Dynamic activation patterns in brain MRI data

Kirell BENZI[1], Benjamin RICAUD[1], Alessandra GRIFFA[2], Pierre VANDERGHEYNST[1], Jean-Phillipe THIRAN[2], Patric HAGMANN[2]

[1]Laboratoire de Traitement des Signaux 2
Ecole Polytechnique Fédérale de Lausanne - EPFL, CH-1015 Lausanne, Switzerland

[2]Laboratoire de Traitement des Signaux 5
Ecole Polytechnique Fédérale de Lausanne, EPFL, CH-1015 Lausanne, Switzerland
first.last@epfl.ch

Résumé – L'IRM fonctionnel enregistrant l'activité cérébrale au cours du temps et l'IRM de diffusion permettant de reconstruire le réseau cérébral sont l'objet d'une attention considérable de la part des neuroscientifiques. Dans cette étude, nous présentons une nouvelle méthode pour la détection de schémas d'activité dans le cerveau en combinant ces 2 techniques IRM. Notre approche est basée sur une organisation originale des données sous la forme d'un graphe spatio-temporel des régions cérébrales. Une connexion entre deux régions est créée si ces régions sont reliées anatomiquement et si elles sont fonctionnellement actives entre deux pas de temps successifs. Nous introduisons la notion de composante activée dynamique qui est un petit sous-graphe connecté du graphe spatio-temporel et qui encode la nature dynamique de l'activité du cerveau. En regroupant les composantes d'activation ensemble, nous retrouvons les réseaux d'état de repos (resting state networks) mis en évidence dans la littérature et nous présentons de nouvelles informations sur l'activité cérébrale qui pourraient être d'un grand intéré pour le diagnostic.

Abstract – Functional MRI (fMRI) giving a recording of brain activity over time at the scale of the second, and Diffusion MRI (dMRI) allowing to retrieve the anatomical connectivity network of the brain, have attracted considerable attention from neuroscientists. In this study, we present a novel method for detecting patterns of activity in the brain by combining fMRI and dMRI data. Our approach is based on an original layout of the data into a spatio-temporal graph of interacting brain regions. Edges from one brain region to another are created if those regions are anatomically linked and if they are functionally activated during successive timesteps. We introduce dynamic activation components that are small connected subgraphs of the spatio-temporal graph which encode the dynamical activity of the brain. By clustering similar activation components together, we recover the resting state networks found in the literature and we present new information on the brain activity that could be of high interest for diagnosis purposes.

1 Introduction

Magnetic resonance imaging is one of the most popular techniques to analyze the brain anatomy and function. By collecting information on the brain's activity over time via functional MRI recordings and its connectivity map through diffusion MRI, these imaging data are expected to be extremely useful for the characterization of pathologies such as Alzheimer's disease, multiple sclerosis, epilepsy and psychiatric disorders [2, 7]. Recently, processing the brain's low-frequency activity during task-free experiments allowed to identify spatial patterns of coherent brain activity : called resting state networks (RSNs) [12, 3]. These RSNs have mostly been extracted using spatial independent component analysis (ICA) on the very noisy data obtained by resting state fMRI. While resting state activity schemes are utterly complex and difficult to visualize, these first steps mark an important progress on the understanding of the brain's behavior. It is important to stress that these methods probing resting state activity do not consider non-stationary aspects of the functional time series. Spatial ICA and correlation approaches output average spatio-temporal patterns over the whole recording time, discarding a large amount of information. Besides, they do not take into account the brain structural connectivity network available trough diffusion MRI which could bring additional and complementary information. Indeed, patterns of strong anatomical connectivity between brain areas have been associated to high functional coherence [5, 6].

Stepping from recent literature, it is more and more evident that dynamical aspects of functional activity during rest contain valuable information for the understanding of brain mechanisms in health and disease [8]. In this study we present such a dynamical analysis of functional activity. We successively (i) suggest a way to combine the spatio-temporal fMRI data with the anatomical brain network, (ii) propose a new graph-based technique able to retrieve dynamical interactions between brain regions, and (iii) show how relevant information can be extracted and visualized from dynamic activation components. To validate our approach, we rely on the presence of RSNs. Among the hundreds of activation components, a significant portion of them display strong similarities and are close to one or several well-known activation schemes corresponding to RSNs.

We explore a dataset composed of fMRI and dMRI recordings of 75 healthy volunteers. Each cortex has been divided

into 68 regions with Freesurfer [1] following the anatomic atlas of [4]. The preprocessing of the fMRI and dMRI data is not detailed here. We assume that our dataset is 1) a brain structural network of 68 nodes with unweighted connections and 2) a time series of fMRI activity associated to each node (filtered and averaged from the set of voxels of each brain region). Each patient recording is about 9 min. long at a sampling rate of TR=1.92 s, giving $L = 276$ points for each time series.

One of the key concepts of our graph approach is to arrange the data as a spatio-temporal graph. The idea is to build a graph based on the locality or closeness of data points in space and time. Closeness *in space* is given by the structural brain network, as directly connected regions are more likely to interact with each other. The neighborhood of a node *in time* at time t is itself together with its one-hop neighbors, at the preceding and following time steps $t-1$ and $t+1$. Those nodes are the ones more susceptible to activate/be activated as a consequence of the state of the node considered at time t. This new mapping of the data is a convenient way to handle localization and correlation both on the network and temporal domain. On this graph each node will have only *one* value associated to it, reducing the analysis of the brain activity to the one of a signal on the graph. Fig. 1 (left) illustrates the construction of this graph. This approach is different from the multilayer approach of [10] as connections between layers involve neighbors (different structure and construction) and the analysis concerns repeated patterns within the network (different analysis).

2 The spatio-temporal graph

In practice, we build the spatio-temporal graph in four steps. First, we use the brain network given by the (preprocessed) dMRI data to establish the spatial connections between regions. Let us call this graph G_0. We have chosen to take a unique brain network for the 75 patients, keeping the most common connections among them. Secondly, we have pooled the 75 fMRI signals giving $T = L \times 75$ time steps. We duplicate G_0 T times, one for each time-step. We define the multilayer graph $\widetilde{G}$ of $N \times T$ nodes to be made of the union of copies of $G_0 : \widetilde{G} = \cup_t G_t$, with $t \in [1, T]$. A node in $\widetilde{G}$ will be labeled i_t if it belongs to G_t and is the copy of the node i on G_0. The spatio-temporal graph G is $\widetilde{G}$ where edges linking nodes of different layers $\{G_t\}$ have been added : the third step consists in connecting the T layers together according to the following time-neighborhood rule. For two nodes i_t and j_{t+1} on two successive layers, there is a connection between them if i and j are neighbors in G_0. We may control the connection strength of the time connection by choosing a weight λ to assign to these types of edges. The default will be $\lambda = 1$. This choice will give equal strength to spatial and temporal connections as G_0 is a binary network where weights are either 1 (connected) or 0 (not connected).

The brain activation function f on G is now a one-dimensional function giving the activation rate of each node of G. It has been normalized to zero mean and unitary variance. We assume that the relevant information on the brain activity is given for high values of $|f|$. Remark that negative values of f are also a sign of activity interesting for neuroscientists. In order to find dynamic activation components we introduce a threshold $\rho > 0$. We define the graph A of activated components as follows : a node i_t belongs to A if $|f(i_t)| > \rho$. The dynamic activation components are the weakly connected components (subgraphs) of A. The choice of ρ is empirical. Too high, it leads to a reduced number of small subgraphs, difficult to interpret. Too low, it gives a huge amount of components where a part of them are due to the noise inside the activation signal.

3 Analysis of dynamic activation components

Due to the diversity of activity inside the brain, it is natural to assume the extracted components to be quite different from each other. However, there is evidence that certain activation patterns related to "resting state networks" regularly appearing during the recording of patients at rest [12]. While RSN activation patterns are not expected to repeat in the exact same way over a patient or *a fortiori* between patients, it should be possible to cluster them. Thus, to retrieve RSNs, the clustering of dynamic activation components should be robust to noise and group components together even if they differ slightly in shape. We expect the main clusters (with the largest number of components) to be closely related to the RSNs.

Dynamic activation components in their raw form cannot be directly used as an input of a machine learning algorithm. To cluster them, we introduce two types of feature vectors. The first one, the dynamic-feature vector, encodes the time evolution of the activations and give information on the dynamics. It contains the full information concerning the dynamic arrangement of the component while being memory friendly and computationally efficient. The second one, the static-feature vector, is used to compare more reliably our approach to the previous studies on (static) RSNs ; it is a time-averaged version of the dynamic-feature vector.

Let P_ℓ be an activation component (a connected subgraph of A), where $\ell \in [1, M]$ is the component label and M is the number of components. The graph $P_\ell = (\mathcal{V}_\ell, \mathcal{E}_\ell)$ has a set of vertices $\mathcal{V}_\ell$ and a set of edges $\mathcal{E}_\ell$.

Dynamic features. Let us assume P_ℓ is activated during k_ℓ time steps from t_1 to t_{k_ℓ}. Let $K = \max_\ell k_\ell$ be the maximal duration in time of a component (in number of time-steps). The dynamic-feature vector d_ℓ describing P_ℓ is of length $N \times K$, with binary entries $d_\ell(j) = 1$ when j corresponds to an activated node i at time step k within the component ($j = i \times k$) and 0 otherwise. Notice that most of the vectors are very sparse : in our case, the average number of nodes per component is 9.4 and the length of a feature vector is comprised between 204 (68 nodes times 3 timesteps) and 544 (68 nodes times 8 timesteps).

1. Software available at http ://surfer.nmr.mgh.harvard.edu/

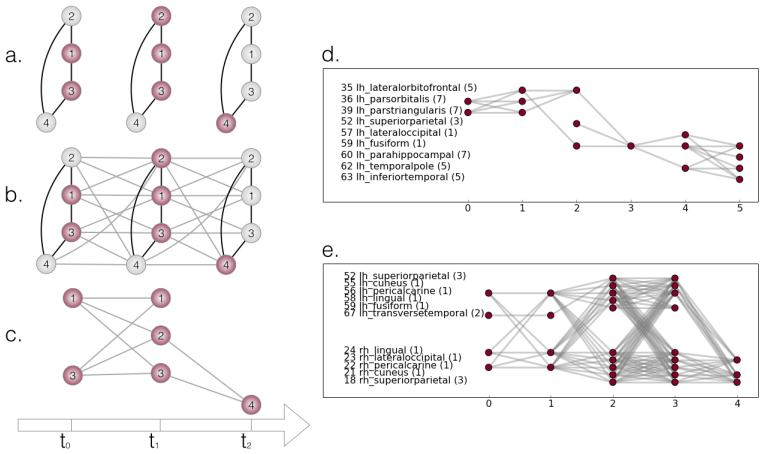

FIGURE 1 – On the left : the successive steps for the creation of the dynamic components. a) The original brain network G_0 is copied for each time step, b) temporal links are added to create the spatio-temporal graph, c) only the nodes with values above the threshold are kept, giving the dynamic activation graph. On the right : two examples of dynamic components of our dataset. The nodes are labeled by their Id (left), name (middle) and RSN (right). The labels lh and rh stand for left and right hemisphere respectively. In the top right figure, the activation moves from one region to another during time (frontal to temporal lobe) but stays in the left hemisphere. In the bottom right figure, the activation spans across several brain regions then reduces before vanishing, mostly involving the visual area. It starts from the peri-calcarine region, which stays active during the whole process.

Static features. To compare to existing RSNs, a time-invariant feature vector is constructed by compressing the component's dynamic into one normalized histogram of brain node occurrence. The static-feature vector $s_\ell \in \mathbb{N}^N$ has the following entries : $s_\ell(j) = \widetilde{s}_\ell(j)/\|\widetilde{s}_\ell\|_2$, where $\widetilde{s}_\ell(j) = \sum_{k=1}^{K} d_\ell(j \times k)$. It is normalized using the ℓ^2-norm to help cluster together components with similar activated brain regions but of different temporal width.

We perform a clustering on the dynamic activation components using the standard k-means algorithm [2] and the static features. We validate our initial assumptions by retrieving the standard RSNs from the obtained clusters. We expect a large number of repeating components (repeated activation of the RSNs), possibly differing by a few nodes (due to noise or other activation behaviors). As a result, in feature space, these slightly altered components should stay close to the one representing a RSN.

To compare our clustering with RSNs, we associate each node of the brain network to a functional resting state network id according to the map given by [12]. There are 7 RSNs : 1) Visual (V), 2) somato-motor (SM), 3) dorsal attention (DA), 4) ventral attention (VA), 5) limbic (L), 6) fronto-parietal (FP) and 7) default mode (DM) networks. Note that our anatomical atlas of 68 nodes does not exactly correspond to the shape of these RSNs given via fMRI data. Moreover, some RSNs are made of small areas scattered all over the brain (dorsal attention and fronto-parietal) which render difficult their description with our atlas. These later two RSNs are made of only 2 and 3 atlas regions respectively, which make them more difficult to detect. As a consequence, some discrepancies appear when comparing our results to the RSNs of the literature.

2. Using the python toolbox scikit-learn available at http ://scikit-learn.org/

Since the number of clusters k has to be set manually, we tested different numbers of clusters (from 7 to 40). Our experiments have shown that the clustering is robust : for each experiment, a large number of clusters could be associated to a single RSN. Moreover, a significant number of clusters stay unaltered from experiment to experiment. We present the results of a clustering with $k = 12$ on Fig. 2. The clusters are in good correspondence with the RSNs. The histograms shows that cluster 11 is associated to the visual RSN, cluster 4 to the somatomotor RSN, cluster 7 to ventral attention and cluster 2 to the default mode. Only a few of the clusters contain a mixture of RSNs.

A closer analysis of the clusters shows two categories of components : symmetric ones with respect to the left and right hemispheres such as cluster 2, 4, 5, 6, 7 and 11 and non-symmetric (unilateral) ones which can be grouped by pairs such as 0-1, 8-9 and 3-10. The coexistence of symmetric/dissymmetric components is similar to what [9] obtains by directly clustering at the voxel level (with 11 clusters). The pair 3-10 corresponds to a fronto-parietal activation, which can be interpreted as a part of the dorsal attention network (known to possess a lateralized behavior) even if some regions are labeled as default mode network. The clusters 0-1 (limbic RSN) and 8-9 contain regions of the temporal lobe while being dissymmetric. Some components appear to be part of more than one resting state : this is what is also obtained in [11] where they performed a temporally-independent component analysis ; for example, some of their dynamic components contain part of the default mode as well as other resting states (see their mode TFM 8, a symmetric version of the mean component of cluster 9). It is also the case in [1] where a sliding window ICA is used to retrieve dynamic patterns.

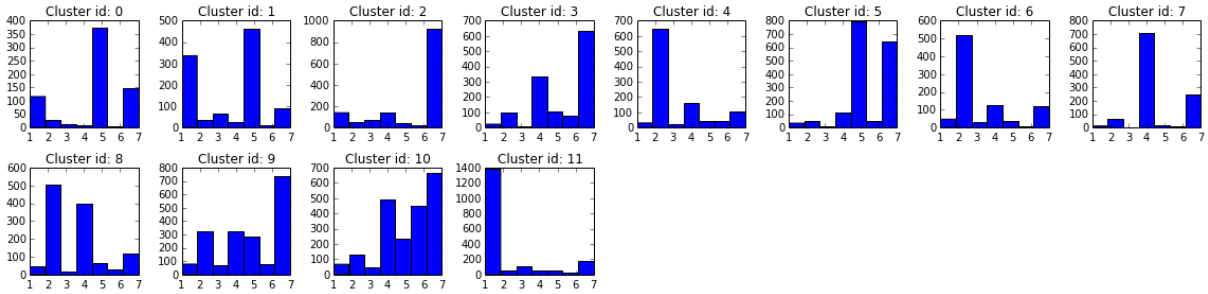

FIGURE 2 – Clustering with $k = 12$. Distribution of the components' nodes among the 7 RSNs for each cluster.

This demonstrates that our method is able to retrieve known RSNs while giving new insights on their dynamics. This is even more convincing when plotting the average component for some of the clusters with $k = 12$ on Fig. 3, 4 and Fig. 1.

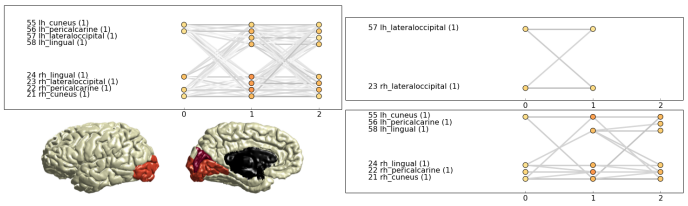

FIGURE 3 – Mean component of cluster 11, corresponding to the visual RSN. It is symmetric between the left and right hemispheres and contains the peri-calcarine, the cuneus, the lingual and the lateral occipital regions. The node colors correspond to the amount of activation for each node, the scale ranges from white to dark red (for the nodes most present in the components). White (low) to dark grey (high) indicates the importance of edges between nodes at successive timesteps. In the right part of the figure, the mean component is shown where only the strongest edges are kept. It is cut in two parts, corresponding to the primary and extrastriate visual cortex.

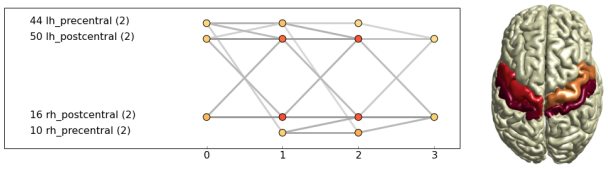

FIGURE 4 – Mean component of cluster 4, the somato-motor RSN. Both left and right pre-central and post-central regions are active, the post-central region being activated more often among the components belonging to the cluster. The relatively similar edge color indicates that all the regions equally interact with each other.

4 Conclusion

Our approach based on a spatio-temporal graph describing brain activation patterns recovers and confirms the existence of RSNs. It also reveals the dynamical interactions between the main anatomical regions of the brain and precises the idea that RSN are made of smaller blocks, interacting dynamically. The activation patterns we obtain show interesting dynamic activity with spreading and propagation over the cortex. This work paves the way for the creation of new tools to further comprehend the brain mechanism and improve diagnosis.

Références

[1] E. A. Allen and et al. Tracking Whole-Brain Connectivity Dynamics in the Resting State. *Cer. Cortex*, 24(3) :663–676, Mar 2014.

[2] F. X. Castellanos and et al. Clinical applications of the functional connectome. *NeuroIm.*, 80(0) :527 – 540, 2013. Mapping the Connectome.

[3] D. M. Cole, S. M. Smith, and Ch. F. Beckmann. Advances and pitfalls in the analysis and interpretation of resting-state FMRI data. *Front. Syst. Neurosci.*, 2010.

[4] R. S. Desikan and et al. An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *NeuroIm.*, 31(3) :968–980, 2006.

[5] J. Goni and et al. Resting-brain functional connectivity predicted by analytic measures of network communication. *PNAS*, 111(2) :833–838, 2014.

[6] M. D. Greicius and et al. Resting-state functional connectivity reflects structural connectivity in the default mode network. *Cer. Cortex*, 19(1) :72–78, 2009.

[7] A. Griffa and et al. Structural connectomics in brain diseases. *NeuroIm.*, 80(0) :515 – 526, 2013. Mapping the Connectome.

[8] R. M. Hutchison and et al. Dynamic functional connectivity : Promise, issues, and interpretations. *NeuroIm.*, 80(0) :360–378, 2013. Mapping the Connectome.

[9] M. H. Lee and et al. Clustering of Resting State Networks. *PLoS ONE*, 7(7) :e40370, Jul 2012.

[10] Peter J Mucha, Thomas Richardson, Kevin Macon, Mason A Porter, and Jukka-Pekka Onnela. Community structure in time-dependent, multiscale, and multiplex networks. *science*, 328(5980) :876–878, 2010.

[11] S. M. Smith and et al. Temporally-independent functional modes of spontaneous brain activity. *PNAS*, 109(8) :3131–3136, 2012.

[12] B.T. Thomas Yeo and et al. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J. Neurophys.*, 106(3) :1125–1165, Sep 2011.