JOINT SUBJECT-IDENTIFICATION AND TASK-DECODING FROM INFERRED FUNCTIONAL BRAIN GRAPHS VIA A MULTI-TASK NEURAL NETWORK

Elif Sema Balcioglu[1], *Berkay Doner*[1], *Ekansh Sareen*[1,2], *Dimitri Van De Ville*[1,2], *Hamid Behjat*[*1,2]

[1] Neuro-X Institute, École Polytechnique Fédérale de Lausanne, Geneva, Switzerland

[2] Department of Radiology and Medical Informatics, University of Geneva, Geneva, Switzerland

ABSTRACT

Functional connectivity (FC) between brain regions as manifested via fMRI entails signatures that can be used to differentiate individuals and decode cognitive tasks. In this work, we use methods from graph structure inference to estimate FC, which is in contrast to the conventional approach of deriving FC via correlation. Moreover, we infer FC graphs from seed-based co-activation patterns instead of raw fMRI data. We also propose a multi-task neural network architecture to jointly perform subject-identification and task-decoding from inferred functional brain graphs. We validate the developed model on data from the Human Connectome Project across eight fMRI tasks. Most importantly, our results show the superior task-decoding performance of FC graphs inferred from seed-based activity maps over graphs inferred from raw fMRI data. Furthermore, via gradient-based back-projection, we derive a significance score for inputs to the neural network, and present results showing the differential role of brain connections in subject-identification and task-decoding.

***Index Terms*—** Functional Connectivity, Graph Structure Inference, Neural Network, Subject-Identification, Task-Decoding, fMRI

1. INTRODUCTION

An increasing number of studies have demonstrated the effectiveness of functional connectivity (FC) as a fingerprint for identifying individuals from a population (subject-identification) [1] and the cognitive tasks they perform (task-decoding) [2]. Conventionally, FC is estimated via computing the correlation between time-courses of pairs of brain regions. Alternatively, functional connectivity estimation via learning-based graph structure inference [3, 4] has shown promising results on both fMRI [5] and EEG [6, 7] data. The key objective of this work is to validate the added predictability power of FC graphs inferred from seed-connectivity maps over FC graphs inferred from raw fMRI data as in prior work. The secondary objective of this work is to present a novel calssification architecture that jointly performs subject-identification and task-decoding using a single trained model. Subject-identification is conventionally performed via building an identifiability matrix via correlations between test and re-test FC graphs [8, 9, 10], whereas for task-decoding, using support vector machines [11, 5] is the more common approach. In this study, we leverage the power of neural networks and develop a multi-task neural network (MTNN) model, which optimizes for multiple losses simultaneously, offering advantages such as enhanced generalizability and reduced overfitting [12]. We assess the robustness and reliability of graphs inferred from seed-connectivity maps using the MTNN within the context of multiple experiments.

*Corresponding author, e-mail: hamid.behjat@epfl.ch

2. METHODS

2.1. Functional learned graphs

Let $\mathcal{G}$ denote an undirected, weighted graph with N vertices, and an edge set $\mathcal{E} \subset \{(i,j) \mid i \neq j,\ i,j \in \{1,\dots,N\}\}$. Let $\mathbf{A}$ denote the $N \times N$ adjacency matrix of $\mathcal{G}$, with elements $A_{i,j} \in \mathbb{R}^+$ denoting the weight of edge (i,j). Let $\mathbf{D}$ denote the degree matrix of $\mathcal{G}$, a diagonal matrix with elements $D_{i,i} = \sum_j A_{i,j}$. The combinatorial Laplacian matrix of $\mathcal{G}$ is given as $\mathbf{L} = \mathbf{D} - \mathbf{A}$. Let $\mathbf{f} \in \mathbb{R}^N$ denote a graph signal whose i-th component $\mathbf{f}[i]$ is the signal value at the i-th vertex of $\mathcal{G}$. The variability of $\mathbf{f}$ can be quantified using a measure of total variation (TV) as [13]:

$$\mathrm{TV}(\mathbf{f}) = \sum_{(i,j)\in\mathcal{E}} \mathbf{A}_{i,j}(\mathbf{f}[i] - \mathbf{f}[j])^2 = \mathbf{f}^T\mathbf{L}\mathbf{f}. \tag{1}$$

Given two signals $\mathbf{f}_1$ and $\mathbf{f}_2$, if $\mathrm{TV}(\mathbf{f}_1) < \mathrm{TV}(\mathbf{f}_2)$, then $\mathbf{f}_1$ is smoother than $\mathbf{f}_2$ as it exhibits less variability on $\mathcal{G}$. Using this notion of smoothness, a sparse graph structure can be inferred from a given set of observations [3]. Let $\mathbf{F}$ denote an $N \times M$ matrix where each column is an observation, a graph signal, and $\mathbf{Z}$ an $N \times N$ matrix with elements $Z_{i,j} = \|\mathbf{F}_{i,:} - \mathbf{F}_{j,:}\|_2$, i.e., Euclidean distance between signal values on vertices i and j. A graph structure can be inferred from $\mathbf{F}$ via the optimization [4]:

$$\begin{aligned} \min_{\mathbf{A}} \quad & \|\mathbf{A}\circ\mathbf{Z}\|_1 - \alpha\mathbf{1}^T\log(\mathbf{A}\mathbf{1}) + \frac{\beta}{2}\|\mathbf{A}\|_F^2, \\ \text{s.t.} \quad & \mathrm{diag}(\mathbf{A}) = 0, \\ & A_{i,j} = A_{j,i} \geq 0,\ i \neq j, \end{aligned} \tag{2}$$

where $\circ$ is the Hadamard product, α and β are regularization parameters, $\|\cdot\|_F$ the Frobenius norm, and $\mathbf{1} = [1,\dots,1]^T$. The first term in (2) enforces smoothness by invoking that [4]:

$$\|\mathbf{A}\circ\mathbf{Z}\|_1 = \frac{1}{2}\mathrm{trace}(\mathbf{F}^T\mathbf{L}\mathbf{F}) = \frac{1}{2}\sum_{k=1}^{M}\overbrace{\mathbf{f}_k^T\mathbf{L}\mathbf{f}_k}^{\mathrm{TV}(\mathbf{f}_k)}, \tag{3}$$

Smooth graph signals reside on strongly-connected vertices, thus, the vertices are expected to have smaller distances. The second term enforces degrees to be positive and improves overall connectivity.

2.2. Multi-task neural network

With a sample size limited to 100 individuals and a relatively high number of input features (specifically, upper triangular entries of the adjacency matrix, i.e., $M = P(P-1)/2$ features), our objective with using the MTNN is to efficiently share the weight-heavy components of the model, which maps the M input features to a lower

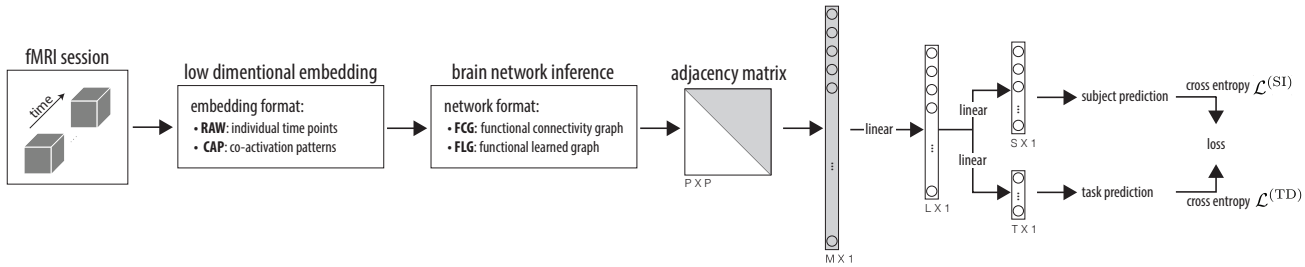

Fig. 1. Pipeline for joint subject-identification and task-decoding from fMRI.

dimensional representation of size $L \mid L << M$. Despite reducing the parameter size, as the initial representations are shared by two objectives, comparable performance can be achieved with other classification methods. The MTNN architecture produces two outputs, mapping the input features to both subject-identification and task-decoding labels. The input to the model is the upper triangular parts of the adjacency matrix of an FC graph. These flattened features are then mapped to an intermediate representation, which splits into two branches, one for subject-identification and one for task-decoding. In each step, we use fully-connected layers, followed by ReLU activation [14], and batch normalization [15] layers. We train the network using the cross-entropy losses of subject-identification (SI) and task-decoding (TD) problems. Specifically, given S subjects and T fMRI tasks, we train the model using a cross-entropy loss given as:

$$H\left(\mathbf{y}^{(\mathrm{SI})}, \mathbf{p}^{(\mathrm{SI})}, \mathbf{y}^{(\mathrm{TD})}, \mathbf{p}^{(\mathrm{TD})}\right) = \mathcal{L}^{(\mathrm{SI})} + \mathcal{L}^{(\mathrm{TD})}, \tag{4}$$

where $\mathcal{L}^{(\mathrm{SI})}$ and $\mathcal{L}^{(\mathrm{TD})}$ denote the cross entropy terms for SI and TD, respectively, defined as:

$$\mathcal{L}^{(\cdot)} = -\sum_{k=1}^{K} \mathbf{y}^{(\cdot)}[k] \log\left(\mathbf{p}^{(\cdot)}[k]\right), \tag{5}$$

where $\mathbf{y}^{(\cdot)}[k] \in \{0, 1\}$ is ground truth label for class k, $\mathbf{p}^{(\cdot)}[k] \in (0, 1)$ is the softmax probability of the output for class k, and K is the number of classes. In our experiments, for subject-identification we had $K = 100$ subjects and for task-decoding we had $K = 8$ tasks. Fig. 1 shows our proposed MTNN architecture.

2.3. Experimental dataset

We obtained structural MRI and fMRI data from the group of 100 unrelated subjects (54% female, mean age = 29.11$\pm$ 3.67, age range = 22-36) of the Human Connectome Project (HCP) dataset [16]. We used the HCP minimally preprocessed data [17], which ensures the necessary co-registration accuracy to apply this method. The fMRI data consisted of two repetitions of seven task-based sessions and a resting-state session per subject, thus, in total 16 sessions per subject. We used the minimally preprocessed volumetric fMRI data as provided by HCP. As for temporal preprocessing, for each voxel, we regressed out head movement, performed de-trending (0, 1st, and 2nd order), and, finally, z-scored across time. For each subject and each fMRI session, we transformed the high-resolution data (3D volumes) into two different low-dimensional embeddings. For both embeddings, we use the Schaefer brain atlas [18] with $P = 400$ parcels, each being a region of the cerebral cortex—with 200 parcels in each of the hemispheres; see Fig. 2(a).

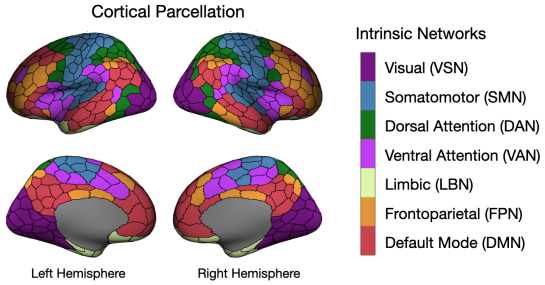

Fig. 2. The Schaefer400 brain atlas [18]; each region is assigned to one of the seven intrinsic brain networks [19].

For the first embedding—denoted in the following as the `RAW` data—each fMRI frame (each frame representing one time point) was parcellated using the atlas, and, subsequently voxels within each parcel were averaged, resulting in $P \times K$ matrices, where K denotes the length of the session. For the second embedding—denoted in the following as the `CAP` data—we extracted P static seed-based co-activation patterns, each associated with one atlas region. Each connectivity map was defined based on the method initially introduced in [20], and more recently used in [21]. Specifically, we extracted the average activity time course of each brain region, and then selected the top 15% of the frames (time-points) that showed maximal activity. These frames were then averaged, resulting in one 3D volume (a co-activation map) per brain region. Each co-activation map was then parcellated with the atlas, and voxels within each region were then averaged. By repeating this procedure across all atlas regions, each fMRI session was transformed to a $P \times P$ matrix, each column representing a seed-connectivity map.

2.4. Experimental setups

Through a series of experiments, we used the MTNN architecture to compare the performances of FC matrices generated using data formats `RAW` and `CAP`, and methods functional learned graph (`FLG`) and functional connectivity graph (`FCG`). As a first experiment, we validated the performance of MTNN compared to two algorithms: Support Vector Machines (SVM) and Random Forest (RF). For this, we found the best parameters (C, γ for SVM with "rbf" kernel, maximum depth, and number of estimators for RF) with grid-search, using a 3-way cross-validation. It is crucial to note that MTNN consists of a single model predicting both subject-identification and task-decoding however, separate models for those objectives are trained for SVM and RF. For all experiments on MTNN, we used the Adam optimizer [22] with a 5e-4 learning rate and a 5e-4 weight decay.

Given the two repetitions available for each subject and each fMRI session type, we used the first repetition data for training and performed testing on the second repetition data. Specifically, we trained our models on 800 first repetition fMRI sessions (100 subjects $\times$ 8 task types), and subsequently tested the trained models on 800 second repetition sessions. This separation of the data was used in all our experiments, that is, for each data format (RAW or CAP) and each graph inference method (FCG or FLG).

In all our experiments on FCG, we used density-matched FCG, wherein we set the smallest edge weights to zero up to the point where the density of the matrix (i.e., number of remaining graph edges) becomes equal to that of FLG of the corresponding session. This enabled us to have a fair comparison between the two methods, as our initial experiments (results not presented here) revealed that fully-connected FCG perform substantially worse than FLG compared to density-matched FCG.

3. RESULTS AND DISCUSSION

In the first experiment we compared the performance of the MTNN on the RAW data against two widely used classification methods: SVM and Random Forest, trained separately for both objectives; see Table 1. The MTNN outperformed both SVM and Random Forest in task-decoding (TD) and exhibits a performance comparable to SVM in subject-identification (SI). Importantly, MTNN's superior performance is obtained using a single trained model in contrast to the use of separate trained models for SI and TD in SVM and RF.

	SVM	RF	MTNN
Subject-Identification	0.998	0.874	0.997
Task-Decoding	0.863	0.436	0.883
Number of trained models	2	2	1

Table 1. Accuracy of MTNN vs alternative methods.

Fig. 3(a) shows subject-identification and task-decoding performances on FCG and FLG derived from the RAW data. The results are presented for varying numbers of frames, ranging from as few as 10 frames to 100 frames. With the exception of 10 frames, FLG outperforms FCG on SI, with accuracy reaching very close to 1.0 after 60 frames, whereas FCG remains at an accuracy of 95%. On TD, FCG catches up to the performance of FLG as we increase the number of frames, however, with only a few frames, FLG shows a substantial advantage, with around 60% accuracy even with as few frames as 10. Fig. 3(b) shows individual performances on eight different tasks on TD. The figure reports the accuracy difference between FLG and FCG.FLG shows superior performance over FCG across the majority of tasks and across different number of frames. In particular, FLG shows as much as 50% improved performance over FCG when a minimal number of frames are used.

We then assessed the performance of the CAP data. In addition to analyzing the data as a whole—i.e., using the entire $P \times P$ matrices as input for graph inference, we inferred graphs from seed-connectivity maps of regions that fall within each of the seven intrinsic brain networks—i.e., using only a subset of the columns from the data for each network, where the number of columns (regions) ranged from 26 (LBN) to 91 (DMN) depending on the network. Fig. 4(a) shows SI and TD performance differences of the FC derived from the CAP data using FLG and FCG. For each network, FLG consistently outperforms FCG in both SI and TD. We then aimed at assessing the SI and TD performance using FLG from the RAW and the

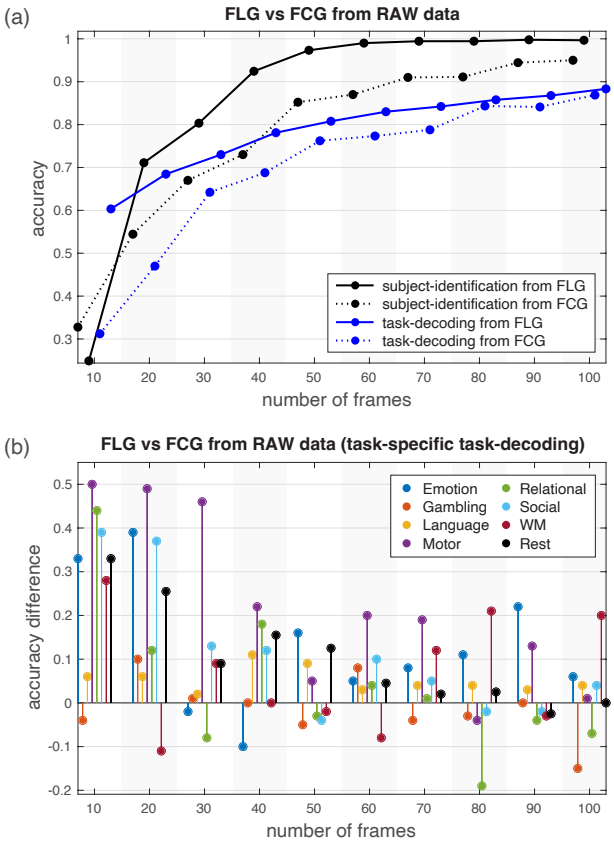

Fig. 3. (a) SI and TD prediction accuracy on data RAW using FCG and FLG graphs. (b) Average across tasks TD accuracies shown in (a) are split into eight different tasks, and the difference in performance between FLG and FCG is reported; values greater than zero indicate superior performance of FLG over FCG, and vice versa.

CAP data. Since both formats have unique characteristics, it is difficult to find a setting where they are treated equally. As a solution, for a network with R regions, we generated a new data format based on the RAW data with the same shape as CAP ($P \times R$), where each column was obtained by randomly selecting 15%—as used for generating the CAP data—of the columns from the RAW data and averaging them. We denote this version of the RAW data generated via random averaging of frames as RAW-RAND. Fig. 4(b) shows the accuracy difference FLG from CAP to FLG from RAW-RAND. FLG from the CAP data provide representations that enable superior TD across all brain networks. In contrast, for SI, FLG from the RAW-RAND data shows superior performance across all networks.

To disentangle the discrepancy between the result on TD and SI of Fig. 4(b), we aimed at assessing the "significance" of input features to the networks—i.e., brain graph edge weights—on the resulting predictions. To this aim, we used the DeepLIFT algorithm [23], which is a gradient-based interpretability method. The significance of each edge is calculated from the difference in neuron activation values from a reference input. A reference input in our case refers to a functional connectome matrix consisting of zeros. The algorithm gives separate consideration to positive and negative contributions to the prediction, where values close to 0 show a lack of significance. Using this method, for each trained MTNN, we constructed two up-

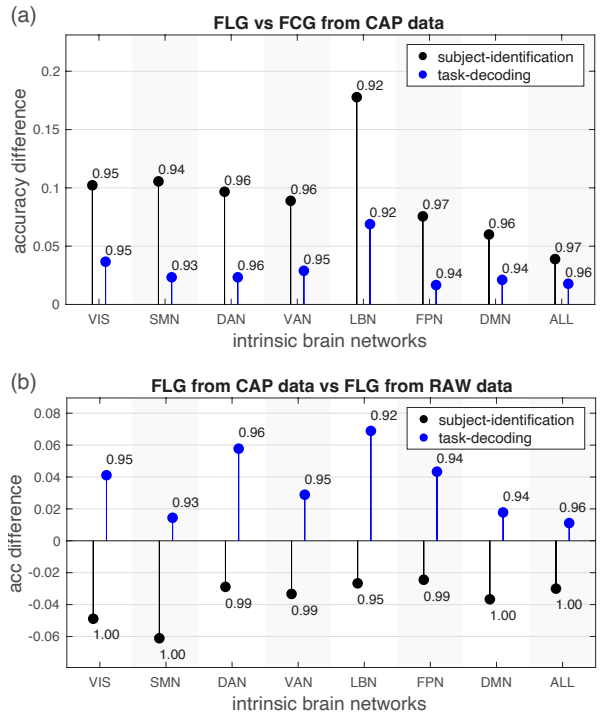

Fig. 4. (a) SI and TD prediction accuracy on the CAP data using FCG and FLG for the seven intrinsic brain networks shown in Fig. 2. Values above zero indicate better performance using FLG, and vice versa. (b) SI and TD prediction accuracy using FLG graphs inferred from CAP data compared to RAW data. Values above zero indicate better performance on CAP data whereas value below zero indicate better performance on RAW data. In both (a) and (b), values at the tip of stem bars show the accuracy of the superior method.

per triangular $P \times P$ matrices, one for SI and one for TD, denoted $\mathbf{S}^{(TD)}$ and $\mathbf{S}^{(SI)}$, respectively. Each upper triangular element (i, j) of these matrices, denoted $S^{(\cdot)}_{i,j}$, represents the significance of the associated graph edge—i.e., the connection between region i and region j—-in the success of the trained model for the task, i.e., TD or SI. Each element $S^{(\cdot)}_{i,j}$ was obtained by averaging T or S values for TD and SI, respectively, each of which is the absolute value of per-class significance of edge (i, j). The averaging enables identifying edges that have the greatest significance across all classes. The values were then taken to the logarithm scale.

We constructed $\mathbf{S}^{(TD)}$ and $\mathbf{S}^{(SI)}$ for the MTNN trained on FLG from the CAP, denoted $\mathbf{S}^{(\cdot)}_1$, and also for that trained on FLG from the RAW-RAND data, denoted $\mathbf{S}^{(\cdot)}_2$. Fig. 5 shows the results. MTNN trained on FLG from the CAP data prioritize inter-network connections more than RAW for TD, shown in the left part of Fig. 5; note the large values found in the off-diagonal blocks relative to the block diagonal values. On the contrary, both networks give similar importance to intra-network connections, which can be seen in values close to zero on block diagonal segments. This observation on TD is similar for SI, as can be inferred from the right-hand plot, which shows the difference between the two prediction tasks (TD minus SI); TD on FLG from the CAP data in general shows lower edge significance than SI, as reflected by the negative values, despite differences being very small (see colorbar range). This similarity is not unexpected, given the shared representation of both objectives in the MTNN network. In contrast to CAP, RAW-RAND places greater weight on inter-network connections between the LBN and other networks (see dark blue regions), which may be a factor to explain the superior SI performance of FLG from the RAW-RAND.

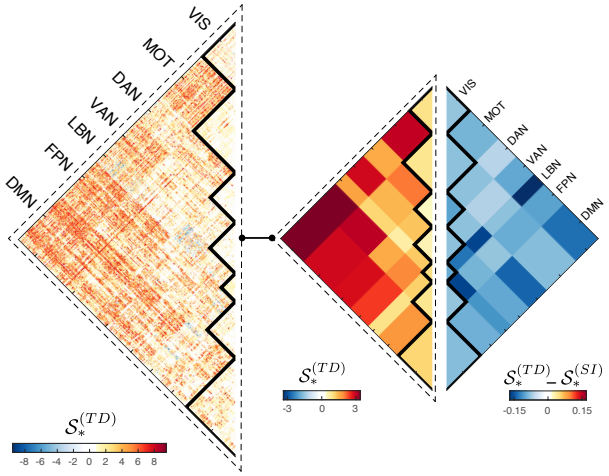

Fig. 5. Difference in the significance of brain connections in task-decoding and subject-identification for MTNN trained on FLG from the CAP data and MTNN trained on FLG from the RAW-RAND data. Specifically, here we are assessing the trained models associated to the results reported in the last column ("ALL") on Fig. 4(b). $\mathcal{S}_*^{(TD)} = \mathcal{S}_1^{(TD)} - \mathcal{S}_2^{(TD)}$; $\mathcal{S}_*^{(SI)} = \mathcal{S}_1^{(SI)} - \mathcal{S}_2^{(SI)}$; formal description of $\mathcal{S}_k^{(\cdot)}$ for $k = 1, 2$ is given in the main text. For the two half-matrices marked with dashed lines and shown as connected, the one on the right is obtained via computing the mean of the intra- and inter-network blocks of the one on the left, to ease visual comparison. The half matrix on the far right hand side was computed via the same method of block averaging.

4. CONCLUSIONS AND OUTLOOK

In this work we proposed a method for joint task-decoding and subject-identification from fMRI data. The proposed method initially embeds 4D fMRI data into a low dimensional space using a cortical atlas, specifically, constructing co-activation patterns linked to individual cortical regions [21]. A functional brain graph is then inferred from the low dimensional embedding [6], and, subsequently fed into the MTNN. The MTNN is trained to jointly predict the identity of the subject and the task being performed. Through a series of experiments, we show the superiority of the proposed method over conventional strategies of performing classification based on correlation-based FC. Our promising results suggest the potential of further expanding the proposed MTNN architecture, on different fronts: i) to take as input higher resolution FC graphs [24] or fMRI spectral features [11] extracted on macro-scale [25] or voxel-wise [26, 27] structural brain graphs; ii) to use graph neural networks [28] or methods of transfer learning to learn individualised task-decoding models [29]; iii) integrating the brain network inference step into the neural network. The present study only briefly touched upon disentangling the significance of different intra/inter network connections in task performance, an avenue which we believe is fertile to be explored in future work.

5. REFERENCES

[1] Emily S Finn, Xilin Shen, Dustin Scheinost, Monica D Rosenberg, Jessica Huang, Marvin M Chun, Xenophon Papademetris, and R Todd Constable, "Functional connectome fingerprinting: identifying individuals using patterns of brain connectivity," *Nature neuroscience*, vol. 18, no. 11, pp. 1664–1671, 2015.

[2] Shima Rastegarnia, Marie St-Laurent, Elizabeth DuPre, Basile Pinsard, and Pierre Bellec, "Brain decoding of the Human Connectome Project tasks in a dense individual fMRI dataset," *NeuroImage*, vol. 283, pp. 120395, 2023.

[3] Xiaowen Dong, Dorina Thanou, Pascal Frossard, and Pierre Vandergheynst, "Learning Laplacian matrix in smooth graph signal representations," *IEEE Trans. Signal Process.*, vol. 64, no. 23, pp. 6160–6173, 2016.

[4] Vassilis Kalofolias, "How to learn a graph from smooth signals," in *Artificial Intelligence and Statistics*. PMLR, 2016, pp. 920–929.

[5] Siyuan Gao, Xinyue Xia, Dustin Scheinost, and Gal Mishne, "Smooth graph learning for functional connectivity estimation," *Neuroimage*, vol. 239, pp. 118289, 2021.

[6] Maliheh Miri, Vahid Abootalebi, Hamid Saeedi-Sourck, Dimitri Van De Ville, and Hamid Behjat, "Spectral representation of eeg data using learned graphs with application to motor imagery decoding," *Biomedical Signal Processing and Control*, vol. 87, pp. 105537, 2024.

[7] Yassine El Ouahidi, Lucas Drumetz, Giulia Lioi, Nicolas Farrugia, Bastien Pasdeloup, and Vincent Gripon, "Spatial graph signal interpolation with an application for merging bci datasets with various dimensionalities," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.

[8] Enrico Amico and Joaquín Goñi, "The quest for identifiability in human functional connectomes," *Scientific Reports*, vol. 8, no. 1, pp. 1–14, 2018.

[9] Dimitri Van De Ville, Younes Farouj, Maria Giulia Preti, Raphaël Liégeois, and Enrico Amico, "When makes you unique: temporality of the human brain fingerprint," *Science advances*, vol. 7, no. 42, pp. eabj0751, 2021.

[10] Maliheh Miri, Vahid Abootalebi, Enrico Amico, Hamid Saeedi-Sourck, Dimitri Van De Ville, and Hamid Behjat, "Brain fingerprinting using EEG graph inference," *bioRxiv*, 2023.

[11] Alessandra Griffa, Enrico Amico, Raphaël Liégeois, Dimitri Van De Ville, and Maria Giulia Preti, "Brain structure-function coupling provides signatures for task decoding and individual fingerprinting," *Neuroimage*, vol. 250, pp. 118970, 2022.

[12] Sebastian Ruder, "An overview of multi-task learning in deep neural networks," *arXiv preprint arXiv:1706.05098*, 2017.

[13] Ulrike Von Luxburg, "A tutorial on spectral clustering," *Stat. Comput.*, vol. 17, pp. 395–416, 2007.

[14] Abien Fred Agarap, "Deep learning using rectified linear units (relu)," *arXiv preprint arXiv:1803.08375*, 2018.

[15] Sergey Ioffe and Christian Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *International conference on machine learning*. pmlr, 2015, pp. 448–456.

[16] D.C. Van Essen, S.M. Smith, D.M. Barch, T.E. Behrens, E. Yacoub, K. Ugurbil, and WU-Minn HCP Consortium., "The WU-Minn human connectome project: an overview.," *Neuroimage*, vol. 80, pp. 62–79, 2013.

[17] Matthew F Glasser, Stamatios N Sotiropoulos, J Anthony Wilson, Timothy S Coalson, Bruce Fischl, Jesper L Andersson, Junqian Xu, Saad Jbabdi, Matthew Webster, Jonathan R Polimeni, et al., "The minimal preprocessing pipelines for the human connectome project," *Neuroimage*, vol. 80, pp. 105–124, 2013.

[18] Alexander Schaefer, Ru Kong, Evan M Gordon, Timothy O Laumann, Xi-Nian Zuo, Avram J Holmes, Simon B Eickhoff, and BT Thomas Yeo, "Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI," *Cerebral cortex*, vol. 28, no. 9, pp. 3095–3114, 2018.

[19] BT Thomas Yeo, Fenna M Krienen, Jorge Sepulcre, Mert R Sabuncu, Danial Lashkari, Marisa Hollinshead, Joshua L Roffman, Jordan W Smoller, Lilla Zöllei, Jonathan R Polimeni, et al., "The organization of the human cerebral cortex estimated by intrinsic functional connectivity," *Journal of neurophysiology*, 2011.

[20] Xiao Liu and Jeff H Duyn, "Time-varying functional network information extracted from brief instances of spontaneous brain activity," *Proceedings of the National Academy of Sciences*, vol. 110, no. 11, pp. 4392–4397, 2013.

[21] Carlo Ferritto, Maria Giulia Preti, Stefano Moia, Dimitri Van De Ville, and Hamid Behjat, "Brain fingerprinting using fMRI spectral signatures on high-resolution cortical graphs," *bioRxiv*, 2023.

[22] Diederik P Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.

[23] Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje, "Learning important features through propagating activation differences," in *International conference on machine learning*. PMLR, 2017, pp. 3145–3153.

[24] Ye Tian, BT Thomas Yeo, Vanessa Cropley, Andrew Zalesky, et al., "High-resolution connectomic fingerprints: Mapping neural identity and behavior," *NeuroImage*, vol. 229, pp. 117695, 2021.

[25] Maria Giulia Preti and Dimitri Van De Ville, "Decoupling of brain function from structure reveals regional behavioral specialization in humans," *Nat. Commun.*, vol. 10, no. 1, pp. 1–7, 2019.

[26] David Abramian, Martin Larsson, Anders Eklund, Iman Aganj, Carl-Fredrik Westin, and Hamid Behjat, "Diffusion-informed spatial smoothing of fMRI data in white matter using spectral graph filters," *NeuroImage*, vol. 237, pp. 118095, 2021.

[27] Hamid Behjat, Anjali Tarun, David Abramian, Martin Larsson, and Dimitri Van De Ville, "Voxel-wise brain graphs from diffusion MRI: Intrinsic eigenspace dimensionality and application to functional MRI," *IEEE Open Journal of Engineering in Medicine and Biology*, 2023.

[28] Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst, "Geometric deep learning: going beyond euclidean data," *IEEE Signal Processing Magazine*, vol. 34, no. 4, pp. 18–42, 2017.

[29] Bruna Junqueira, Bruno Aristimunha, Sylvain Chevallier, and Raphael Y de Camargo, "A systematic evaluation of euclidean alignment with deep learning for eeg decoding," *arXiv preprint arXiv:2401.10746*, 2024.