

A Deep Graph Neural Network Architecture for Modelling Spatio-temporal Dynamics in resting-state functional MRI Data

Tiago Azevedo^{a,*}, Alexander Campbell^a, Rafael Romero-Garcia^b, Luca Passamonti^{c,d}, Richard A.I. Bethlehem^{e,b}, Pietro Liò^a, Nicola Toschi^{f,g}

^a*Department of Computer Science, University of Cambridge, UK*

^b*Brain Mapping Unit, Department of Psychiatry, University of Cambridge, UK*

^c*Department of Clinical Neurosciences, University of Cambridge, UK*

^d*Istituto di Bioimmagini e Fisiologia Molecolare (IBFM), Consiglio Nazionale delle Ricerche (CNR), Milano, Segrate, Italy*

^e*Autism Research Centre, Department of Psychiatry, University of Cambridge, UK*

^f*Department of Biomedicine and Prevention, University of Rome “Tor Vergata”, Italy*

^g*A. A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital and Harvard Medical School, US*

Abstract

Resting-state functional magnetic resonance imaging (rs-fMRI) has been successfully employed to understand the organisation of the human brain. For rs-fMRI analysis, the brain is typically parcellated into regions of interest (ROIs) and modelled as a graph where each ROI is a node and pairwise correlation between ROI blood-oxygen-level-dependent (BOLD) time series are edges. Recently, graph neural networks (GNNs) have seen a surge in popularity due to their successes in modelling unstructured relational data. The latest developments with GNNs, however, have not yet been fully exploited for the analysis of rs-fMRI data, particularly with regards to its spatio-temporal dynamics. Herein we present a novel deep neural network architecture, combining both GNNs and temporal convolutional networks (TCNs), which is able to learn from the spatial and temporal components of rs-fMRI data in an end-to-end fashion. In particular, this corresponds to intra-feature learning (i.e., learning temporal dynamics with TCNs) as well as inter-feature

*Corresponding author

Email addresses: tiago.azevedo@cst.cam.ac.uk (Tiago Azevedo), pietro.lio@cst.cam.ac.uk (Pietro Liò), toschi@med.uniroma2.it (Nicola Toschi)

learning (i.e., leveraging spatial interactions between ROIs with GNNs). We evaluate our model with an ablation study using 35,159 samples from the UK Biobank rs-fMRI database. We also demonstrate explainability features of our architecture which map to realistic neurobiological insights. We hope our model could lay the groundwork for future deep learning architectures focused on leveraging the inherently and inextricably spatio-temporal nature of rs-fMRI data.

Keywords: deep learning, graph neural networks, UK Biobank, time series, temporal convolutional network, rs-fMRI, spatio-temporal dynamics

¹ 1. Introduction

² Resting-state functional magnetic resonance imaging (rs-fMRI) is one of
³ the most commonly used noninvasive imaging techniques employed to gain
⁴ insight into human brain function. The use of rs-fMRI data has proven ex-
⁵ tremely useful as an investigative tool in neuroscience and, to some extent,
⁶ as a biomarker of brain disease diagnosis and progression [1]. Typical use of
⁷ rs-fMRI data involves using graph-theoretical measures (such as centrality
⁸ measures and community structure) to summarise high-dimensional, whole-
⁹ brain data for use in downstream tasks. As part of this process, it is common
¹⁰ practice to reduce the dimensionality of the data in one of three main ways:
¹¹ (1) by collapsing the temporal dimension (e.g., into brain region connectivity
¹² matrices based on similarity metrics between time series), (2) by reducing
¹³ the spatial dimension (e.g., in global signal regression for physiological noise
¹⁴ modelling [2]), and (3) by employing approaches that collapse both the tem-
¹⁵ poral and spatial dimensions (e.g., in independent component analyses [3]).
¹⁶ This feature engineering step is performed mostly due to the considerable vol-
¹⁷ ume of data in a typical rs-fMRI dataset and its relatively low signal-to-noise
¹⁸ ratio [4]. Although computationally beneficial, such dimensionality reduction
¹⁹ steps inevitably involve disregarding large amounts of information which can
²⁰ potentially be useful depending on the analysis task. For instance, collaps-
²¹ ing the temporal dimension of rs-fMRI data reduces the brain to a static
²² volume where the interactions between different brain regions are fixed over
²³ time. This stands in contrast to a growing body of research showing that the
²⁴ functional connectivity of the brain is dynamic and constantly changing over
²⁵ time [5, 6]. As another example, association measures most commonly used
²⁶ are still based on linear models, while it is well known that neuromonitoring

27 data and brain signal in particular interact nonlinearly [7, 8].

28 To overcome such limitations, a different approach to the analysis of rs-
29 fMRI data would be to devise a model that is able to combine both feature
30 engineering and the learning of a low-dimensional representation of the brain.
31 In order to do this, such a model would need to be able to accommodate both
32 the spatial as well as the temporal complexities of rs-fMRI data. To date,
33 deep learning architectures have had great success at leveraging specific in-
34 ductive biases from complex high-dimensional data. Convolutional neural
35 networks (CNNs), for instance, are extremely effective at extracting shared
36 spatial features such as corners and edges from grid-like data (e.g., 2D and
37 3D images). These features can then be combined into more complex con-
38 cepts deeper within the network architecture [9]. Recurrent neural networks
39 (RNNs), on the other hand, are able to learn features from data that are
40 temporally organised as a sequence of steps [10]. In contrast to both CNNs
41 and RNNs, graph neural networks (GNNs) can learn from data that does not
42 have a rigid structure like a grid or a sequence, and can be depicted in the
43 form of unordered entities and relations such as graphs. The formulation of
44 GNN models that deal with complex data structures has seen fast develop-
45 ments in the past years [11, 12] and are therefore strong candidates for the
46 analysis of rs-fMRI data.

47 Previous work has attempted to leverage deep learning architectures in
48 order to model rs-fMRI data. In particular, GNNs have been used to classify
49 binary sex [13], and CNNs have been successfully employed in the diagnosis
50 of cognitive impairment [14]. Other pioneering studies have devised *ad-hoc*
51 deep learning models for fMRI data such as the classification of brain dis-
52 orders using Siamese-inspired neural networks [15], but the spatial informa-
53 tion was not represented using GNNs. Learning from spatial and temporal
54 components of data using deep learning can also be seen in various non-
55 biological domains [16, 17, 18]; such approaches, however, usually rely on a
56 single spatio-temporal convolutional block that creates low-dimensional em-
57 beddings at each timestep, instead of performing a prediction task for the
58 entire graph at once. To improve these drawbacks, we formulate a novel deep
59 neural network architecture that exploits the advantages of GNNs and CNNs
60 in order to effectively model the linear and non-linear temporal and spatial
61 components of rs-fMRI data. We engineered our architecture to specifically
62 retain edge weights and contain elements of explainability [19, 20], in or-
63 der to provide advantages when a neuroscientific explanation of the inner
64 model workings is desirable. Our proposed model uses GNNs to account

for spatial inter-relationships between brain regions, and temporal convolutional networks (TCNs) to capture the intra-temporal dynamics of blood-oxygenated-level dependent (BOLD) time series. By incorporating GNNs and CNNs in the same end-to-end architecture we essentially combine intra- and inter-feature learning. In particular, GNNs can tackle a limitation of some graph representations of rs-fMRI data, in which association measures between different regions of interest (ROIs) of the brain are based on linear models; instead, GNNs can capture higher-order interactions between ROIs. A very preliminary version of this work with a 30-fold smaller dataset was recently presented as a conference contribution [21]. However, further work was needed regarding a larger dataset, wider choices of the graph threshold hyperparameter, and analysis on inclusion of edge weights.

We test our architecture on the publicly available UK Biobank dataset, which at the time of writing provides rs-fMRI scans from more than 30,000 distinct people. This dataset offers a unique opportunity to formulate novel architectures, while supporting the need of large datasets for reproducible findings with minimal statistical errors [22]. We also conducted an ablation analysis on a proof-of-concept binary sex prediction task to better evaluate the different contributions of each component of our model. We release all the code and artifacts used to develop this work in a public repository for easier adoption by the community (see “Data and Code Availability” section).

2. Methods

2.1. Problem Definition

To represent rs-fMRI data as an undirected weighted graph, the brain is spatially parcellated into N regions of interest (ROIs) representing graph nodes indexed by the set $\mathcal{V} = \{1, \dots, N\}$. Let $\mathbf{x}_i \in \mathbb{R}^T$ represent the features of node i corresponding to the BOLD timeseries of length T . The connections between each ROI are represented by an edge set $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ composed of $|\mathcal{E}| = E$ unordered pairs (i, j) , where for every edge k connecting two nodes $(i, j) \in \mathcal{E}$ the connection strength is defined as $\mathbf{e}_k \in \mathbb{R}$. Let the tuple $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ denote the resulting graph. Given the graph structure \mathcal{G} , let $\mathbf{X} \in \mathbb{R}^{N \times T}$, $\mathbf{E} \in \mathbb{R}^{E \times 1}$, and $\mathbf{A} \in \mathbb{R}^{N \times N}$ denote the nodes features, edge features and adjacency matrix, respectively.

98 *2.2. Temporal Convolutional Networks*

99 In order to learn a representation of the temporal dynamics contained
100 in rs-fMRI time series, we use temporal convolutional networks (TCNs) [23].
101 These are a simplification over the original *WaveNet* architecture used for au-
102 dio synthesis [24], which has been seen to provide significantly better results
103 for sequence modelling in comparison to more traditional RNN architectures
104 (e.g., LSTMs) across a range of tasks and datasets. In particular, Bai et al.
105 [23] posit that convolutional networks should be seen as the natural starting
106 point for sequence modelling tasks, which makes them ideal for extracting
107 information from rs-fMRI time series.

TCNs are based on dilated causal convolutions [25], which are special 1D filters where the size of the receptive field exponentially increases over the temporal dimension of the data as the depth of the network increases. The padding of the convolution is ‘causal’ in the sense that an output at a specific time step is convolved only with elements from earlier time steps from the previous layers, thus preserving temporal order. More formally, given a single ROI timeseries $\mathbf{x}_i \in \mathbb{R}^T$ and a filter $\mathbf{f} \in \mathbb{R}^K$, the dilated causal convolution operation of \mathbf{x} with \mathbf{f} at time t is represented as

$$\mathbf{x}_i * \mathbf{f}(t) = \sum_{s=0}^{K-1} \mathbf{f}(s) \mathbf{x}_i(t - d \times s) \quad (1)$$

108 where d is the dilation factor which controls the number of time steps suc-
109 cessively skipped. In contrast to the original TCN architecture [23], we use
110 batch normalisation instead of weight normalisation because it empirically
111 provided a more stable training procedure in terms of loss evolution.

112 *2.3. Graph Network Block*

113 Battaglia et al. [26] formalise a graph network (GN) framework through
114 the definition of functions that work on graph-structured representations.
115 The main unit of computation in the GN framework is called the *GN block*,
116 which contains three update functions and three aggregation functions work-
117 ing on the edge, node, and global level.

118 The first operation of this GN block, which can be broadly defined as the
119 *edge model*, concerns the update function ϕ^e , which computes updated edge
120 attributes for each edge k based on the original edge’s attributes \mathbf{e}_k and the
121 features of the connected nodes i and j :

148 However, assuming that distinct nodes (i.e., brain regions in this work)
 149 have different levels of importance for the downstream prediction task [27,
 150 28], we assumed that a hierarchical (as opposed to flat) pooling mechanism
 151 would create richer embeddings. To this end, we employ the differentiable
 152 pooling operator introduced by Ying et al. [29], commonly called DiffPool,
 153 which learns how to sequentially collapse nodes in smaller clusters until only
 154 a single node exists with the final embedding.

155 When describing a Graph Network (GN) block, a sparse representation
 156 of nodes and edges is used to describe the operations that a GN block can
 157 have; however, DiffPool works on dense representations of a graph. In other
 158 words, a graph \mathcal{G} is represented by a dense adjacency matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ and
 159 a feature matrix $\mathbf{X} \in \mathbb{R}^{N \times F}$, where N is the number of nodes and F the
 160 number of features in each node.

161 The DiffPool operator, at layer l , thus receives both an adjacency matrix
 162 and a node embedding matrix, and computes updated versions of both:

$$\mathbf{A}^{(l+1)}, \mathbf{X}^{(l+1)} = \text{DiffPool}(\mathbf{A}^{(l)}, \mathbf{X}^{(l)}) \quad (5)$$

163 To achieve this, the DiffPool operator uses a graph neural network (GNN)
 164 architecture. Specifically, the same GNN architecture is duplicated to com-
 165 pute two distinct representations: a new embedding $\mathbf{Z} \in \mathbb{R}^{N_{(l)} \times F'}$ and an
 166 assignment matrix $\mathbf{S} \in \mathbb{R}^{N_{(l)} \times N_{(l+1)}}$:

$$\mathbf{Z}^{(l)} = \text{GNN}_{l, \text{embed}}(\mathbf{A}^{(l)}, \mathbf{X}^{(l)}) \quad (6)$$

$$\mathbf{S}^{(l)} = \text{softmax}(\text{GNN}_{l, \text{pool}}(\mathbf{A}^{(l)}, \mathbf{X}^{(l)})), \quad (7)$$

167 where $N_{(l)}$ is the number of nodes in layer l , $N_{(l+1)}$ the new number of
 168 nodes, each corresponding to a cluster ($N_{(l+1)} < N_{(l)}$), and F' the number of
 169 features per node, which can be different from the original size F from the
 170 matrix \mathbf{X} .

171 The operator ends with the creation of the new node embedding matrix
 172 and adjacency matrix, to be inputted to the next layer:

$$\mathbf{X}^{(l+1)} = \mathbf{S}^{(l)T} \mathbf{Z}^{(l)} \quad (8)$$

$$\mathbf{A}^{(l+1)} = \mathbf{S}^{(l)T} \mathbf{A}^{(l)} \mathbf{S}^{(l)}, \quad (9)$$

173 where $\mathbf{X}^{(l+1)} \in \mathbb{R}^{N_{(l+1)} \times F'}$ and $\mathbf{A}^{(l+1)} \in \mathbb{R}^{N_{(l+1)} \times N_{(l+1)}}$.

174 **3. Experiments Overview**

175 *3.1. Dataset - UK Biobank*

176 Subject-level structural T1 and T2-FLAIR data as well as ICA-FIX [30]
177 denoised rs-fMRI data was obtained from UK BioBank (application 20904) [31]¹.
178 All data were acquired on a standard Siemens Skyra 3T scanner running
179 VD13A SP4, with a standard Siemens 32-channel RF receive head coil.
180 The structural data was further preprocessed with Freesurfer (v6.0)² us-
181 ing the T2-FLAIR weighted image to improve pial surface reconstruction,
182 similarly to Glasser et al. [32]’s pipeline. Reconstruction included bias field
183 correction, registration to stereotaxic space, intensity normalisation, skull-
184 stripping, and white matter segmentation. When no T2-FLAIR data was
185 available, Freesurfer reconstruction was done using the T1 weighted image
186 only. Following surface reconstruction, the Desikan-Killiany atlas [33] was
187 aligned to each individual structural image and and ROIs extracted. The
188 same atlas was aligned to the functional denoised rs-fMRI data (490 volumes
189 TR/TE = 735/39.00 ms, multiband factor 8, voxel size: $2.4 \times 2.4 \times 2.4$, FA=52
190 deg, FOV 210x210 mm) using the warping parameters from the structural to
191 functional alignment obtained using FSL’s linear registration (FLIRT), and
192 mean BOLD time series (490 timepoints per scan) were extracted for each
193 ROI. The time series were then scaled subject-wise using the median and
194 interquartile range according to the *RobustScaler* implementation from the
195 *scikit-learn* [34] python package. Edge weights were defined as full correla-
196 tions calculated with the Ledoit Wolf covariate estimator using the *nilearn*
197 python package³. Figure 1 shows an example scaled time series and the re-
198 sulting example graph from a single subject. The total number of subjects
199 used from the UK Biobank was 35,159, in which 18,649 were females and
200 16,510 were males ($18,649/16,510 \approx 1.13$). The median age was 64 with a
201 minimum age of 44 and a maximum of 81.

202 *3.2. Model Implementation*

203 The neural network architecture depicted in Figure 2 was implemented
204 using Pytorch [35], and Pytorch Geometric [36] for the specific graph neural

¹https://biobank.ctsu.ox.ac.uk/crystal/crystal/docs/brain_mri.pdf

²<http://surfer.nmr.mgh.harvard.edu/>

³<https://nilearn.github.io/>

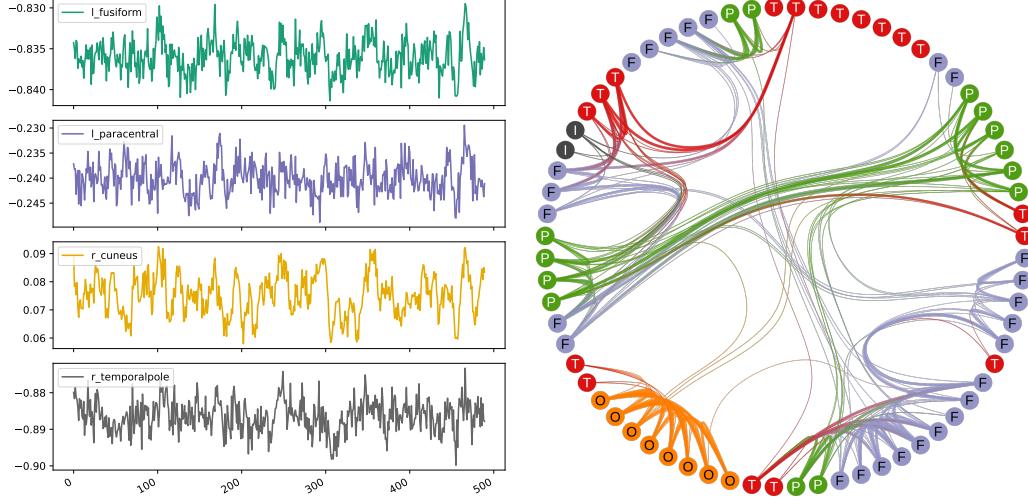


Figure 1: Left: Mean BOLD time series of four brain regions from the same subject, after scaling. Right: Graph representation of one subject’s data, at 10% threshold as described in Section 3.2. Thicker edges represent a stronger correlation between nodes, in this case with values between approximately 0.54 and 0.87. Each node is labelled and coloured according to its brain region (i.e., T/F/O/P/I correspond to Temporal, Frontal, Occipital, Parietal, and Insula).

205 network components. The edge feature matrix $\mathbf{E} \in \mathbb{R}^{E \times 1}$ defined in Section
 206 2.1 was implemented as two sparse matrices: a sparse representation of
 207 the adjacency matrix $\mathbf{E}_i \in \mathbb{R}^{2 \times E}$, and a sparse representation of the edge
 208 features $\mathbf{E}_a \in \mathbb{R}^{E \times 1}$ (i.e., there was only one feature per edge corresponding
 209 to the correlation value). The number of nodes N was 68 (corresponding
 210 to each brain region from the Desikan-Killiany atlas), the number of node
 211 features F was the number of timepoints (i.e., 490), and E is the number of
 212 edges in the graph. The number of edges depends on the threshold percent-
 213 age used to retain only the strongest correlations. Given the non-conclusive
 214 evidence on the optimal threshold percentage in the vast majority of func-
 215 tional connectivity literature [37], in this work this threshold was included
 216 in the hyperparameters to be optimised.

217 The full list of hyperparameters to be optimised and respective value
 218 range was the following:

- 219 • **dropout:** $[0, 0.9]$ (uniform distribution)
- 220 • **threshold:** $\{5, 10, 20, 30, 40\}$ (categorical)

221 • learning rate: $[\ln(1e-7), \ln(1e-2)]$ (log uniform distribution)
222 • weight decay: $[\ln(1e-7), 0]$ (log uniform distribution)

223 To extract information from the rs-fMRI time series in each node, our
224 model starts by employing a strided temporal convolutional network (TCN)
225 architecture. Such architecture was implemented by using two blocks where
226 each contained two layers of *1D convolutions*, *1D batch normalisation*, *ReLU*
227 *activation*, and *dropout*. Each block used a stride of 2, a kernel with size
228 7 (i.e., $K = 7$ in Equation 1), and contained a skip connection. While
229 the first block used no dilation, the second one used a dilation of 2. The
230 four *1D convolution* filters increased the numbers of output channels at each
231 layer, specifically 8, 16, 32, and 64. After these two blocks (i.e., four layers),
232 node features from all channels are flattened out and inputted to a linear
233 transformation to reduce each node representation to a fixed embedding of
234 size 128. These transformations thus reduce the original node feature matrix
235 from size $N \times T$ to size $N \times 64 \times 31$ after the two blocks, and finally to size
236 $N \times 128$ corresponding to the final embedding.

237 Then, the Graph Network (GN) block is applied, in which the update
238 functions are single-layered multi-layer perceptrons (MLPs), and aggrega-
239 tion function is an edge-wise averaging around each node. The original di-
240 mensions of \mathbf{X} , \mathbf{E}_i , and \mathbf{E}_a right before the GN block are kept after these
241 transformations.

242 We employed two types of pooling mechanisms in our analysis, both of
243 which reduce the node feature matrix from a size of $N \times 128$ to 1×128 : a
244 global average pooling mechanism, and the hierarchical pooling mechanism
245 (i.e., DiffPool). For DiffPool, which expects a dense graph representation,
246 data is first transformed into a symmetric adjacency matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$,
247 which is a weighted matrix when considering edge features, and binary oth-
248 erwise. As recommended by the original paper [29], we employed three layers
249 of GraphSAGE [38] followed by a *1D batch normalisation*, with a final skip
250 connection.

251 A summary conceptual architecture of the whole model is shown in Fig-
252 ure 2.

253 3.3. *Training Procedure*

254 In order to assess the validity of our model, we performed a proof-of-
255 concept through the well-known binary sex prediction task [39, 40]. We

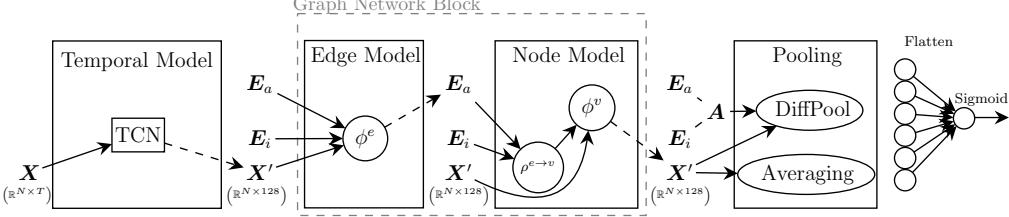


Figure 2: Three main working blocks of the spatio-temporal model. The temporal model creates an initial representation from original node features \mathbf{X} (i.e., temporal dynamics). It is followed by transformations in the Graph Network Block which leverages the structure of data represented in edge features \mathbf{E}_a and sparse connectivity \mathbf{E}_i . Finally, a Pooling mechanism (either DiffPool or global averaging) creates a final graph representation which is flatten for a final prediction task.

256 used a 5-fold stratified cross validation procedure: the UK Biobank dataset
 257 was divided into training and test sets five times, in which each test set
 258 corresponds to 20% of the original size, and a sample would only belong to
 259 a test set once (i.e., all test sets are mutually exclusive). This division was
 260 done in a stratified fashion considering the sex label, bucketised age, and
 261 bucketised BMI measures (for each variable we created 8 equal-sized buckets
 262 based on sample quantiles). For each test set, the training set is further
 263 divided once to generate a single inner training and validation sets, using the
 264 same stratification strategy as for the training/test case.

265 The neural network was trained over 100 epochs with the Adam optimiser [41] and Binary Cross Entropy loss function. The training procedure
 266 was set to stop earlier if the validation loss did not reduce further after
 267 33 consecutive epochs. A hyperparameter search was included in the inner
 268 training/validation sets, in which 25 random runs were launched exploring
 269 random values of dropout, edge threshold, learning rate, and weight decay
 270 (see Section 3.2 for values range). In each random run, the model with the
 271 smallest validation loss was saved, and the model with the smallest validation
 272 loss across the 25 runs was selected to be evaluated in the test set. This pro-
 273 cedure is done separately for each test set, and metrics are averaged across
 274 the five test sets.

276 We used *Weights & Biases* [42] to log our training procedure and generate
 277 the random hyperparameters for all the 25 models in each inner sweep. These
 278 inner sweeps were run across two different servers, and each model took
 279 between 20 minutes and 11 hours to train depending on GPU type and early
 280 stopping. All these details are stored using *Weights & Biases*, and can be

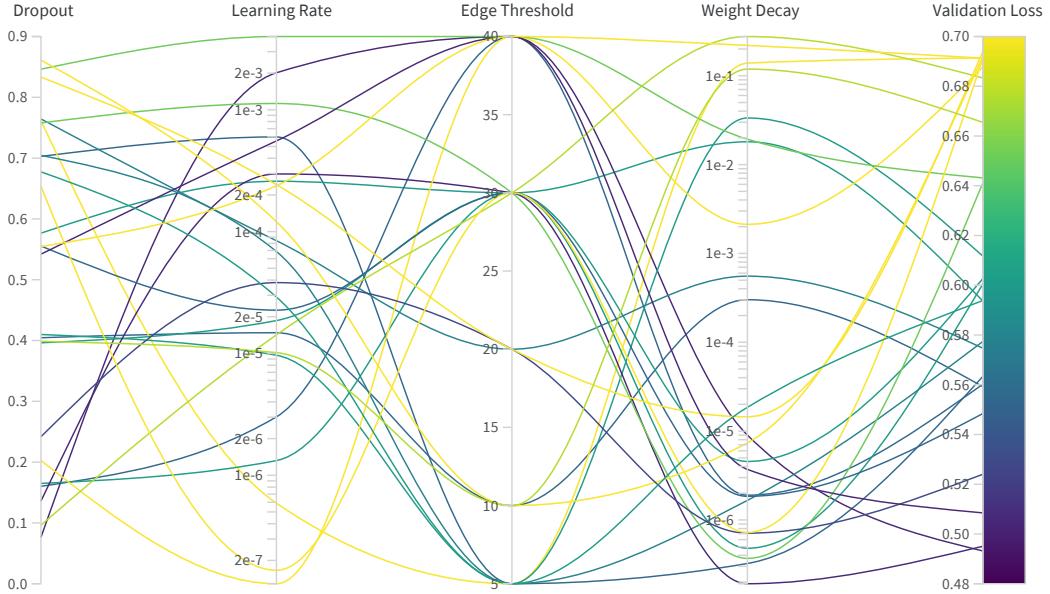


Figure 3: Values of hyperparameters corresponding to each validation loss achieved for one illustrative inner sweep of one fold. For each one of the 25 training runs (represented by each curved line), a set of random values is chosen for dropout, learning rate, edge threshold and weight decay, which ultimately will result in the model’s validation loss.

281 accessed through our public repository (see “Data and Code Availability”).
 282 Figure 3 shows the results for the inner sweep of one of the folds for illustrative
 283 purposes. While a certain amount of variability is visible, some trends are
 284 evident in this particular split: the best models (i.e., with lower validation
 285 loss) tend to be achieved with higher edge thresholds, higher learning rates,
 286 smaller weight decays, and smaller dropout rates. We highlight that different
 287 sweeps could result in very different trends.

288 *3.4. Evaluation*

289 As shown in Figure 2, our model consists of (1) a TCN block that learns
 290 intra-temporal features from the mean BOLD time series of each ROI, fol-
 291 lowed by (2) a GN block which leverages the spatial inter-relationships be-
 292 tween ROIs, and finally (3) a hierarchical pooling mechanism which leverages
 293 all the information in the input, from the temporal rs-fMRI dynamics, to the
 294 graph structure and the edge features of that graph.

295 In order to understand the inner workings of this combination, we con-
 296 ducted an ablation analysis to quantify the contributions of each component

297 of our model for the specific prediction task. Firstly, we consider two cases
298 where the GN block is not used, hence practically evaluating the importance
299 of edge weights for this prediction task. In one case the graph structure is
300 completely ignored (i.e., no GN block and global average pooling), and in
301 another case a binary graph is used only for the final hierarchical pooling
302 part (i.e., no GN block and DiffPool applied to a binary graph).

303 In order to investigate the influence of the different GN components, we
304 consider not only the case where both *node model* and *edge model* are used
305 in the GN Block, but also a case where only the *node model* is applied. For
306 each one of these two cases, we consider both a global average pooling, as
307 well as DiffPool with a weighted adjacency matrix.

308 **4. Results**

309 *4.1. General Results*

310 Table 1 shows the results of our ablation analysis across three different
311 backbones - no graph block, only *node model*, and full graph network
312 block - each with two different aggregators (i.e., global average pooling and
313 DiffPool). For notation purposes, we identify each one of these cases using
314 “Backbone → Aggregator”, in which *Aggregator* can be “Average” or “Diff-
315 Pool”, and *Backbone* can be “N” for only *node model*, “N + E” for both *node*
316 *model* and *edge model* (i.e., full GN Block), and empty otherwise.

Table 1: Ablation analysis, with metrics averaged across the five test sets, with standard deviation in parenthesis. Aggregator on the right-hand side of the arrow, “N” corresponds to only *node model*, and “N + E” corresponds to full Graph Network block. **Params** stands for number of parameters.

Model	AUC	Accuracy	Sensitivity	Specificity	Params
N + E → Average	0.83 (0.013)	0.76 (0.012)	0.74 (0.035)	0.77 (0.032)	474,639
N + E → DiffPool	0.84 (0.009)	0.76 (0.010)	0.78 (0.042)	0.74 (0.055)	849,057
N → Average	0.80 (0.006)	0.73 (0.007)	0.71 (0.037)	0.74 (0.030)	441,486
N → DiffPool	0.84 (0.009)	0.76 (0.008)	0.77 (0.014)	0.75 (0.016)	815,904
→ DiffPool	0.81 (0.021)	0.73 (0.019)	0.73 (0.035)	0.73 (0.049)	648,979
→ Average	0.77 (0.004)	0.70 (0.005)	0.68 (0.004)	0.73 (0.008)	274,561

317 Using a GNN component achieves better results overall when compared
318 with the “ \rightarrow Average” case (i.e., no GNN), with clearer gains for AUC and
319 accuracy metrics. Using DiffPool as an aggregator appears to deliver the
320 best averaged metrics for AUC, accuracy, and sensitivity. The simplest GNN
321 model in terms of number of parameters (i.e., “ $N \rightarrow$ Average”) yields com-
322 parable results to other models, thus likely being the best compromise in
323 terms of model complexity and performance power. Using the *edge model*
324 did not bring significantly better results when compared to only using the
325 *node model*, thus indicating that the information contained in the edge at-
326 tributes is successfully leveraged by the *node model* alone for this particular
327 prediction task.

328 The results presented so far consider an adjacency matrix threshold be-
329 low 50% as a hyperparameter at training time, a common data reduction
330 practice in the connectivity analysis field. We further analysed the results
331 of using no threshold at all, and explored the type of activation function as
332 a hyperparameter instead (i.e., *ReLU* or *tanh* activations). This choice was
333 made explicitly since retaining 100% of the adjacency matrix elements results
334 in a share of negative correlation elements, whose physiological significance is
335 likely to be important in brain connectivity [43]. The results of this analysis
336 are presented in Table 2.

Table 2: Results with no thresholded graphs, with metrics averaged across the five test sets, with standard deviation in parenthesis. Aggregator on the right-hand side of the arrow, “ N ” corresponds to only *node model*, and “ $N + E$ ” corresponds to full Graph Network block. **Params** stands for number of parameters.

Model	AUC	Accuracy	Sensitivity	Specificity	Params
$N + E \rightarrow$ Average	0.79 (0.017)	0.72 (0.013)	0.71 (0.039)	0.73 (0.027)	474,639
$N + E \rightarrow$ DiffPool	0.79 (0.045)	0.69 (0.077)	0.55 (0.264)	0.81 (0.098)	849,057
$N \rightarrow$ Average	0.80 (0.032)	0.72 (0.028)	0.68 (0.077)	0.76 (0.059)	441,486
$N \rightarrow$ DiffPool	0.81 (0.009)	0.74 (0.009)	0.73 (0.060)	0.74 (0.042)	815,904

337 The performance was lower for all cases which did not implement a thresh-
338 old. A possible explanation would be the excessive “noise” not allowing the
339 dominating spatial structure of the graph to be successfully leveraged in a
340 practical timeframe, possibly generating some overfitting. However, metrics
341 are still comparable to the ones in Table 1, suggesting that these models are

342 still able to extract some information from the data within only 100 epochs
343 of training.

344 *4.2. Explainability*

345 Although deep neural networks are usually regarded as “black boxes”, in
346 this paper we strived to inject some elements of explainability by inspecting
347 selected learnt mechanisms after training. For instance, the weights of the
348 TCN layers can be visually inspected. We visualised the first two layers of one
349 of the trained $N + E \rightarrow$ DiffPool model trained on unthresholded matrices.
350 For example, Figure 4 shows the weights learnt from the first TCN layer
351 (each row corresponding to one of the 8 output channels of that layer), while
352 Figure 5 depicts the same for the second TCN layer (each row corresponding
353 to one of the 16 output channels and the columns corresponding to the 8
354 kernels of size 7 coming from the previous 8 channels).

355 In both figures, and with just a
356 few exceptions, it can be seen that
357 the output channels in the first two
358 TCN convolutional layers are a non-
359 trivial weighted multiplication of in-
360 put channels. Given the qualitative
361 variability observed in these learnt
362 weights, we argue that these kernels
363 are likely filtering and selecting dif-
364 ferent, non-mutually redundant pat-
365 terns presented in the original time
366 series. One possible counterexample
367 is the kernel for the 7th output chan-
368 nel in the first TCN convolutional
369 layer illustrated in Figure 4, which,
370 in practice, is applying a simple low
371 pass filter by smoothing the original
372 time series from the input channel.
373 We posit that quantitative analysis
374 and comparison of the kernel weights
375 across time series (which is out of the scope of the present paper), has the
376 potential to yield interpretable information on which brain dynamics may
377 contribute the most to the final prediction. It can also potentially high-
378 light what frequency components (likely attached to different physiological

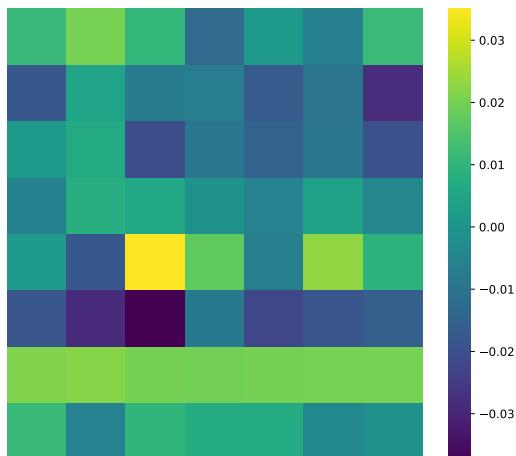


Figure 4: Weights of the kernels in the first TCN convolutional layer in a $N + E \rightarrow$ DiffPool model trained on unthresholded graphs. Rows correspond to the 8 output channels of this layer, and each column is a position in the kernel array of size 7

379 significance [44]) are selected most often.

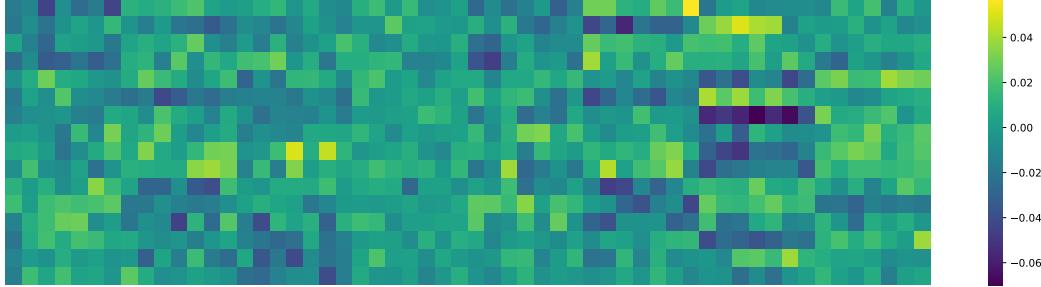


Figure 5: Weights of the kernels in the second TCN convolutional layer in a $N + E \rightarrow$ DiffPool model trained on unthresholded graphs. Rows correspond to the 16 output channels of this layer, and each column is a position in the 8 kernels of size 7 that come from the 8 input channels (56 columns in total).

380 We further designed a strategy to inspect the hierarchical spatial pooling
381 mechanism provided by the DiffPool architecture. To this end, we analysed
382 the assignment matrices from the first Diffpool layer $\mathbf{S}^{(1)}$ (see Equation 7),
383 over all participants across all test sets. This is of particular interest be-
384 cause it corresponds to an aggregation of subsets of brain regions which our
385 architecture has considered optimal while learning a particular prediction
386 task. These aggregations can therefore be considered “optimal” for that task
387 within this architecture, and provide hints to merge with neurophysiological
388 knowledge. An assignment matrix corresponds to how the original nodes in
389 the graph will be mapped into new nodes, thus, a direct way of summarising
390 this effect across individuals is to count how many times two ROIs have ended
391 up in the same cluster, regardless of cluster size and number. More formally,
392 we create an association matrix $\mathbf{S}' \in \mathbb{R}^{68 \times 68}$, where each element $S'_{i,j}$ is the
393 number of times brain regions i and j have been aggregated in the same
394 DiffPool cluster. This means that the higher the value of $S'_{i,j}$, the more often
395 information from brain regions i and j is pooled when learning to predict
396 binary sex. It is important to point out that, given that matrix thresholding
397 can potentially disconnect nodes from the rest of the graph and hence yield
398 subject-wise graphs with variable numbers of connected nodes, this inspec-
399 tion strategy is only possible when working with unthresholded matrices (i.e.,
400 $N \rightarrow$ DiffPool and $N + E \rightarrow$ DiffPool models, see “Experiments Overview”
401 section).

402 Figure 6 depicts the association matrix \mathbf{S}' for the $N + E \rightarrow$ DiffPool
403 model trained on unthresholded matrices, with dendrograms resulting from

404 hierarchical clustering of the summary association matrix elements itself.
 405 For visualisation purposes in a more traditional brain connectivity style, we
 406 selected the four main clusters defined by the dendrograms for the N +
 407 E → DiffPool model and overlaid these clusters on a brain surface in Figure 7.

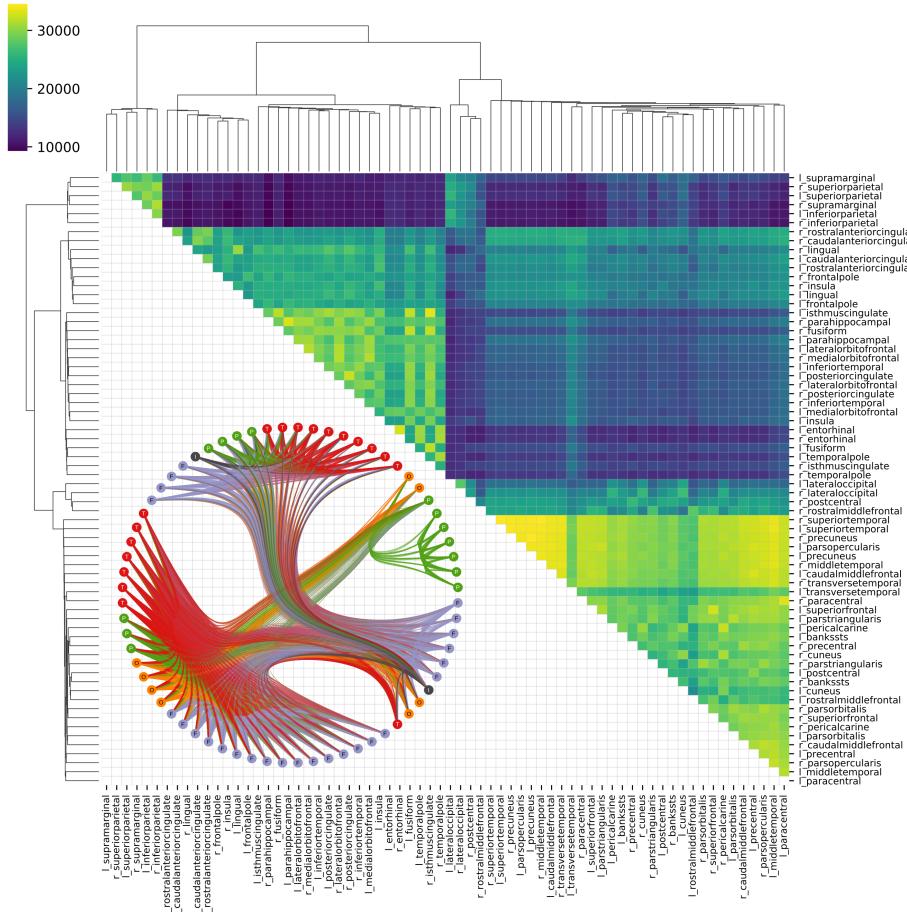


Figure 6: Upper-triangle of the association matrix \mathbf{S}' for $N + E \rightarrow$ DiffPool model generated when predicting binary sex on unthresholded matrices, with dendrograms from hierarchical clustering. Each element $S'_{i,j}$ indicates how many times brain regions i and j are pooled together. On the lower left corner, a graph representation of the same association matrix \mathbf{S}' , thresholded at 50% with nodes identified and coloured according to their general brain region (i.e., T/F/O/P/I correspond to Temporal, Frontal, Occipital, Parietal, and Insula); thicker edges represent a higher $S'_{i,j}$ value, in this graph representation ranging from 20,256 to 34,565.

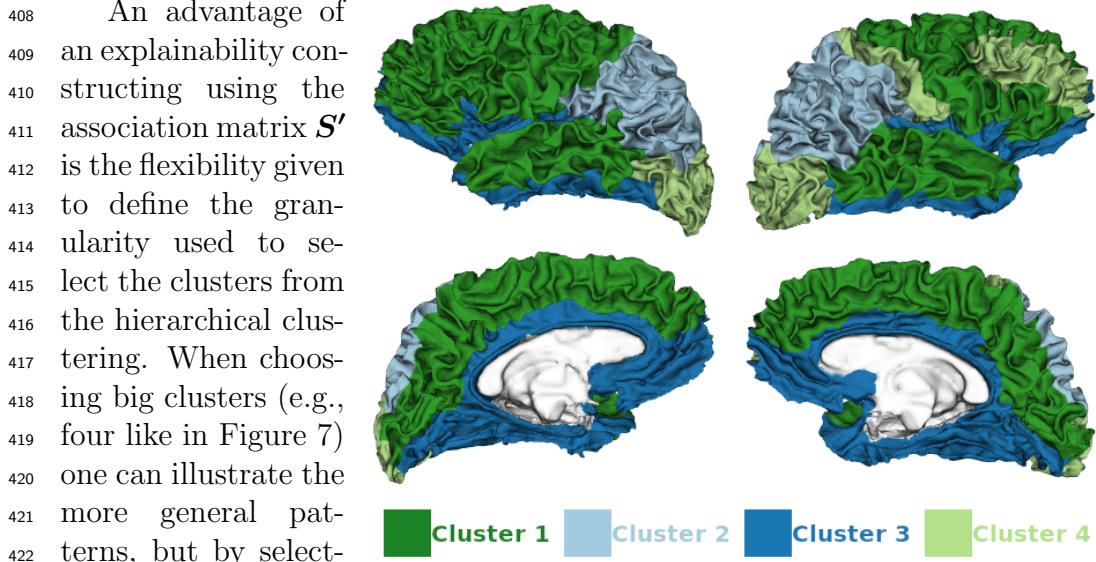


Figure 7: Four main brain clusters on association matrix S' generated from $N + E \rightarrow$ DiffPool model predicting binary sex on unthresholded matrices. Each colour corresponds to one cluster.

432 When looking at how the GNNs clustered the brain regions to optimise
433 and achieve best sex prediction, we found that a clustering into 4 sets of
434 brain regions showed interesting properties in terms of neurobiological inter-
435 pretability. More specifically, the brain regions were grouped in a manner
436 that mirrors to a certain degree the well-known cytoarchitectural and func-
437 tional properties of the cerebral cortex. For example, in Figure 7 cluster 1
438 (dark green) included high order associative brain areas such the prefrontal
439 and temporal cortices that have been consistently involved in complex cog-
440 nitive functions such as language, working memory, and decision-making. In
441 cluster 2, represented by the light blue colour, the GNN grouped together
442 the left and right posterior parietal cortices which have a well-known role in
443 visuo-spatial processing and navigation skills, amongst many other cognitive
444 functions. Cluster 3 (in dark blue) included midline cortical areas which
445 are part of the classic limbic or emotional system. The final cluster 4, in

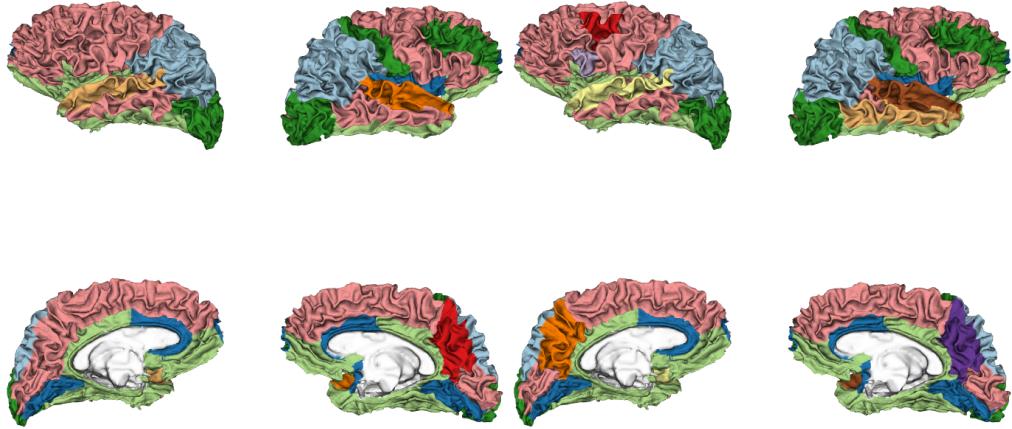


Figure 8: Main brain clusters on association matrix S' generated from $N + E \rightarrow$ DiffPool model predicting binary sex on unthresholded matrices. Each colour corresponds to one cluster. Left: Eight main clusters. Right: Twelve main clusters.

446 light green, was mainly localised in the right hemisphere and in part in the
447 left visual cortex; it grouped together sensory-motor cortices and the right
448 dorsolateral prefrontal cortex.

449 We do not wish to overinterpret our results or make “reverse neuro-
450 science” inferences in the sense of interpreting *post hoc* the behavioural mean-
451 ing of a set of regions without having directly analysed their behavioural rel-
452 evance. However, we speculatively note that the clusters that discriminated
453 binary sex the most may have some neurobiological relevance in terms of ex-
454 plaining the well-known behavioural differences between males and females
455 in terms of cognitive, motor and emotional skills [45, 46]. Future work, par-
456 ticularly directly at investigating the links between brain and behavioural
457 measures, is warranted to confirm whether the clustering of regions that our
458 model has performed to achieve optimal sex classification is fully relevant to
459 mechanistically describe the sex differences that are seen at the behavioural
460 level.

461 These results demonstrate the explainability capacity of our model when
462 using the DiffPool aggregator, which is able to cluster brain regions in a
463 specific way for the classification task at hand.

464 **5. Conclusion**

465 In this paper we presented a novel deep learning architecture which can
466 successfully use the high-dimensional and noisy rs-fMRI data, by leveraging
467 not only their temporal dynamics, but also their spatial associations repre-
468 sented by what is commonly called the connectivity between brain locations.
469 In contrast to previous work, we use TCNs to model temporal intra-relations
470 and combine them with GNNs to model inter-regional relations. We illus-
471 trated and analysed the effectiveness of our model in a proof-of-concept bi-
472 nary sex prediction task which also included an ablation analysis with vari-
473 ations of the spatial pooling mechanisms. This work is, to the best of our
474 knowledge, the first to leverage both the spatial and temporal information
475 in rs-fMRI data in a single, end-to-end framework which includes temporal
476 convolutions and graph neural networks, while also providing the flexibil-
477 ity to extract human readable explainability. Importantly, we included edge
478 features (i.e., weights) when leveraging the graph structure in the network.
479 This information is often ignored in the few papers which currently apply
480 GNNs to the study of fMRI data [47]. Our ablation study showed how the
481 graph network block was successful in leveraging the weights of the spatial
482 dynamics, indicating the importance of designing an architecture specifically
483 targeted for spatio-temporal rs-fMRI data. We also showed the explainabil-
484 ity capacities of our models by analysing the clusters created by the graph
485 hierarchical pooling mechanism, and the non-linear patterns learnt from the
486 rs-fMRI time series.

487 We hope this paper can lay future groundwork on exploring flexible ar-
488 chitectures which are able to leverage the entirety of neuromonitoring data
489 that arise from the extremely complex spatio-temporal interplay of groups of
490 firing neurons. Our architecture can very easily include other types of data
491 for future work (e.g., multimodal structural and temporal data), and be ex-
492 tended to include possible confounds that could drive the prediction task in
493 other brain disorders. Another exciting recent trend that can be included in
494 our architecture is to allow the network to learn the underlying connectivity
495 from scratch [48, 49] instead of computing associations or other handcrafted
496 features like the ones used in this and other works [50, 51].

497 **Data and Code Availability**

498 The code used to process the data from the UK Biobank is publicly
499 available at <https://github.com/ucam-department-of-psychiatry/UKB>.

500 The code used to conduct the analysis described in this paper is publicly
501 available at <https://github.com/tjiagoM/spatio-temporal-brain>.

502 CRediT Authorship Contribution Statement

503 **T.A.:** Conceptualization, Methodology, Software, Validation, Investigation, Writing - Original Draft, Writing - Review & Editing, Visualization.
504 **A.C.:** Conceptualization, Methodology, Writing - Review & Editing.
505 **R.R.G.:** Resources, Data Curation. **L.P.:** Validation, Writing - Review &
506 Editing **R.A.I.B.:** Resources, Data Curation, Writing - Review & Editing,
507 Visualization. **P.L.:** Conceptualization, Resources, Writing - Review & Edit-
508 ing, Supervision, Funding acquisition. **N.T.:** Conceptualization, Methodol-
509 ogy, Resources, Writing - Review & Editing, Supervision, Funding acquisi-
510 tion.
511

512 Declaration of Competing Interest

513 The authors declare that they have no known competing financial inter-
514 ests or personal relationships that could have appeared to influence the work
515 reported in this paper.

516 Acknowledgements

517 T.A. is funded by the W. D. Armstrong Trust Fund, University of Cam-
518 bridge, UK. R.A.I.B. is funded by a British Academy Post-Doctoral fellow-
519 ship and the Autism Research Trust. R.R.G is funded by the Guarantors
520 of Brain. L.P. is funded by the Medical Research Council (MRC) grant
521 (MR/P01271X/1) at the University of Cambridge, UK.

522 The UK Biobank data (application 20904) were curated and analysed
523 using a computational facility funded by an MRC research infrastructure
524 award (MR/M009041/1) and supported by the NIHR Cambridge Biomed-
525 ical Research Centre and a Marmaduke Shield Award to Dr. Richard A.I.
526 Bethlehem and Varun Warrier. The views expressed are those of the authors
527 and not necessarily those of the NHS, the NIHR or the Department of Health
528 and Social Care.

529 The models in this work were developed and evaluated on two differ-
530 ent servers. Runs were performed using resources provided by the Cam-
531 bridge Service for Data Driven Discovery (CSD3) operated by the Uni-
532 versity of Cambridge Research Computing Service (www.cs3.cam.ac.uk),

533 provided by Dell EMC and Intel using Tier-2 funding from the Engineering
534 and Physical Sciences Research Council (capital grant EP/P020259/1),
535 and DiRAC funding from the Science and Technology Facilities Council
536 (www.dirac.ac.uk). We also used Titan V GPUs generously donated to
537 N.T. by NVIDIA.

538 **References**

539 [1] A. Fornito, A. Zalesky, M. Breakspear, The connectomics of brain dis-
540 orders, *Nature Reviews Neuroscience* 16 (2015) 159–172.

541 [2] J. Li, R. Kong, R. Liégeois, C. Orban, Y. Tan, N. Sun, A. J. Holmes,
542 M. R. Sabuncu, T. Ge, B. T. Yeo, Global signal regression strengthens
543 association between resting-state functional connectivity and behavior,
544 *NeuroImage* 196 (2019) 126–141.

545 [3] C. F. Beckmann, M. DeLuca, J. T. Devlin, S. M. Smith, Investigations
546 into resting-state connectivity using independent component analysis,
547 *Philosophical Transactions of the Royal Society B: Biological Sciences*
548 360 (2005) 1001–1013.

549 [4] S. M. Smith, T. E. Nichols, Statistical challenges in “big data” human
550 neuroimaging, *Neuron* 97 (2018) 263–268.

551 [5] A. Avena-Koenigsberger, B. Misic, O. Sporns, Communication dynam-
552 ics in complex brain networks, *Nature Reviews Neuroscience* 19 (2017)
553 17–33.

554 [6] X. Liao, M. Cao, M. Xia, Y. He, Individual differences and time-varying
555 features of modular brain architecture, *NeuroImage* 152 (2017) 94–107.

556 [7] A. Duggento, L. Passamonti, G. Valenza, R. Barbieri, M. Guerrisi,
557 N. Toschi, Multivariate granger causality unveils directed parietal to
558 prefrontal cortex connectivity during task-free MRI, *Scientific Reports*
559 8 (2018).

560 [8] G. Goelman, R. Dan, T. Keadan, Characterizing directed functional
561 pathways in the visual system by multivariate nonlinear coherence of
562 fMRI data, *Scientific Reports* 8 (2018). URL: <https://doi.org/10.1038/s41598-018-34672-5>. doi:10.1038/s41598-018-34672-5.

564 [9] S. Spasov, L. Passamonti, A. Duggento, P. Liò, N. Toschi, A parameter-
565 efficient deep learning approach to predict conversion from mild cognitive
566 impairment to alzheimer's disease, *NeuroImage* 189 (2019) 276–287.

567 [10] A. Duggento, M. Guerrisi, N. Toschi, Recurrent neural networks for
568 reconstructing complex directed brain connectivity, in: 2019 41st Annual
569 International Conference of the IEEE Engineering in Medicine and
570 Biology Society (EMBC), IEEE, 2019.

571 [11] J. Zhou, G. Cui, Z. Zhang, C. Yang, Z. Liu, L. Wang, C. Li, M. Sun,
572 Graph neural networks: A review of methods and applications, *arXiv* preprint
573 arXiv:1812.08434 (2018).

574 [12] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, P. S. Yu, A comprehensive
575 survey on graph neural networks, *arXiv* preprint arXiv:1901.00596
576 (2019).

577 [13] S. Arslan, S. I. Ktena, B. Glocker, D. Rueckert, Graph saliency maps
578 through spectral convolutional networks: Application to sex classification
579 with brain connectivity, in: *Lecture Notes in Computer Science*, Springer
580 International Publishing, 2018, pp. 3–13.

581 [14] D. Wen, Z. Wei, Y. Zhou, G. Li, X. Zhang, W. Han, Deep learning
582 methods to process fMRI data and their application in the diagnosis of
583 cognitive impairment: A brief overview and our opinion, *Frontiers in
584 Neuroinformatics* 12 (2018).

585 [15] A. Riaz, M. Asad, E. Alonso, G. Slabaugh, DeepFMRI: End-to-end
586 deep learning for functional connectivity and classification of ADHD
587 using fMRI, *Journal of Neuroscience Methods* 335 (2020) 108506.

588 [16] S. Yan, Y. Xiong, D. Lin, Spatial temporal graph convolutional net-
589 works for skeleton-based action recognition, in: *Thirty-second AAAI*
590 conference on artificial intelligence, 2018.

591 [17] Z. Diao, X. Wang, D. Zhang, Y. Liu, K. Xie, S. He, Dynamic spatial-
592 temporal graph convolutional neural networks for traffic forecasting,
593 in: *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33,
594 2019, pp. 890–897.

595 [18] S. Guo, Y. Lin, N. Feng, C. Song, H. Wan, Attention based spatial-
596 temporal graph convolutional networks for traffic flow forecasting, in: Proceedings of the AAAI Conference on Artificial Intelligence, volume 33, 2019, pp. 922–929.

599 [19] A. B. Arrieta, N. Díaz-Rodríguez, J. D. Ser, A. Bennetot, S. Tabik,
600 A. Barbado, S. Garcia, S. Gil-Lopez, D. Molina, R. Benjamins,
601 R. Chatila, F. Herrera, Explainable artificial intelligence (XAI): Concepts,
602 taxonomies, opportunities and challenges toward responsible AI,
603 Information Fusion 58 (2020) 82–115.

604 [20] W. Samek, G. Montavon, A. Vedaldi, L. K. Hansen, K.-R. Müller, Explainable AI: interpreting, explaining and visualizing deep learning, volume 11700, Springer Nature, 2019.

607 [21] T. Azevedo, L. Passamonti, P. Lio, N. Toschi, A deep spatiotemporal
608 graph learning architecture for brain connectivity analysis, in: 2020 42nd
609 Annual International Conference of the IEEE Engineering in Medicine
610 & Biology Society (EMBC), IEEE, 2020.

611 [22] S. Marek, B. Tervo-Clemmens, F. J. Calabro, D. F. Montez, B. P. Kay,
612 A. S. Hatoum, M. R. Donohue, W. Foran, R. L. Miller, E. Feczko,
613 O. Miranda-Dominguez, A. M. Graham, E. A. Earl, A. J. Perrone,
614 M. Cordova, O. Doyle, L. A. Moore, G. Conan, J. Uriarte, K. Snider,
615 A. Tam, J. Chen, D. J. Newbold, A. Zheng, N. A. Seider, A. N. Van,
616 T. O. Laumann, W. K. Thompson, D. J. Greene, S. E. Petersen, T. E.
617 Nichols, B. T. Yeo, D. M. Barch, H. Garavan, B. Luna, D. A. Fair, N. U.
618 Dosenbach, Towards reproducible brain-wide association studies (2020).

619 [23] S. Bai, J. Z. Kolter, V. Koltun, An empirical evaluation of generic con-
620 volutional and recurrent networks for sequence modeling, arXiv preprint
621 arXiv:1803.01271 (2018).

622 [24] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals,
623 A. Graves, N. Kalchbrenner, A. Senior, K. Kavukcuoglu, Wavenet: A
624 generative model for raw audio, 2016. [arXiv:arXiv:1609.03499](https://arxiv.org/abs/1609.03499).

625 [25] F. Yu, V. Koltun, Multi-scale context aggregation by dilated convolu-
626 tions, in: ICLR, 2016.

627 [26] P. W. Battaglia, J. B. Hamrick, V. Bapst, A. Sanchez-Gonzalez,
628 V. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro,
629 R. Faulkner, C. Gulcehre, F. Song, A. Ballard, J. Gilmer, G. Dahl,
630 A. Vaswani, K. Allen, C. Nash, V. Langston, C. Dyer, N. Heess,
631 D. Wierstra, P. Kohli, M. Botvinick, O. Vinyals, Y. Li, R. Pascanu,
632 Relational inductive biases, deep learning, and graph networks, 2018.
633 [arXiv:arXiv:1806.01261](https://arxiv.org/abs/1806.01261).

634 [27] S. J. Kiebel, J. Daunizeau, K. J. Friston, A hierarchy of time-scales and
635 the brain, PLoS Computational Biology 4 (2008) e1000209.

636 [28] C. C. Hilgetag, A. Goulas, ‘hierarchy’ in the organization of brain net-
637 works, Philosophical Transactions of the Royal Society B: Biological
638 Sciences 375 (2020) 20190319.

639 [29] Z. Ying, J. You, C. Morris, X. Ren, W. Hamilton, J. Leskovec, Hierar-
640 chical graph representation learning with differentiable pooling, in: Ad-
641 vances in neural information processing systems, 2018, pp. 4800–4810.

642 [30] G. Salimi-Khorshidi, G. Douaud, C. F. Beckmann, M. F. Glasser,
643 L. Griffanti, S. M. Smith, Automatic denoising of functional MRI data:
644 Combining independent component analysis and hierarchical fusion of
645 classifiers, NeuroImage 90 (2014) 449–468.

646 [31] C. Bycroft, C. Freeman, D. Petkova, G. Band, L. T. Elliott, K. Sharp,
647 A. Motyer, D. Vukcevic, O. Delaneau, J. O’Connell, A. Cortes, S. Welsh,
648 A. Young, M. Effingham, G. McVean, S. Leslie, N. Allen, P. Donnelly,
649 J. Marchini, The UK biobank resource with deep phenotyping and
650 genomic data, Nature 562 (2018) 203–209.

651 [32] M. F. Glasser, S. N. Sotiropoulos, J. A. Wilson, T. S. Coalson, B. Fis-
652 chl, J. L. Andersson, J. Xu, S. Jbabdi, M. Webster, J. R. Polimeni,
653 D. C. V. Essen, M. Jenkinson, The minimal preprocessing pipelines for
654 the human connectome project, NeuroImage 80 (2013) 105–124.

655 [33] R. S. Desikan, F. Ségonne, B. Fischl, B. T. Quinn, B. C. Dickerson,
656 D. Blacker, R. L. Buckner, A. M. Dale, R. P. Maguire, B. T. Hyman,
657 M. S. Albert, R. J. Killiany, An automated labeling system for subdivid-
658 ing the human cerebral cortex on MRI scans into gyral based regions
659 of interest, NeuroImage 31 (2006) 968–980.

660 [34] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion,
661 O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vander-
662 plas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay,
663 Scikit-learn: Machine learning in Python, *Journal of Machine Learning
664 Research* 12 (2011) 2825–2830.

665 [35] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan,
666 T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf,
667 E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner,
668 L. Fang, J. Bai, S. Chintala, Pytorch: An imperative style, high-
669 performance deep learning library, in: H. Wallach, H. Larochelle,
670 A. Beygelzimer, F. d'Alché-Buc, E. Fox, R. Garnett (Eds.), *Advances
671 in Neural Information Processing Systems 32*, Curran Associates, Inc.,
672 2019, pp. 8024–8035.

673 [36] M. Fey, J. E. Lenssen, Fast graph representation learning with PyTorch
674 Geometric, in: *ICLR Workshop on Representation Learning on Graphs
675 and Manifolds*, 2019.

676 [37] K. A. Garrison, D. Scheinost, E. S. Finn, X. Shen, R. T. Constable,
677 The (in)stability of functional brain network measures across thresholds,
678 *NeuroImage* 118 (2015) 651–661.

679 [38] W. Hamilton, Z. Ying, J. Leskovec, Inductive representation learning
680 on large graphs, in: *Advances in neural information processing systems*,
681 2017, pp. 1024–1034.

682 [39] R. Jiang, V. D. Calhoun, L. Fan, N. Zuo, R. Jung, S. Qi, D. Lin,
683 J. Li, C. Zhuo, M. Song, Z. Fu, T. Jiang, J. Sui, Gender differences
684 in connectome-based predictions of individualized intelligence quotient
685 and sub-domain scores, *Cerebral Cortex* (2019).

686 [40] S. Weis, K. R. Patil, F. Hoffstaedter, A. Nostro, B. T. T. Yeo, S. B.
687 Eickhoff, Sex classification by resting state brain connectivity, *Cerebral
688 Cortex* (2019).

689 [41] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization,
690 arXiv preprint arXiv:1412.6980 (2014).

691 [42] L. Biewald, Experiment tracking with weights and biases, 2020. URL:
692 <https://www.wandb.com/>, software available from wandb.com.

693 [43] L. Zhan, L. M. Jenkins, O. E. Wolfson, J. J. GadElkarim, K. Noc-
694 ito, P. M. Thompson, O. A. Ajilore, M. K. Chung, A. D. Leow, The
695 significance of negative correlations in brain connectivity, *Journal of*
696 *Comparative Neurology* 525 (2017) 3251–3265.

697 [44] N. H. Yuen, N. Osachoff, J. J. Chen, Intrinsic frequencies of the resting-
698 state fMRI signal: The frequency dependence of functional connectivity
699 and the effect of mode mixing, *Frontiers in Neuroscience* 13 (2019).

700 [45] R. Nieuwenhuys, J. Voogd, C. van Huijzen, *The Human Central Nervous*
701 *System*, Springer Berlin Heidelberg, 2008.

702 [46] D. Poeppel, G. R. Mangun, M. S. Gazzaniga (Eds.), *The cognitive neu-*
703 *rosciences*, sixth edition ed., The MIT Press, Cambridge, MA, 2020.

704 [47] B.-H. Kim, J. C. Ye, Understanding graph isomorphism network for
705 rs-fMRI functional connectivity analysis, *Frontiers in Neuroscience* 14
706 (2020).

707 [48] A. Kazi, L. Cosmo, N. Navab, M. Bronstein, Differentiable
708 graph module (dgm) for graph convolutional networks, 2020.
709 [arXiv:arXiv:2002.04999](https://arxiv.org/abs/2002.04999).

710 [49] B. Jie, M. Liu, C. Lian, F. Shi, D. Shen, Designing weighted correla-
711 tion kernels in convolutional neural networks for functional connectivity
712 based brain disease diagnosis, *Medical Image Analysis* 63 (2020) 101709.

713 [50] X. Li, Y. Zhou, S. Gao, N. Dvornek, M. Zhang, J. Zhuang, S. Gu,
714 D. Scheinost, L. Staib, P. Ventola, J. Duncan, BrainGNN: Interpretable
715 brain graph neural network for fMRI analysis (2020).

716 [51] X. Li, Y. Zhou, N. C. Dvornek, M. Zhang, J. Zhuang, P. Ventola, J. S.
717 Duncan, Pooling regularized graph neural network for fmri biomarker
718 analysis (2020). [arXiv:arXiv:2007.14589](https://arxiv.org/abs/2007.14589).