the Cognitive Fluid Intelligence Score (CFIS) [105] adjusted for age, which is obtained via a battery of tests measuring cognitive reasoning (dynamic range: 70 – 150). We used the Automatic Anatomical Labelling (AAL) atlas with 116 ROIs as our parcellation. From here, we compute the Pearson correlation between the mean rs-fMRI time course in the two regions; we subtract the contribution of the top (roughly constant) eigenvector to obtain the input Γ_n . The input \mathbf{A}_n are obtained by dividing the number of tracts between ROIs by the total fiber tracts across all ROI pairs.

Predicting Behavioral Phenotypes: Table 1 (and Fig. 1 in Supplementary) compares the model against the baselines when predicting CFIS in a five-fold cross validated setting. Lower Median Absolute Error (MAE) and higher Normalized Mutual Information (NMI) and correlation coefficient (R) signify improved performance. Our framework outperforms the baselines during testing, though the model of [11] comes in a close second. This suggests that the Matrix Autoencoder faithfully models subject-specific variation even in unseen patients.

Functional to Structural Association: We evaluate three aspects of our functional to structural manifold alignment. First is our ability to recover structural

Table 1. CFIS prediction on the HCP dataset against the baselines using Median Absolute Error (MAE), Normalized Mutual Information (NMI) for training and testing, and Correlation Coefficient (R) for the test set. Best performance is highlighted in bold.

Method	MAE Train	MAE Test	NMI Train	NMI Test	R
No rs-fMRI dec.	6.31 ± 5.61	16.42 ± 12.41	0.85	0.61	0.07
No DTI dec.	6.30 ± 5.80	15.44 ± 13.00	0.86	0.61	0.11
Decoupled.	2.53 ± 2.41	14.90 ± 13.60	0.87	0.59	0.10
BrainNetCNN	6.80 ± 6.25	14.95 ± 12.74	0.88	0.59	0.12
Dict. Learn. + ANN	3.19 ± 2.19	15.26 ± 13.99	0.89	0.66	0.29
Our Framework	3.19 ± 2.47	14.08 ± 11.85	0.86	0.69	0.30

connectivity matrices during testing. Here, we compare two distance metrics: (1) F_{self} is the Frobenius norm between a test example \mathbf{A}_n and the model prediction for the same example $\hat{\mathbf{A}}_n$, and (2) F_{other} is $\hat{\mathbf{A}}_n$ and other SC matrices \mathbf{A}_m , ($m \neq n$). As shown in the left of Fig. 2(a), F_{self} is consistently smaller than F_{other} , with statistical significance determined using the Wilcoxon rank sum test. This indicates that individual differences in SC are preserved by our framework. In the same plot, we also benchmark the recovery performance of our framework against a baseline Matrix encoder-decoder (gray box in Fig. 1) with the SC matrices as *input and output*. We also compare against a linear regression between the vectorized upper diagonal FC features (input) and SC features (output) to help evaluate the benefit of our matrix decomposition. As seen, our function → structure decoding achieves similar performance as directly encoding/decoding the structural connectivity. At the same time, the linear regression baseline performs worse than both of these techniques. This suggests that the ability to directly leverage the low rank matrix structure is key to preserving individual differences during reconstruction.

Second, we use t-SNE to visualize the symmetric FC and SC embeddings, \mathbf{F}_n and \mathbf{S}_n , respectively. Fig. 2(b) displays the 2D t-SNE representation computed

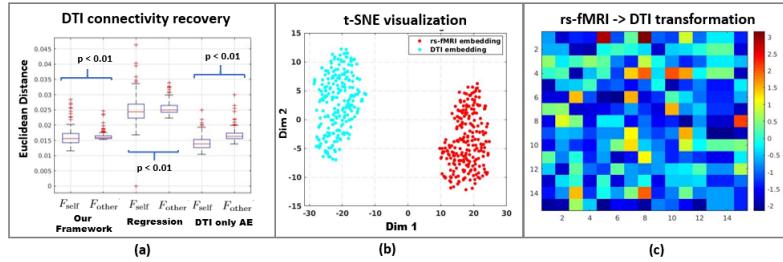


Fig. 2. (a) Recovery of SC for (L): Our Framework (M): Linear Regression (R): DTI only Autoencoder (b) t-SNE visualization for FC and SC embeddings (c) Coefficient of Variation (C_v) (log scale) for the weights of $\Phi_{\text{align}}(\cdot)$. Cold colors imply small deviations, i.e. better stability

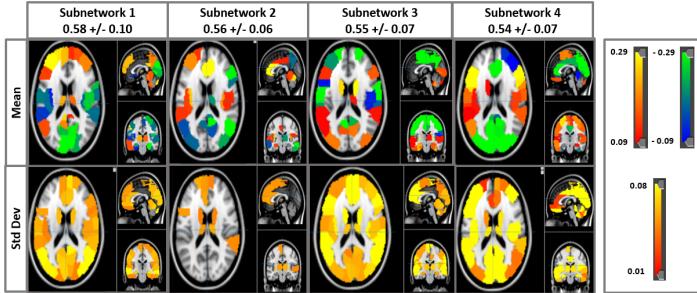


Fig. 3. Top four bases learned by the Matrix Autoencoder measured by the absolute correlation coefficient across cross validation folds and initializations.

from the upper-triangle entries of the embedding. As seen, the FC and SC are clustered in two different locations within this space. Interestingly, the learned representations are non-overlapping without explicit enforcement. This suggests that the alignment decoder $\Phi_{\text{align}}(\cdot)$ is learning a conversion between manifolds.

Third, we examine the stability of the transformation learned by the alignment decoder, i.e. the weights $\mathbf{W}_{\text{align}} \in \mathcal{R}^{K \times K}$ of $\Phi_{\text{align}}(\cdot)$. We first match the columns of $\mathbf{W}_{\text{align}}$ across cross validation folds according to correspondences between the functional brain basis. For each entry of $\mathbf{W}_{\text{align}}$, we compute the coefficient of variation (C_v), i.e. the ratio of the standard deviation to the mean (in absolute value). Lower values of C_v indicate smaller deviations from the mean values, i.e. better stability. Fig. 2(c) displays the log coefficient of variation $\log(C_v)$, where the cool colors indicate smaller C_v . As seen, a majority of the entries of $\mathbf{W}_{\text{align}}$ have low variation over the mean pattern value. Overall, our results suggest that our framework learns a stable mapping across the manifolds that explains individual patterns of structural connectivity faithfully.

Examining the FC Biomarkers: We explore the functional connectivity patterns learned by our framework by first matching the brain bases (i.e., columns of \mathbf{W}) across the cross validation folds based on the absolute correlation coefficient. We run this experiment five times with different initializations for the ANN branch to check for consistency in the learned representation. Fig 3 displays the four most consistent bases, as projected onto the brain using the region definitions of the AAL atlas. We notice that while there is spatial overlap between the bases, the standard deviations are small, which indicates that our framework is learning stable patterns in the data. Subnetwork 1 highlights regions from the default mode network, which is widely inferred within the resting state literature, and known to play a critical role in consolidating working memory [34]. Subnetworks 1, 3 and 4 highlight regions from the somatomotor network and visual cortex, together believed to be important functional biomarkers of cognitive intelligence [7]. Finally, Subnetwork 2 and 4 displays contributions from the frontoparietal network and the medial prefrontal network. These areas are

believed to play a role in working memory, attention, and decision making, all of which are associated with cognitive intelligence [30].

Application to a Secondary Dataset: We evaluate our framework on a second clinical cohort of 57 children with high-functioning Autism Spectrum Disorder (ASD). Rs-fMRI and DTI scans were acquired on a Philips 3T Achieva scanner (**rs-fMRI**: EPI, TR/TE = 2500/30ms, res = $3.05 \times 3.15 \times 3$ mm, duration = 128 or 156 time samples; **DTI**: EPI, TR/TE = 6356/75ms, res = $0.8 \times 0.8 \times 2.2$ mm, b-value = 700s/mm^2 , 32 gradient directions). Rs-fMRI pre-processing includes motion correction, normalization to MNI, spatial and temporal filtering, and nuisance regression with CompCorr [3]. DTI data was preprocessed using the FDT pipeline in FSL [18], with tractography performed using the BEDPOSTx and PROBTRACKx functions in FSL [2]. We define \mathbf{y}_n using three phenotypes that characterize various impairments associated with ASD: (1) the Autism Diagnostic Observation Schedule (ADOS) [33], (2) the Social Responsiveness Scale (SRS) [33], and (3) Praxis [32]. We carry forward the same parcellation scheme and model parameters as used for the HCP dataset. Table 2 compares the *multi-score* prediction testing performance of ADOS, SRS, and Praxis in a five fold cross validation setting. We observe that only our framework and the model of [11] can *simultaneously predict all three measures*. In contrast, the other baselines achieve good testing performance on one or two of the measures (for example, No DTI decoder baseline for ADOS and SRS) but cannot generalize all three. Our supplementary document includes additional results for SC recovery and FC biomarker extraction, and scatter plots for phenotypic prediction on the ASD datatset. Overall, our experiments on both healthy (HCP) and clinical (ASD) populations suggest that our model is robust across cohorts and generalizes effectively even with modest dataset sizes.

4 Conclusion

We have introduced a novel matrix autoencoder to map the manifold of rs-fMRI functional connectivity to the manifold of DTI structural connectivity. Our framework is strategically designed to leverage the underlying geometry of the data spaces and robustly recover brain biomarkers that are simultaneously explanatory of behavioral phenotypes. We demonstrate that our framework offers both interpretability and generalizability, even for multi-score prediction on modest sized datasets. Finally, our framework makes minimal assumptions, and can potentially find application both within and outside the medical realm.

Acknowledgements. This work was supported by the National Science Foundation CRCNS award 1822575, National Science Foundation CAREER award 1845430, National Institute of Mental Health (R01 MH085328-09, R01 MH078160-07, K01 MH109766 and R01 MH106564), National Institute of Neurological Disorders and Stroke (R01NS048527-08), and the Autism Speaks foundation.

Table 2. Multi-score performance on the ASD dataset using Median Absolute Error (MAE), Normalized Mutual Information (NMI), and Correlation Coefficient (R) for testing. Best performance is highlighted in bold. Near misses are underlined

Measure	Method	MAE Test	NMI Test	R
ADOS	No rs-fMRI dec.	3.11 ± 2.74	<u>0.46</u>	0.089
	No DTI dec.	2.61 ± 2.59	0.41	0.14
	Decoupled	2.64 ± 2.30	0.49	0.35
	BrainNetCNN	3.89 ± 2.80	0.35	0.05
	Dict. Learn.+ANN	<u>2.71 ± 2.40</u>	0.43	<u>0.50</u>
SRS	Our Framework	<u>2.71 ± 1.84</u>	0.49	0.51
	No rs-fMRI dec.	16.84 ± 16.01	0.77	0.039
	No DTI dec.	15.65 ± 12.69	0.81	0.24
	Decoupled	17.40 ± 14.16	0.74	0.02
	BrainNetCNN	17.50 ± 15.18	0.73	0.15
Praxis	Dict. Learn.+ANN	<u>16.79 ± 13.83</u>	0.89	0.37
	Our Framework	<u>16.04 ± 13.40</u>	<u>0.83</u>	<u>0.34</u>
	No rs-fMRI dec.	14.03 ± 10.80	0.74	0.02
	No DTI dec.	19.65 ± 13.18	0.81	0.23
	Decoupled	17.08 ± 12.23	0.76	0.09
Praxis	BrainNetCNN	19.35 ± 12.56	0.74	0.20
	Dict. Learn.+ANN	13.19 ± 10.75	<u>0.82</u>	0.37
	Our Framework	13.14 ± 10.78	0.86	<u>0.32</u>

References

1. Assaf, Y., Pasternak, O.: Diffusion tensor imaging (dti)-based white matter mapping in brain research: a review. *J. molecular neuroscience* **34**(1), 51–61 (2008)
2. Behrens, T.E., et al.: Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *NeuroImage* **34**(1), 144–155 (2007)
3. Behzadi, Y., et al.: A component based noise correction method (compcor) for bold and perfusion based fmri. *NeuroImage* **37**(1), 90–101 (2017)
4. Bessadok, A., Mahjoub, M.A., Rekik, I.: Topology-aware generative adversarial network for joint prediction of multiple brain graphs from a single brain graph. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 551–561. Springer (2020)
5. Bilker, W.B., et al.: Development of abbreviated nine-item forms of the raven’s standard progressive matrices test. *Assessment* **19**(3), 354–369 (2012)
6. Castellazzi, G., et al.: A machine learning approach for the differential diagnosis of alzheimer and vascular dementia fed by mri selected features. *Frontiers in neuroinformatics* **14**, 25 (2020)
7. Chén, O.Y., et al.: Resting-state brain information flow predicts cognitive flexibility in humans. *Scientific reports* **9**(1), 1–16 (2019)
8. Chu, S.H., et al.: Function-specific and enhanced brain structural connectivity mapping via joint modeling of diffusion and functional mri. *Sci. Rep.* **8**(1) (2018)
9. Dong, Z., et al.: Deep manifold learning of symmetric positive definite matrices with application to face recognition. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 31 (2017)

10. Duncan, J.: Frontal lobe function and general intelligence: why it matters. *Cortex: A Journal Devoted to the Study of the Nervous System and Behavior* (2005)
11. D'Souza, N.S., Nebel, M.B., Wymbs, N., Mostofsky, S., Venkataraman, A.: Integrating neural networks and dictionary learning for multidimensional clinical characterizations from functional connectomics data. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 709–717. Springer (2019)
12. Essen, V., et al.: The wu-minn human connectome project: an overview. *Neuroimage* **80**, 62–79 (2013)
13. Fox, M.D., Raichle, M.E.: Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nat. rev. neuro.* **8**(9), 700–711 (2007)
14. Fukushima, M., et al.: Structure–function relationships during segregated and integrated network states of human brain functional connectivity. *Brain Structure and Function* **223**(3), 1091–1106 (2018)
15. Hahn, K., et al.: Selectively and progressively disrupted structural connectivity of functional brain networks in alzheimer's disease—revealed by a novel framework to analyze edge distributions of networks detecting disruptions with strong statistical evidence. *Neuroimage* **81**, 96–109 (2013)
16. Honey, C.J., et al.: Predicting human resting-state functional connectivity from structural connectivity. *Proc. of the Nat. Acad. of Sci.* **106**(6), 2035–2040 (2009)
17. Huang, Z., Van Gool, L.: A riemannian network for spd matrix learning. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 31 (2017)
18. Jenkinson, M., et al.: Fsl. *NeuroImage* **62**(2), 782–790 (2012)
19. Karras, T., et al.: Training generative adversarial networks with limited data. arXiv preprint arXiv:2006.06676 (2020)
20. Kawahara, J., et al.: Brainnetcnn: Convolutional neural networks for brain networks; towards predicting neurodevelopment. *NeuroImage* **146**, 1038–1049 (2017)
21. Kiar, G., et al.: ndmg: Neurodata's mri graphs pipeline. Zenodo (2016)
22. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
23. Le Bihan, D., et al.: Diffusion tensor imaging: concepts and applications. *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine* **13**(4), 534–546 (2001)
24. LeCun, Y.A., Bottou, L., Orr, G.B., Müller, K.R.: Efficient backprop. In: *Neural networks: Tricks of the trade*, pp. 9–48. Springer (2012)
25. Lee, M.H., et al.: Resting-state fmri: a review of methods and clinical applications. *American Journal of neuroradiology* **34**(10), 1866–1872 (2013)
26. Lin, L., et al.: Predicting healthy older adult's brain age based on structural connectivity networks using artificial neural networks. *Computer methods and programs in biomedicine* **125**, 8–17 (2016)
27. Liu, J., et al.: Community-preserving graph convolutions for structural and functional joint embedding of brain networks. In: 2019 IEEE International Conference on Big Data (Big Data). pp. 1163–1168. IEEE (2019)
28. Liu, S., et al.: Multimodal neuroimaging computing: a review of the applications in neuropsychiatric disorders. *Brain informatics* **2**(3), 167–180 (2015)
29. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
30. Menon, V.: Large-scale brain networks and psychopathology: a unifying triple network model. *Trends in cognitive sciences* **15**(10), 483–506 (2011)

31. Messé, A., et al.: Predicting functional connectivity from structural connectivity via computational models using mri: an extensive comparison study. *NeuroImage* **111**, 65–75 (2015)
32. Mostofsky, S.H., et al.: Developmental dyspraxia is not limited to imitation in children with autism spectrum disorders. *Journal of the International Neuropsychological Society: JINS* **12**(3), 314 (2006)
33. Payakachat, N., et al.: Autism spectrum disorders: a review of measures for clinical, health services and cost–effectiveness applications. *Expert review of pharmacoeconomics & outcomes research* **12**(4), 485–503 (2012)
34. Sestieri, C., et al.: Episodic memory retrieval, parietal cortex, and the default mode network: functional and topographic analyses. *J. Neuroscience* **31**(12) (2011)
35. Smith, S.M., et al.: Resting-state fmri in the human connectome project. *Neuroimage* **80**, 144–168 (2013)
36. Sui, J., et al.: Combination of resting state fmri, dti, and smri data to discriminate schizophrenia by n-way mcca+jica. *Frontiers in human neuroscience* **7**, 235 (2013)
37. Wang, C., et al.: Manifold alignment using procrustes analysis. In: *Proceedings of the 25th international conference on Machine learning*. pp. 1120–1127 (2008)
38. Wong, E., Anderson, J.S., Zielinski, B.A., Fletcher, P.T.: Riemannian regression and classification models of brain networks applied to autism. In: *International Workshop on Connectomics in Neuroimaging*. pp. 78–87. Springer (2018)
39. Zhang, L., Wang, L., Zhu, D.: Recovering brain structural connectivity from functional connectivity via multi-gcn based generative adversarial network. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 53–61. Springer (2020)
40. Zhang, W., Zhan, L., Thompson, P., Wang, Y.: Deep representation learning for multimodal brain networks. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 613–624. Springer (2020)