Vision, Modeling, and Visualization (2011)
Peter Eisert, Konrad Polthier, and Joachim Hornegger (Eds.)

Data-Driven Visualization of Functional Brain Regions from Resting State fMRI Data

Alessandro Crippa[†]

Jos B.T.M. Roerdink[‡]

Johann Bernoulli Institute for Mathematics and Computer Science, University of Groningen, The Netherlands

Abstract

Functional parcellation of the human cortex plays an important role in the understanding of brain functions. Traditionally, functional areas are defined according to anatomical landmarks. Recently, new techniques were proposed that do not require a priori *segmentation of the cortex. Such methods allow functional parcellation by functional information alone. We propose here a data-driven approach for the exploration of functional connectivity of the cortex. The method extends a known parcellation method, used in multichannel EEG analysis, to define and extract functional units (FUs), i.e., spatially connected brain regions that record highly correlated fMRI signals. We apply the method to the study of fMRI data and provide a visualization, inspired by the EEG case, that uses linked views to facilitate the understanding of both the location and the functional similarity of brain regions. Initial feedback on our approach was received from four domain experts, researchers in the field of neuroscience.*

Categories and Subject Descriptors (according to ACM CCS): Data [E.1]: Graphs and networks—; Life and Medical Sciences [J.3]: Health—

1. Introduction

Parcellation of the gray matter of the human cortex into functionally distinct areas is a key aspect of neuroscience. In the standard approach [FvdKD*04, GGS*99], anatomical landmarks such as sulci and gyri are used to delineate boundaries between areas of interest for the interpretation of functional neuroanatomy, with the default assumption that brain structure reflects functional specialization. Nevertheless, several studies showed that there is a limited correspondence between anatomical boundaries and functional specialization, and often cytoarchitectonic analysis is used as a means to access functional parcellation (e.g., [ASB*99]). Another criterion for the identification of functionally different brain areas is the connectivity with other brain regions [JB76]. A recently developed branch of research uses diffusion imaging to perform parcellation of functional brain regions under the assumption that differences in connectivity patterns should play an important role in identifying functionally different cortical areas [BJBW*03, JBBR*04, CCNR10].

Several approaches have been proposed in the literature for performing gray matter parcellation solely based on functional similarity, i.e., comparing functional activation of different brain regions and studying their interconnections. The most insightful method for such studies consists of analyzing the signals from single brain cells, as often performed in animal studies [MIT91]. This method has a very high spatial resolution, but its application to the simultaneous analysis of several brain regions is difficult. To analyze the whole brain, functional magnetic resonance imaging (fMRI) is the method of choice. This technique has a much lower spatial resolution but allows us to explore the whole brain at once. fMRI analysis detects variations in the BOLD signal, which is an indirect indicator of neurological activity, and is usually based on the study of brain responses to certain stimuli [Log08]. A functional similarity between brain regions is assumed when these regions show a similar response to the same stimulus. This assumption is the basis of functional comparison of brain regions.

† e-mail: a.crippa@rug.nl
‡ e-mail: j.b.t.m.roerdink@rug.nl

Resting state fMRI, introduced by Biswal *et al.* [BYHH95], is the study of brain activity during cognitive rest [SFM*09]. Connectivity studies in resting state conditions revealed that the functional network of the human brain shows a complex structure even when no brain exercise is performed, and shows stable patterns of activity [Rai06, RM06]. Variations in BOLD signals that reflect spontaneous physiological activity have been observed, usually below the frequency of 0.1 Hz, in resting state functional networks. Resting state fMRI has recently provided prospects for the understanding of the functional brain.

In this paper, we present a functional parcellation technique that is based on a visualization approach previously developed for EEG data. The original method, introduced by ten Caat *et al.* [tCMR08], aims for the clustering of multichannel EEG networks and has its main strength in the fact that it is totally data driven. Data-driven approaches have the advantage that they do not require any *a priori* hypotheses about the effects to be expected. Therefore, they are particularly suitable in exploratory phases of research. Also, they can better deal with individual variations, since no assumptions on the location of functional brain areas are made. We extend the method of [tCMR08] to the fMRI domain. Currently, our method is limited to fMRI activation patterns in the cortex, i.e., subcortical nuclei are not included in the analysis. The method provides the possibility to locate FUs and investigate inter-FU similarities. Compared to the EEG case, major computational and visualization challenges arise due to the fact that much larger graphs have to be processed. Finally, we report on an informal evaluation of the results with medical domain experts.

2. Related work

Seed-based methods have been the first approach to functional connectivity analysis. Such methods investigate the network of brain regions whose activation correlates with that of a given region of interest (ROI). Although successful, seed-based methods have been shown to provide significant variability with respect to the choice of the ROIs, due to the lack of standard ROI selection procedures [CHA*00, CHA*01, CHC*02].

To avoid this experimenter bias, several clustering methods have been applied to automatically identify ROIs; examples include hierarchical clustering [CHC*02, FBM99], unsupervised segmentation [GGM07], multivariate analysis [TDP06], and others. Salvador *et al.* [SSC*05] use a brain parcellation based on a prior anatomical template, dividing each cerebral hemisphere into 45 anatomical ROIs.

So far, independent component analysis (ICA) [KKJ*03] of resting state data seems to be the most successful technique for localizing functional regions. ICA has the advantage to be data driven, so neither ROI selection nor an *a priori* defined anatomical template is required. ICA is a method for separating a signal into its independent components, and it has been successfully applied to resting-state as well as task-related fMRI analysis. ICA gained consensus in the neuroscience community and is nowadays the most common technique for the analysis of resting state fMRI. An interesting result achieved by ICA is the ability to extract the so called "default-mode network", a functional set of brain regions that are active when a subject is not focused on performing tasks. Note, however, that in [KKJ*03] there is no concept of links between regions, hence the method is not directly comparable to the graph-based technique proposed in this paper.

3. Methods

As our method is based on previous work on extracting FU maps for EEG data, we first give an outline of that approach.

3.1. Functional Unit Maps for EEG recordings

Functional unit (FU) maps have been introduced by ten Caat *et al.* [tCMR08] as a multiscale approach to visualize both local and global similarities in the electrode signals recorded during an EEG experiment.

The method is based on the analysis of a graph, called "coherence graph", where nodes represent EEG electrodes and edges connecting nodes are weighted by the coherence (correlation at a given frequency band) of the signals recorded at their nodes. The coherence graph is pruned and only statistically significant edges are kept. FUs are computed as maximal cliques in the pruned coherence graph, i.e., as spatially connected sets of nodes that correspond to electrodes registering similar signals. The reason to define FUs as sets of spatially connected nodes is based on the assumption that EEG electrodes which are spatially close usually record signals from a single source as the result of volume conduction [LRMV99].

For the computation of the FUs and for the visualization of the results, a notion of neighborhood among electrodes is necessary along with the coherence of the registered signals. Electrode positions are thus projected from the scalp to a horizontal 2D plane, and neighborhoods between electrodes are computed on this plane using a Voronoi tessellation. FUs are computed by an algorithm based on the watershed method [BM93], whose time complexity is $O(n^2 \log(n))$, where n is the number of electrodes.

The visualization of the FUs and their clustering in an EEG recording is called a *FU map*. An example of such a FU map is shown in Figure 1. A FU map is the 2D representation of the scalp, where cells, representing electrodes, are colored according to the FU they belong to. Four-coloration of the FUs is used, so that two contiguous FUs always have different colors.

A FU map shows two kinds of information, both at local

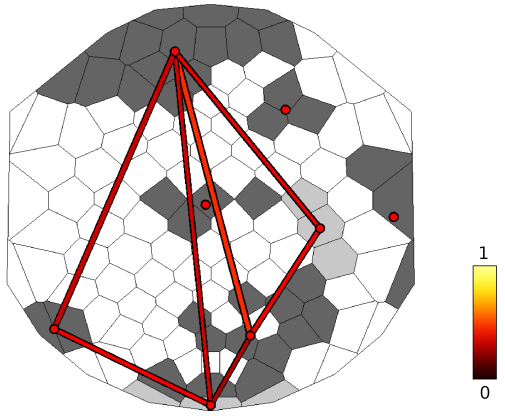

Figure 1: *FU map for a multichannel EEG recording [tCMR08]. Voronoi cells represent EEG electrodes, colored according to the FU they belong to (four-coloration of the FUs is used). White Voronoi cells do not belong to any FU. Circles represent the barycenters of the FUs, and two circles are connected by a link if the average coherence between the corresponding FUs is significant.*

and global scale. At a local scale, it shows which neighboring electrodes register similar signals and form the FUs. At a global scale, it shows the similarity between the FUs, defined as the coherence between the average signals recorded in two FUs. Similarity between FUs is visualized by using links connecting the barycenters of the FUs, color-coded according to the strength of their similarity. Links are only shown if the average coherence between the corresponding FUs is significant.

Notice that because of the graph pruning, cells exist that do not belong to any FU. These are colored in white. Please refer to [tCMR08] for a complete description of the clustering method and its application to multichannel EEG.

3.2. FU in fMRI

We provide here first a summary of the steps needed to compute a FU map using functional MRI data. We will describe the details of each step in the next subsections.

The algorithm creates a graph where each node represents a fMRI time series, recorded at a certain position on the brain cortex, and edges are weighted according to the correlations between the time series. We call this a "correlation graph". The correlation graph corresponds to the coherence graph in the EEG case. FUs are computed on the correlation graph by using the same method as in the EEG case. Results are displayed in two ways. First, the FUs are mapped on a 3D cortical mesh, as this is the representation most familiar to neuroscientists, which are the potential end users of the method. However, drawing edges between FUs on the brain surface would lead to a very cluttered display, due to the very large number of FUs present (note that the spatial resolution of fMRI data is much higher than the spatial resolution of EEG data). Therefore, to ease the exploration of the locations of the FUs, we draw the edges between FUs in a 2D visualization, similar to that of the EEG case, where FUs registering highly similar signals are connected by edges. The 2D and 3D visualizations are presented in a linked view.

The steps of our algorithm for computing FU maps of fMRI recordings are as follows.

1. Downsample an available high-resolution (HR) mesh of the brain cortex to obtain a low-resolution (LR) mesh. Then map the fMRI time series on the LR surface mesh.
2. On the LR mesh:
   a. Create the correlation graph;
   b. Compute the fMRI FU map;
   c. Label the FUs by using four-coloration;
   d. Map the colored FUs on the cortical mesh.
3. Map the labels of the LR mesh back to the HR mesh.
4. Map the colored FUs to a flattened 2D version of the HR cortical mesh. Draw links between FUs in which, on average, highly correlated (or anti-correlated) fMRI signals were recorded.
5. Display the 3D cortex mesh and the flattened 2D mesh in a linked visualization.

We now discuss these steps in more detail.

3.3. Mapping of the fMRI data on the cortex mesh

The first step of the algorithm consists of mapping the fMRI signal to the nodes of two cortical meshes, one per hemisphere. We use the cortical meshes provided by CARET, an application for functional and structural analysis of the brain cortex [VEDD*01]. There are two reasons to use hemicortical meshes instead of one single mesh for the whole brain. First, CARET provides both 3D and 2D versions of the hemicortical meshes. The 2D version is helpful to visualize all the cortical surface, as the brain cortex is highly folded. CARET is a well known application and its cortical meshes are extensively used in the neuroscience community. Second, using two meshes can speed up computations in case the investigator is only interested in the analysis of a single hemisphere.

The meshes available in CARET represent population-averaged [VE05] hemicortical surface reconstructions in standard MNI space (see section 4.1 for an explanation of MNI space). The meshes consist of $15 \cdot 10^4$ nodes in total, the average distance between two nodes being 1 mm (the value of the mean distance between nodes is computed by CARET, no information of the standard deviation is available).

Voxels in fMRI data typically have a resolution of 5 cubic mm. The cortical meshes provided by CARET are therefore

downsampled to obtain low resolution meshes with $5 \cdot 10^3$ nodes and an average distance between nodes of 4.8 mm. The downsampling allows a better match between the resolution of the mesh and the resolution of the fMRI data, and eases the computation of the FUs since the time requirements of our algorithm increase as $O(n^2 \log(n))$, with n the number of mesh nodes (see section 3.1). The downsampling is performed using a tool available in CARET. It is important to note that this downsampling does not affect the topology of the mesh. Note that the CARET mesh has a lot of detail and it precisely follows the cortical folds (sulci and gyri). Also, during downsampling the folds have a higher sampling rate than flat areas, so that more precision is obtained in these critical regions.

Functional MRI data are mapped on the downsampled hemicortical surfaces by using the Gaussian mapping implemented in CARET. This algorithm uses a spatial Gaussian kernel to assign an fMRI time series to each node of the mesh, which is the average time series of the voxels in a neighborhood of the mesh node. We used a Gaussian kernel of 10mm.

Coregistration of single subject's data to the MNI standard template and spatial Gaussian smoothing of the data are both fundamental steps in the standard fMRI analysis, to be performed before any analysis on the data. The registration steps allow the comparison of multiple datasets and the analysis of results in a standard space. Gaussian smoothing reduces the signal-to-noise ratio in single voxel time series, and it is a critical step for the comparison of results among subjects. Nevertheless, since both functions and structures are unique in each subject's data, these operations are approximate and could reduce the accuracy of the subsequent analysis steps. There is no standard way for performing coregistration and smoothing, and different analysis programs use different techniques and/or parameters. A Gaussian kernel of ca. 10mm is, on average, the standard operator for spatial smoothing in fMRI. An alternative to Gaussian smoothing is wavelet denoising [WR04].

3.4. Creation of the correlation graph

In this step we create two graphs, one per hemisphere, in which the vertices represent the nodes in the downsampled cortical meshes and the weight of an edge equals the correlation between the fMRI signals of the nodes of the edge. Edges whose correlation is not statistically significant ($p > 0.05$) are pruned. Cross correlations of fMRI time series, and p-values associated with the correlations, are computed in Matlab©. Furthermore, edges whose correlation is lower, in absolute value, than 0.975 (95% of the range in which the correlation is defined) are also pruned. The pruning is necessary to be able to create results consistent with anatomical and functional evidence. Other threshold values were tested and their results are discussed in section 5.

The main difference in the creation of the correlation graph, with respect to the creation of the coherence graph in the EEG scenario, is that correlation can assume negative values. The choice of using correlation was driven by the fact that it is possible to register negative correlated (anti-correlated) fMRI signals at two different brain regions. The computation of the FUs in the fMRI domain is performed solely using positive correlations, since the aim of the method is to locate brain regions that work in synchrony. But similarities between FUs are computed by using either positive or negative correlation, with the aim to locate which FUs registered highly correlated, or highly anti-correlated, average signals.

3.5. Computation of the fMRI FU map

In this step we compute the FUs, still on the LR meshes. To do this we need both the correlation graphs and a notion of neighborhood for the graph nodes. Neighborhood is implicitly defined in the cortical meshes: two graph nodes are neighbors if the corresponding mesh nodes are vertices of the same polygon (note that CARET provides triangular meshes). The computation of the FUs is performed in Matlab, and takes approximately 15 minutes in total for both hemispheres.

3.6. Four-coloration of the FU map

FUs are labelled using a four-coloration of the FU map of each hemisphere, so that neighboring FUs are labelled differently. Four different gray values are used as labels, and are assigned to the nodes belonging to the FUs so that all mesh nodes in a FU are colored with the same gray value.

Like in the EEG scenario, there could be nodes that do not belong to any FU. Although cells unassigned to FUs are colored in white in the EEG case, in the fMRI case we use black to color such unassigned mesh nodes. The reason for this choice is that we will eventually use Phong lighting in the visualization of the 3D cortical meshes, and the detection of anatomical landmarks such as sulci and gyri (both used to provide context and ease the detection of the FUs) could be difficult in black regions.

To improve the visualization of the results, the next step is to map the FU labels to the original high resolution cortical meshes.

3.7. High-resolution FU map

In this step we map the labels of the mesh nodes from the low resolution meshes back to the original high resolution hemicortical meshes provided by CARET. This operation is performed in two steps. The first step is a "one to one" mapping from each node in the low-resolution mesh to its closest node in the high-resolution mesh. Step two is a labeling of the remaining nodes in the high resolution mesh. This step is performed by using a watershed method: the

flooding algorithm uses each node which was labeled in the first step as a basin marker and propagates the labels along the mesh surface [tCMR08]. The mapping of the FUs to the high-resolution map is not strictly necessary, but it produces a visualization that the final users are more comfortable with and in which the locations of cortical regions are more easily accessible.

3.8. Creation of the 2D map with FU edges

As for the EEG scenario, an important feature of our visualization is the possibility to investigate inter-FUs similarities. This goal is achieved by drawing edges that connect those FUs that registered highly correlated average signals. We draw the edges only in a 2D view, because drawing them in the 3D view would lead to a cluttered display. The 2D representation of the cortical meshes is a flattened version of the 3D meshes produced by CARET. An edge between two FUs is drawn when the correlation of the average signal in the nodes of the first FU and the average signal in the nodes of the second FU is bigger than the threshold used for pruning the correlation graph.

Edges are colored according to the value of the correlation using two colormaps: one for the positive correlations and one for the negative ones. The vertices of the edges between FUs are the barycenters of the FUs, which are visualized by small red circles that serve as identifiers for the FUs (cf. Figure 4). If the barycenter falls outside the FU, which may happen for FUs with a non-convex shape, the FU is represented by a cell of the FU that is nearest to the position of the barycenter. In this case, edges connecting that FU to other FUs leave from that cell.

The visualization of the flattened version of the meshes is paired with a linked visualization of the 3D cortical meshes, and interaction between the two visualizations is possible. The user can interactively check which areas in the flattened cortical meshes correspond to which areas in the 3D meshes by means of a pointer that shows the correspondence.

4. Experimental results

4.1. Data acquisition and preprocessing

A resting-state fMRI dataset (single subject, healthy, male, 30 years old) was selected from an experiment in which the participants were requested to rest with their eyes open in the MRI scanner and fixate a crosshair projected on the monitor. The data were registered to the MNI brain [BJO02] (the standard average brain used as a reference for comparing brain scans of different subjects). A low-pass filter was subsequently applied to remove fMRI signals with frequency higher than 0.1 Hz. Both registration and filtering are standard preprocessing steps in fMRI resting state analysis, and were performed using FSL [SJW*04], a library of analysis tools for brain imaging data.

4.2. Results

The first step of the algorithm is illustrated in Figure 2, which shows the differences between the original cortex mesh (on the left) and the downsampled mesh (on the right).

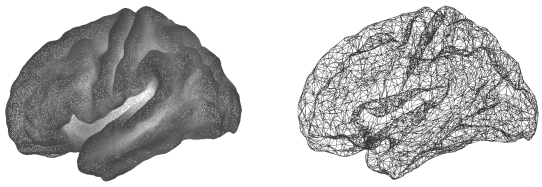

Figure 2: *The original CARET mesh (left) for the left hemisphere, and the corresponding downsampled version used to map the fMRI signal (right). The mesh on the left has* $7.5 \cdot 10^4$ *nodes, the mesh on the right has has* $2.5 \cdot 10^3$ *nodes. Edges between mesh nodes are represented by segments whose greyscale values vary according to the sulcal depth.*

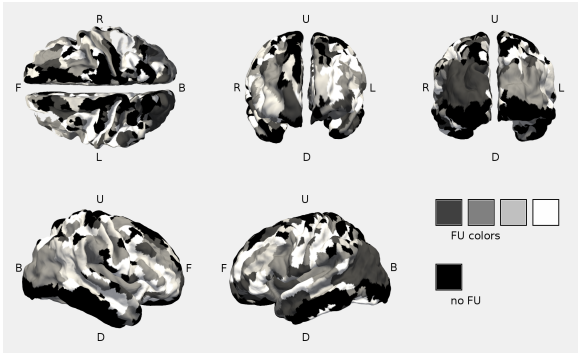

Figure 3: *High-resolution FU maps, seen from different viewpoints. Top row, from left to right: top view, front view, back view; bottom row, from left to right: right view, left view, color labels (black is used for nodes that do not belong to any FU). Mapping the FU labels from the low-resolution meshes to the original high resolution meshes allows the user to better localize the FUs in the anatomy. The high resolution cortical meshes are colored with the FU labels and shaded using Phong lighting on the GPU.*

Figure 3 shows an example of a high-resolution FU map seen from different viewpoints. Figure 4 shows the linked view of the 2D and 3D representations of the FU maps. On top, the meshes are visualized in 3D and shaded using Phong lighting for optimal quality. FUs, sulci and gyri are clearly visible, as the user is able to change the point of view. It is also possible to visualize one single hemisphere at a time and examine the FUs in the interhemispheric fissure (the gap between two hemispheres). The bottom part of the figure shows the two (individually) flattened brain hemispheres and the

edges between the FUs. Colormaps for positive and negative correlations are displayed on the right. Note that for this dataset all the edges carry positive correlation values, since after thresholding no negative correlations between signals registered in the FUs remained.

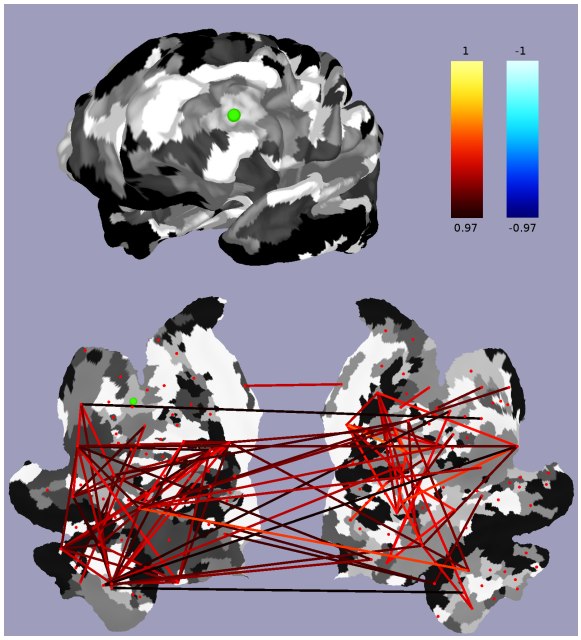

Figure 4: *Linked visualizations of the FU map. Top: a 3D rendering of the cortical mesh, shaded using Phong lighting to enhance both the FUs and the anatomical landmarks consisting of sulci and gyri, which are recognizable by neuroscientists. Bottom: the flattened meshes and the edges between FUs that recorded highly correlated fMRI signals. FUs are represented by gray values using the same colormap as in figure 3, edges between FUs are colored according to the shown colormaps for positive and negative correlations. In this dataset, no significant negative correlations between FUs were found. The green spheres represent the same spot in the two meshes.*

5. Feedback from domain experts

To evaluate the potential of our method we conducted interviews with four medical domain experts, researchers in the field of neuroscience. One of the participants was familiar with the visualization of FU maps in the EEG domain. The aim of the study was the validation of the results via comparison of the location of the FUs to the location of known functional or cytoarchitectonic brain regions, and the validation of the found similarities between FUs. Evaluation of the visualization and its comparison with the visualization of FU maps in the EEG case was also part of the interview.

As a general comment, the domain experts found the results interesting, and they reported that the results were "reasonably symmetrical in the two hemispheres" and that "several FUs follow anatomical landmarks". They also reported that "several FUs reflect actual functional regions", but "there are also brain regions where the results diverge from the current functional evidence".

The experts commented on the consistency of the results with functional and cytoarchitectonic evidence mainly for the following regions. The subdivision in small FUs in the posterior parietal lobe, especially in the left hemisphere, appeared to be consistent both with the functional and cytoarchitectonic differentiation of this region. In the inferior frontal gyrus, the FU map detects a tripartition of Broca's area (Brodmann's areas 44, 45, 47) and, consistent with functional evidence, the tripartition is more evident in the left hemisphere (these areas are located around the green dot in figure 4, top). A tripartition appears also in the Insula (the grey region in figure 4, bottom right, from which five FU edges leave), consistent with functional, cytoarchitectonic, and connectivity evidence. In the occipital lobe, the FUs locate the secondary visual areas consistently with its cytoarchitectonic properties, but the location of the FUs in the medial aspects of the occipital lobe is not fully consistent with functional evidence, especially in the right hemisphere. The subdivision in anterior and posterior regions of the cingulum is also interesting, although this only occurs in the right hemisphere and the FUs in the interhemispheric fissure generally disagree with functional evidence. The middle temporal area is identified in both hemispheres, especially in the left hemisphere. The subdivision in dorsal and ventral parts of the premotor area was also found to be consistent with functional results, although a more defined differentiation along the central sulcus was expected.

The experts stressed the fact that a task-related fMRI analysis would be necessary to completely assess the benefits of the method, since only with such an analysis it is possible to reliably locate functional regions in a single subject analysis. One of the experts suggested that the lack of FUs in the motor areas could be caused by the fact that no motor task was performed. There is apparently no explanation for the lack of FUs in the ventral part of the occipital and temporal lobes. Another expert proposed that once the real benefits of the method are assessed, a possible application could be a pre-surgery assessment of functional areas when task-related fMRI is not possible, e.g., in the case of coma.

The domain experts were more cautious when commenting on the links between the FUs. They agreed that the information provided by these links could be used as an exploratory tool, since it is known that the brain is dynamically active even during resting state conditions, but the extent to which brain regions interact during resting state is still not clear. One of the experts positively commented on the edge connecting the two FUs located on the left and right part of the cingulum, which reflects the structural connectivity of

this region. Another one commented on the relatively high number of edges leaving from the Insula in the right hemisphere, suggesting that it could reflect the high structural connectivity this region is known to have with several brain areas. The experts were also expecting to find links between the medial prefrontal cortex, the posterior cingulum and precuneus, the inferior parietal cortex and the medial temporal lobe. This network represents the "default mode network", as detected by ICA [KKJ*03] and other methods in resting state analysis. Our algorithm does not find this "network". As the ICA-based approach of [KKJ*03] is not directly comparable to our method, further analysis and consideration of more subjects would be necessary to understand the differences between the outcomes of the two methods.

The experts were satisfied with the possibility, provided by the linked visualization, to compare the 2D and the 3D maps, and all of them agreed that the linked view is essential to provide an intuitive anatomical context. The participant familiar with the visualization of FU maps in the EEG domain found the 2D visualization very intuitive.

A current limitation of the method is the need of setting a threshold for the pruning of the correlation graph. We set this threshold to 95% of the range of correlation, but noticed that changes in the threshold can produce different results. We asked the medical experts to explore FU maps computed using different thresholds. Bigger and fewer FUs are found if the threshold is set to lower values (e.g., lower than 0.95), smaller and more numerous ones are found if the threshold is set higher (e.g., bigger than 0.98). In both cases, the domain experts thought that the FU map did not reflect the functional or cytoarchitectonic evidence. They reported that only small variations occurred when varying threshold values in the range [0.95-0.98].

6. Discussion

In this paper we proposed a method, building upon a method previously used in multichannel EEG analysis, that allows the visualization of brain functional connectivity from resting state fMRI data. Since only one subject was used in the experimental part, our method should be considered a "proof of principle " at this stage.

We evaluated the potential of our method by performing interviews with four medical domain experts, whose main interest was the possibility to localize functional and cytoarchitectonic regions by means of resting state analysis alone, i.e., without any focused task and without a priori hypothesis. They were less interested in the possibility to investigate functional similarities between functional units, since they could not compare the results to a ground truth, which currently is not available. A relation to the "default mode network" obtained by the ICA method could not be established.

As a general conclusion, the domain experts agreed that this method could provide insights in the localization of functional and cytoarchitectonic brain areas via resting state fMRI. Because the precise location of functional areas changes among subjects, a task-related fMRI study could provide further validation of the technique by allowing a better comparison between the location of the FUs and the exact locations of functional brain areas.

Several opportunities for future work exist. First, a continuation of this line of research could be the quantitative comparison of FU maps, using the method proposed in [CMR10]. This method allows the comparison of FU maps among several subjects by detecting both the average functional behavior in the group of subjects and the individual differences. Second, a reimplementation of the algorithm for computing the FUs, e.g., in C/C++, would result in a substantial speedup. Another possibility for future work is to study visualization of functional brain regions by FU maps computed for *joint* EEG/fMRI data. Such simultaneous measurements have become feasible in the last decade. Another possible improvement could consist of implementing interaction facilities in the linked visualization. These could give emphasis to certain aspects of the visualization, such as brain regions, graph edges, connected graph components, or symmetry across the hemispheres. A user study involving neuroscientists and medical domain experts would be necessary to identify which key features would improve our visualization tool.

References

[ASB*99] AMUNTS K., SCHLEICHER A., BURGEL U., MOHLBERG H., UYLINGS H. B., ZILLES K.: Broca's region revisited: cytoarchitecture and intersubject variability. *J. Comp. Neurol. 412* (Sep 1999), 319–341. 1

[BJBW*03] BEHRENS T. E., JOHANSEN-BERG H., WOOLRICH M. W., SMITH S. M., WHEELER-KINGSHOTT C. A., BOULBY P. A., BARKER G. J., SILLERY E. L., SHEEHAN K., CICCARELLI O., THOMPSON A. J., BRADY J. M., MATTHEWS P. M.: Non-invasive mapping of connections between human thalamus and cortex using diffusion imaging. *Nat. Neurosci. 6* (Jul 2003), 750–757. 1

[BJO02] BRETT M., JOHNSRUDE I. S., OWEN A. M.: The problem of functional localization in the human brain. *Nature Reviews Neuroscience 3* (2002), 243–249. 5

[BM93] BEUCHER S., MEYER F.: The morphological approach to segmentation: the watershed transformation. In *Mathematical Morphology in Image Processing*, Doherty E., (Ed.). Marcel Dekker, New York, 1993, ch. 12, pp. 433–481. 2

[BYHH95] BISWAL B., YETKIN F. Z., HAUGHTON V. M., HYDE J. S.: Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn Reson Med 34* (Oct 1995), 537–541. 2

[CCNR10] CRIPPA A., CERLIANI L., NANETTI L., ROERDINK J. B. T. M.: Heuristics for connectivity-based parcellation of SMA/preSMA through Force-Deirected Graph Layout. *Neuroimage 1* (2010), 1. 1

[CHA*00] CORDES D., HAUGHTON V. M., ARFANAKIS K., WENDT G. J., TURSKI P. A., MORITZ C. H., QUIGLEY M. A., MEYERAND M. E.: Mapping functionally related regions of

brain with functional connectivity MR imaging. *AJNR Am J Neuroradiol 21* (Oct 2000), 1636–1644. 2

[CHA*01] CORDES D., HAUGHTON V. M., ARFANAKIS K., CAREW J. D., TURSKI P. A., MORITZ C. H., QUIGLEY M. A., MEYERAND M. E.: Frequencies contributing to functional connectivity in the cerebral cortex in "resting-state" data. *AJNR Am J Neuroradiol 22* (Aug 2001), 1326–1333. 2

[CHC*02] CORDES D., HAUGHTON V., CAREW J. D., ARFANAKIS K., MARAVILLA K.: Hierarchical clustering to measure connectivity in fMRI resting-state data. *Magn Reson Imaging 20* (May 2002), 305–317. 2

[CMR10] CRIPPA A., MAURITS N. M., ROERDINK J. B. T. M.: Graph Averaging as a Means to Compare Multichannel EEG Coherence Networks. In *Eurographics Workshop on Visual Computing for Biology and Medicine* (Leipzig, Germany, 2010), Bartz D., Botha C., Hornegger J., Machiraju R., Wiebel A., Preim B., (Eds.), Eurographics Association, pp. 33–40. 7

[FBM99] FILZMOSER P., BAUMGARTNER R., MOSER E.: A hierarchical clustering method for analyzing functional MR images. *Magn Reson Imaging 17* (Jul 1999), 817–826. 2

[FvdKD*04] FISCHL B., VAN DER KOUWE A., DESTRIEUX C., HALGREN E., SEGONNE F., SALAT D. H., BUSA E., SEIDMAN L. J., GOLDSTEIN J., KENNEDY D., CAVINESS V., MAKRIS N., ROSEN B., DALE A. M.: Automatically parcellating the human cerebral cortex. *Cereb. Cortex 14* (Jan 2004), 11–22. 1

[GGM07] GOLLAND P., GOLLAND Y., MALACH R.: Detection of spatial activation patterns as unsupervised segmentation of fMRI data. *Med Image Comput Comput Assist Interv 10* (2007), 110–118. 2

[GGS*99] GOLDSTEIN J. M., GOODMAN J. M., SEIDMAN L. J., KENNEDY D. N., MAKRIS N., LEE H., TOURVILLE J., CAVINESS V. S., FARAONE S. V., TSUANG M. T.: Cortical abnormalities in schizophrenia identified by structural magnetic resonance imaging. *Arch. Gen. Psychiatry 56* (Jun 1999), 537–547. 1

[JB76] JONES E. G., BURTON H.: Areal differences in the laminar distribution of thalamic afferents in cortical fields of the insular, parietal and temporal regions of primates. *J. Comp. Neurol. 168* (Jul 1976), 197–247. 1

[JBBR*04] JOHANSEN-BERG H., BEHRENS T. E., ROBSON M. D., DROBNJAK I., RUSHWORTH M. F., BRADY J. M., SMITH S. M., HIGHAM D. J., MATTHEWS P. M.: Changes in connectivity profiles define functionally distinct regions in human medial frontal cortex. *Proc. Natl. Acad. Sci. U.S.A. 101* (Sep 2004), 13335–13340. 1

[KKJ*03] KIVINIEMI V., KANTOLA J. H., JAUHIAINEN J., HYVARINEN A., TERVONEN O.: Independent component analysis of nondeterministic fMRI signal sources. *Neuroimage 19* (Jun 2003), 253–260. 2, 7

[Log08] LOGOTHETIS N. K.: What we can do and what we cannot do with fMRI. *Nature 453* (Jun 2008), 869–878. 1

[LRMV99] LACHAUX J. P., RODRIGUEZ E., MARTINERIE J., VARELA F. J.: Measuring phase synchrony in brain signals. *Hum Brain Mapp 8* (1999), 194–208. 2

[MIT91] MUSHIAKE H., INASE M., TANJI J.: Neuronal activity in the primate premotor, supplementary, and precentral motor cortex during visually guided and internally determined sequential movements. *J. Neurophysiol. 66* (Sep 1991), 705–718. 1

[Rai06] RAICHLE M. E.: Neuroscience. The brain's dark energy. *Science 314* (Nov 2006), 1249–1250. 2

[RM06] RAICHLE M. E., MINTUN M. A.: Brain work and brain imaging. *Annu. Rev. Neurosci. 29* (2006), 449–476. 2

[SFM*09] SMITH S. M., FOX P. T., MILLER K. L., GLAHN D. C., FOX P. M., MACKAY C. E., FILIPPINI N., WATKINS K. E., TORO R., LAIRD A. R., BECKMANN C. F.: Correspondence of the brain's functional architecture during activation and rest. *Proc. Natl. Acad. Sci. U.S.A. 106* (Aug 2009), 13040–13045. 2

[SJW*04] SMITH S. M., JENKINSON M., WOOLRICH M. W., BECKMANN C. F., BEHRENS T. E., JOHANSEN-BERG H., BANNISTER P. R., LUCA M. D., DROBNJAK I., FLITNEY D. E., NIAZY R. K., SAUNDERS J., VICKERS J., ZHANG Y., STEFANO N. D., BRADY J. M., MATTHEWS P. M.: Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage 23*, Supplement 1 (2004), S208 – S219. Mathematics in Brain Imaging. 5

[SSC*05] SALVADOR R., SUCKLING J., COLEMAN M. R., PICKARD J. D., MENON D., BULLMORE E.: Neurophysiological architecture of functional magnetic resonance images of human brain. *Cereb. Cortex 15* (Sep 2005), 1332–1342. 2

[tCMR08] TEN CAAT M., MAURITS N. M., ROERDINK J. B. T. M.: Data-driven visualization and group analysis of multichannel EEG coherence with functional units. *IEEE Trans. Vis. Comput. Graph. 14*, 4 (2008), 756–771. 2, 3, 5

[TDP06] THIRION B., DODEL S., POLINE J. B.: Detection of signal synchronizations in resting-state fMRI datasets. *Neuroimage 29* (Jan 2006), 321–327. 2

[VE05] VAN ESSEN D. C.: A Population-Average, Landmark- and Surface-based (PALS) atlas of human cerebral cortex. *Neuroimage 28* (Nov 2005), 635–662. 3

[VEDD*01] VAN ESSEN D. C., DRURY H. A., DICKSON J., HARWELL J., HANLON D., ANDERSON C. H., VAN ESSEN D. C.: An integrated software suite for surface-based analyses of cerebral cortex. *J Am Med Inform Assoc 8* (2001), 443–459. 3

[WR04] WINK A. M., ROERDINK J. B. T. M.: Denoising functional MR images: a comparison of wavelet denoising and Gaussian smoothing. *IEEE Trans. Med. Imaging 23*, 3 (2004), 374–387. 4