A Connectivity-Based Method for Defining Regions-of-Interest in fMRI Data

Filip Deleus and Marc M. Van Hulle, *Senior Member, IEEE*

Abstract—In this paper, we describe a new methodology for defining brain regions-of-interset (ROIs) in functional magnetic resonance imaging (fMRI) data. The ROIs are defined based on their functional connectivity to other ROIs, i.e., ROIs are defined as sets of voxels with similar connectivity patterns to other ROIs. The method relies on 1) a spatially regularized canonical correlation analysis for identifying maximally correlated signals, which are not due to correlated noise; 2) a test for merging ROIs which have similar connectivity patterns to the other ROIs; and 3) a graph-cuts optimization for assigning voxels to ROIs. Since our method is fully connectivity-based, the extracted ROIs and their corresponding time signals are ideally suited for a subsequent brain connectivity analysis.

Index Terms—fMRI, functional connectivity, image segmentation.

I. INTRODUCTION

THE fMRI signal is generally regarded as a proxy for the underlying neural activity. Most fMRI studies are concerned with the detection of foci of activation in the brain and with modelling the relationship to the experimental paradigm. These studies are called *activation* studies. However, disjunct regions of the brain do not operate in isolation so that there is a growing interest in studying the interactions and connectivity patterns between these regions, yielding an augmented understanding of the functional organization of the brain. These studies are called *connectivity* studies. Two main approaches exist for the analysis of fMRI data, a *voxel-based approach* in which each individual voxel is treated separately, and a *ROI-based approach* in which one first defines regions-of-interest and then performs the analysis on the level of these regions. The most popular voxel-based approach for connectivity analysis is the seed region approach. A connectivity map is calculated in which the value at each voxel represents the correlation between the time signal of this voxel and that of the

seed voxel. This technique is subjective since the exploration of the brain connectivity heavily depends on the choice of the seed voxel. Other voxel-based approaches include principal components analysis (PCA) [1] and independent components analysis (ICA) [2]. These techniques maximize a mathematical criterion, based on all voxels. However, the physiological interpretation of these results is often difficult, and prior anatomical or neurophysiological knowledge cannot be easily incorporated in the analysis. Finally, we have the ROI-based approaches and their advantage is that they compute connectivities between brain regions, thus anatomically-related entities, leading to a much easier interpretation of the connectivity patterns found. The most popular ROI-based approaches are structural equations modelling (SEM) [3] and dynamic causal modelling (DCM) [4], which are used for the model fitting of an assumed connectivity pattern between a set of ROIs (i.e., between their representative signals). However, ROIs and their representative signals can also be used in exploratory (not confirmatory) effective connectivity studies. These methods include Bayesian networks (BN) and dynamic Bayesian networks (DBN) which have been used to achieve learning functional structure from fMRI data [5], [6], and multivariate analyses using the Granger Causality concept [7]–[10].

ROIs are frequently defined based on their cytoarchitectonic structure or anatomical landmarks such as sulci [11], [12], or on their functional specialization. However, the delineation of a ROI is a very difficult task because the automatic labeling of many sulci is still an open problem, and tools for defining 3-D volumes on the cortex are difficult to use (a manual parcellation can take many hours or days) [13]. Furthermore, the exact parcellations of the cortex and the precise boundaries of many brain areas are still a matter of debate. There exist also data-driven methods to assign voxels to ROIs. These methods are cluster analysis [14]–[17] and region-growing methods [18]–[20]. The objective of cluster analysis is to maximize the intragroup homogeneity within the ROIs, based on a distance measure between the voxels of a ROI. The most applied distance measures in fMRI are the correlation coefficient or Kendall's concordance coefficient [21] between the voxels' time signals. In order to obtain contiguous clusters of spatially connected voxels, the spatial distance between the voxels can be incorporated in the distance measure, as for example in [22]. On the other hand, the objective of the region growing methods is also to maximize the intragroup homogeneity within the ROIs, but they account directly for the spatial structure in the fMRI data by starting from a set of seed voxels and adding only neighboring voxels to the ROIs.

After the assignment of voxels to ROIs, a representative signal can be defined for the further connectivity analysis. Three methods are frequently used for this. In the first method,

The authors are with the K. U. Leuven, Laboratorium Neuro- en Psychofysiologie, Campus Gasthuisberg, O&N 2, Herestraat 49, bus 1021, 3000 Leuven, Belgium (e-mail: filip.deleus@med.kuleuven.be).

the time signal of the peak voxel of fMRI activation is taken as the representative signal for the ROI. This method is biased by which contrast between the experimental conditions has been used. In the second method, the average signal is taken accross all the voxels of a ROI, which, however, may be distorted by other (noise) signals present in the ROI. In the third method, a PCA is first performed on the time signals of the individual ROIs, and the first principal components are taken as the representative signals. This method is mostly used in fMRI connectivity studies. However, PCA does only account for the variance within one region and not for the relations between regions.

We propose a new method which defines ROIs based on their (functional) connectivities to other ROIs, i.e., as *sets of voxels with similar connectivity patterns* to other ROIs. Our method is mainly based on canonical correlation analysis (CCA). CCA [23] is a multivariate statistical technique commonly used to identify and quantify the correlation between two sets of variables. More in detail, given two data sets X and Y, CCA seeks to find a linear transformation of the variables in X and a linear transformation of the variables in Y such that the resulting two new variables are maximally correlated. CCA has already been used in fMRI for detecting neural activity [24]–[26]. This activation detection has been shown to work well and the reason is that CCA enables simultaneous temporal modeling and adaptive spatial filtering of the data [27]. We will use CCA to define functionally connected ROIs in fMRI data.

Our method starts from a given set of ROIs, which may be based on prior neurophysiological knowledge. Constrained canonical correlation analysis (CCA) has been performed on these sets of voxels in order to find subsets that are maximally connected. We modify the classical CCA-algorithm in two ways: first by introducing spatial regularization constraints to get results that are more compliant with the spatial structure of fMRI data, second by changing the goal function of the optimization to make the algorithm more indifferent to spatially-correlated noise. In the next step, local maxima are identified on the weight vectors of the CCA. We merge neighbouring local maxima if they have similar connectivity patterns to the other ROIs. The local maxima are now centers of new ROIs, and voxels are assigned to these regions using graph-cuts optimization. After the ROIs are converged, a connectivity analysis can be performed, based on any of the current methodologies. Our method can be summarized as follows:
— Input: ROIs = sets of voxels.
— Iterate until a stable solution has been reached.
  1) Perform CCA on the given sets of voxels (Section II).
  2) Identify local maxima in the weights patterns of the CCA.
  3) Merge local maxima based on their connectivities to other ROIs (Section III-A).
  4) Assign voxels to local maxima to define new ROIs (Section III-B).
  5) Go back to step 1.
— Perform connectivity analysis (SEM, DCM, BN, DBN, Granger causality-based approaches, etc.) on the representative signals of the ROIs

We demonstrate our method in Section IV on a synthetic and a real fMRI data set and we compare our results with clusters that
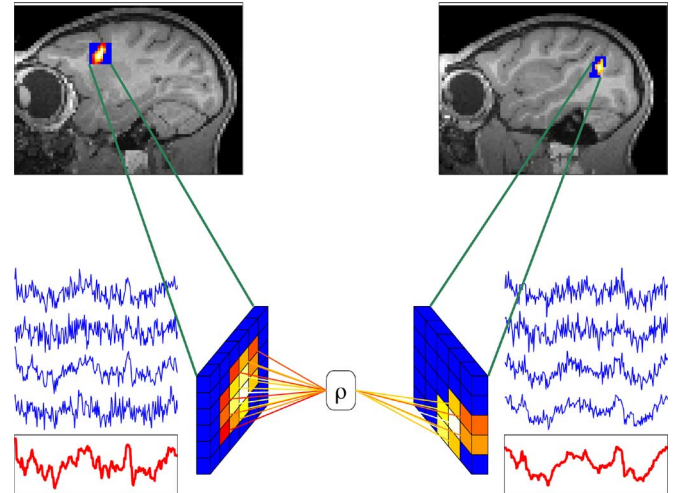

Fig. 1. Illustration of spatially regularized CCA: subsets of given sets of voxels are selected to delineate ROIs in which the representative signals (red) have maximum correlation: the blue voxels have zero weight and the strength of the weight increases from orange to white.

are found by K-means clustering, first on the set of voxels that were selected by our method, and second, on the set of voxels that were selected based on Kendall's concordance coefficient.

II. SPATIALLY REGULARIZED CANONICAL CORRELATION ANALYSIS

In this section, we introduce our spatially regularized CCA algorithm. The aim is to find subsets of voxels from the given regions, the representative signals of which are maximally correlated. As an illustration, consider Fig. 1. Two sets of voxels are given, one based on prior neurophysiological knowledge, to include the motion area MT in the temporal sulcus (right panel), and one broadly defined cube in the frontal lobe to include the frontal eye field, FEF (left panel). Spatially regularized CCA delineates the two ROIs in both predefined sets of voxels such that the representative signals of both ROIs are maximally correlated. In the next subsections, this idea will be worked out more formally. We start with the case of two regions, and extend it later to multiple regions.

A. CCA on Two ROIs

Let $\mathbf{X} = [\mathbf{x}_1 \ \mathbf{x}_2 \ \cdots \ \mathbf{x}_n]$ be a $(T \times n)$-dimensional data matrix consisting of the n fMRI signals of length T. Each signal corresponds to one voxel in ROI X. Let $\mathbf{Y} = [\mathbf{y}_1 \ \mathbf{y}_2 \ \cdots \ \mathbf{y}_m]$ be a $(T \times m)$-dimensional data matrix consisting of the m fMRI signals of length T of region Y. All the time signals are standardized to zero mean and unit variance. With CCA, we obtain the weight vectors $\mathbf{w}_x$ and $\mathbf{w}_y$ so that the correlation $\rho_{x,y}$ between the two resulting linear combinations is maximal

$$\mathbf{x} = w_{x1}\mathbf{x}_1 + w_{x2}\mathbf{x}_2 + \cdots + w_{xn}\mathbf{x}_n = \mathbf{X}\mathbf{w}_x \quad (1)$$

$$\mathbf{y} = w_{y1}\mathbf{y}_1 + w_{y2}\mathbf{y}_2 + \cdots + w_{yn}\mathbf{y}_n = \mathbf{Y}\mathbf{w}_y. \quad (2)$$

Due to [23] this problem can be reduced to two eigenvalue problems: the weight vector $\mathbf{w}_x$ is found as the eigenvector belonging to the largest eigenvalue of $(\mathbf{X}'\mathbf{X})^{-1}(\mathbf{X}'\mathbf{Y})(\mathbf{Y}'\mathbf{Y})^{-1}(\mathbf{Y}'\mathbf{X})$ and $\mathbf{w}_y$ is found as the eigenvector belonging

to the largest eigenvalue of $(\mathbf{Y}'\mathbf{Y})^{-1}(\mathbf{Y}'\mathbf{X})\,(\mathbf{X}'\mathbf{X})^{-1}(\mathbf{X}'\mathbf{Y})$. The resulting signals $\mathbf{X}\mathbf{w}_x$ and $\mathbf{Y}\mathbf{w}_y$ can then be taken as the representative time signals for the ROIs X and Y, respectively. However, these representative signals cannot be interpreted as a weighted mean of the original time signals since the weight vectors could be positive or negative. Therefore, we impose nonnegativity constraints $\mathbf{w}_x \geq \mathbf{0}$ and $\mathbf{w}_y \geq \mathbf{0}$ on the weight vectors. Furthermore, it is known from fMRI analysis that true neural activation typically tends to occur in a functional cluster [28], [29] or, in other words, that true fMRI activation is more likely to occur in clusters of spatially connected voxels than in a single voxel or in small groups of isolated voxels. However, due to the presence of spatial correlations in fMRI signals of neighbouring voxels, the weight vectors $\mathbf{w}_x$ and $\mathbf{w}_y$ will not reflect this spatial organisation, but will tend to yield scattered weight patterns with high variances on each individual weight value. This phenomenon is analogous to the multicollinearity problem in regression analysis: the matrices $(\mathbf{X}'\mathbf{X})$ and $(\mathbf{Y}'\mathbf{Y})$ become more singular with higher spatial correlations, and the inverse operations on them, $(\mathbf{X}'\mathbf{X})^{-1}$ and $(\mathbf{Y}'\mathbf{Y})^{-1}$, lead to increased variances. In order to have weight vectors which are more compliant with the spatial structure of the data, we impose spatial regularization constraints. For each voxel v_i, we basically impose that $(\sum_{j \in N_i} w_j / S_{N_i}) \geq \kappa w_i$, with N_i the neighbourhood of voxel v_i and S_{N_i} the size of this neighbourhood. The regularization parameter κ lies in the interval $[0\,1]$ and controls the degree of spatial regularization: with $\kappa = 0$ there is no spatial smoothing, while with $\kappa = 1$ all the weight values in one weight vector are equal and the resulting signals are the unweighted means of the original signals. The spatial regularization constraints on the voxels of one region, say region X, can be summarized as $\mathbf{R}_x\mathbf{w}_x \leq \mathbf{0}$. For the $(n \times n)$ square matrix $\mathbf{R}_x$, the diagonal elements are equal to κ, and $\forall j \in N_i$, the element on (i,j) is equal to $-1/S_{N_i}$. The optimization problem for the spatially regularized CCA can then be formulated as

$$\max_{\mathbf{w}_x, \mathbf{w}_y} \quad \rho_{x,y} = \frac{\mathbf{w}_x'\mathbf{X}'\mathbf{Y}\mathbf{w}_y}{\sqrt{\mathbf{w}_x'\mathbf{X}'\mathbf{X}\mathbf{w}_x}\sqrt{\mathbf{w}_y'\mathbf{Y}'\mathbf{Y}\mathbf{w}_y}} \tag{3}$$

$$\text{s.t.} \quad \mathbf{R}_x\mathbf{w}_x \leq \mathbf{0} \tag{4}$$

$$\mathbf{R}_y\mathbf{w}_y \leq \mathbf{0} \tag{5}$$

$$\mathbf{w}_x \geq \mathbf{0} \tag{6}$$

$$\mathbf{w}_y \geq \mathbf{0}. \tag{7}$$

We developed an active-set algorithm that exploits the special characteristics of this optimization problem. An active set algorithm is based on binding and nonbinding constraints: if the solution does not violate a given constraint after it has been deleted from the formulation, this constraint is called nonbinding, otherwise it is a binding constraint. The general principles of an active set algorithm are as follows (for a more general description of active set algorithms we refer to [30]).

- A binding inequality constraint becomes an equality constraint (it is called an *active* constraint).
- A nonbinding inequality constraint has no influence and can be safely removed from the problem formulation (it is called a *passive* constraint).

- An active constraint (= equality constraint) can become passive in a next iteration, depending on the value of its Lagrange-multiplier.
- A passive constraint can be violated in a next iteration. A line search is then performed between the previous solution and the current solution, until the solution is back at the borders of the feasible region. The violated inequality constraint then becomes an equality constraint (the constraint becomes active).

For one region, say X, there are $2n$ inequality constraints: n nonnegativity constraints and n regularization constraints. However, if a nonnegativity constraint becomes active, then the corresponding value becomes 0, which is the same as removing the corresponding time signal (and voxel) from the analysis. Furthermore, if a time signal has been removed, also the corresponding row and column has to be deleted in the regularization matrix $\mathbf{R}_x$. Thus, the equality constraints (active constraints) can only come from the regularization constraints. The optimization problem with active constraints can now be formulated as

$$\max_{\mathbf{w}_x, \mathbf{w}_y} \quad \rho_{x,y} = \mathbf{w}_x'\mathbf{X}'\mathbf{Y}\mathbf{w}_y \tag{8}$$

$$\text{s.t.} \quad \mathbf{R}_x\mathbf{w}_x = \mathbf{0} \tag{9}$$

$$\mathbf{R}_y\mathbf{w}_y = \mathbf{0} \tag{10}$$

$$\mathbf{w}_x'\mathbf{X}'\mathbf{X}\mathbf{w}_x = 1 \tag{11}$$

$$\mathbf{w}_y'\mathbf{Y}'\mathbf{Y}\mathbf{w}_y = 1 \tag{12}$$

in which the denominator of the objective function has been replaced by unit variance constraints. The Lagrange-formulation is then

$$L = \mathbf{w}_x'\mathbf{X}'\mathbf{Y}\mathbf{w}_y - \lambda_x'\mathbf{R}_x\mathbf{w}_x - \lambda_y'\mathbf{R}_y\mathbf{w}_y$$
$$- \frac{\rho_x}{2}(\mathbf{w}_x'\mathbf{X}'\mathbf{X}\mathbf{w}_x - 1) - \frac{\rho_y}{2}(\mathbf{w}_y'\mathbf{Y}'\mathbf{Y}\mathbf{w}_y - 1) \tag{13}$$

which we solve by setting its partial derivatives to zero

$$\frac{\partial L}{\partial \mathbf{w}_x} = \mathbf{X}'\mathbf{Y}\mathbf{w}_y - \mathbf{R}_x'\lambda_x - \rho_x\mathbf{X}'\mathbf{X}\mathbf{w}_x = \mathbf{0} \tag{14}$$

$$\frac{\partial L}{\partial \mathbf{w}_y} = \mathbf{Y}'\mathbf{X}\mathbf{w}_x - \mathbf{R}_y'\lambda_y - \rho_y\mathbf{Y}'\mathbf{Y}\mathbf{w}_y = \mathbf{0}. \tag{15}$$

Subtracting $\mathbf{w}_x'$ times (14) from $\mathbf{w}_y'$ times (15) and given the constraints (9)–(12) one must conclude that $\rho_x = \rho_y = \rho$. After premultiplying (14) by $\mathbf{R}_x'(\mathbf{X}'\mathbf{X})^{-1}$ and given (9)–(12), one obtains

$$\lambda_x = (\mathbf{R}_x(\mathbf{X}'\mathbf{X})^{-1}\mathbf{R}_x')^{-1}\mathbf{R}_x(\mathbf{X}'\mathbf{X})^{-1}(\mathbf{X}'\mathbf{Y})\mathbf{w}_y. \tag{16}$$

We define $\mathbf{M}_x$ such that

$$\lambda_x = \mathbf{M}_x\mathbf{w}_y \tag{17}$$

and we define λ_y and $\mathbf{M}_y$ in a similar way. Putting λ_x and λ_y in (14)–(15), we find, after some algebraic manipulations, that

$$(\mathbf{\Phi}_x \cdot \mathbf{\Phi}_y)\mathbf{w}_x = \rho^2\mathbf{w}_x \tag{18}$$

$$(\mathbf{\Phi}_y \cdot \mathbf{\Phi}_x)\mathbf{w}_y = \rho^2\mathbf{w}_y \tag{19}$$

with

$$\mathbf{\Phi}_x = (\mathbf{X'X})^{-1}(\mathbf{X'Y} - \mathbf{R}_x'\mathbf{M}_x) \tag{20}$$

$$\mathbf{\Phi}_y = (\mathbf{Y'Y})^{-1}(\mathbf{Y'X} - \mathbf{R}_y'\mathbf{M}_y). \tag{21}$$

Equations (18)–(19) are two eigenvalue problems. Hence, the optimization problem with active constraints has a closed-form solution. In the active set algorithm, (14) and (15) determine whether an active nonnegativity constraint becomes passive, i.e., whether a variable (voxel) enters the formulation, while (16) determines whether a spatial regularization constraint becomes passive, since the Lagrange multipliers λ_x and λ_y can be considered as shadow costs of the regularization constraints.

B. CCA on Multiple ROIs

Thus far, we have applied CCA between two brain regions. Formally, generalizations to three or more sets have been proposed by [31] and [32], for unconstrained CCA. We adopt the `maxvar` approach in which the first eigenvalue of the correlation matrix of the resulting signals has been maximized, subject to the nonnegativity and spatial regularization constraints. The intuition behind the `maxvar` approach is the same as in PCA on the correlation matrix of a set of signals: the first principal component (first eigenvector) makes a linear combination of the signals so that the new signal after projection has maximum variance. This maximum variance will be obtained if the signals that are most correlated also have the largest values in the eigenvector. However, the `maxvar` approach goes a step deeper than PCA, since not only the eigenvector of the correlation matrix has to be determined (to make the linear combination that yields a signal with maximum variance), but also the signals on which the correlation matrix is calculated, have to be determined, by calculating optimal weight vectors $\mathbf{w}_x$. More formally, this approach can be formulated as follows.

Let $X_1, X_2, \ldots, X_p$ be p brain regions with data matrices $\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_p$, then the optimization problem is

$$\max_{\mathbf{v}, \mathbf{w}_1 \ldots \mathbf{w}_p} \quad \mathbf{v'Qv} \tag{22}$$

$$\text{s.t.} \quad \mathbf{w}_i'\mathbf{X}_i'\mathbf{X}_i\mathbf{w}_i = 1 \quad i = 1 \ldots p \tag{23}$$

$$\mathbf{R}_i\mathbf{w}_i \leq \mathbf{0} \quad i = 1 \ldots p \tag{24}$$

$$\mathbf{w}_i' \geq \mathbf{0} \quad i = 1 \ldots p \tag{25}$$

$$\mathbf{v'v} = 1 \tag{26}$$

with

$$
\begin{aligned}
\mathbf{Q} &= [\mathbf{x}_1 \quad \mathbf{x}_2 \quad \cdots \quad \mathbf{x}_p]'[\mathbf{x}_1 \quad \mathbf{x}_2 \quad \cdots \quad \mathbf{x}_p] \\
&= \begin{bmatrix}
\mathbf{w}_1'\mathbf{X}_1'\mathbf{X}_1\mathbf{w}_1 & \mathbf{w}_1'\mathbf{X}_1'\mathbf{X}_2\mathbf{w}_2 & \cdots & \mathbf{w}_1'\mathbf{X}_1'\mathbf{X}_p\mathbf{w}_p \\
\mathbf{w}_2'\mathbf{X}_2'\mathbf{X}_1\mathbf{w}_1 & \mathbf{w}_2'\mathbf{X}_2'\mathbf{X}_2\mathbf{w}_2 & \cdots & \mathbf{w}_2'\mathbf{X}_2'\mathbf{X}_p\mathbf{w}_p \\
\vdots & & \ddots & \\
\mathbf{w}_p'\mathbf{X}_p'\mathbf{X}_1\mathbf{w}_1 & \mathbf{w}_p'\mathbf{X}_p'\mathbf{X}_2\mathbf{w}_2 & \cdots & \mathbf{w}_p'\mathbf{X}_p'\mathbf{X}_p\mathbf{w}_p
\end{bmatrix}.
\end{aligned}
\tag{27}
$$

This problem is solved by iterating the above described active set algorithm in which, at each iteration, one weight vector $\mathbf{w}_i$ is optimized given the eigenvector $\mathbf{v}$ and the other $(p-1)$ weight vectors. If $p = 2$, this formulation is equal to the CCA formulation in the previous section. For $p > 2$, the eigenvector $\mathbf{v}$ gives an indication of the importance of each region and, hence, can be interpreted as a measure of functional connectivity between the p brain regions.

C. Removing Correlated Noise

Nonwhite noise in fMRI has been a known problem for a long time. The sources are many: low-frequency drift due to hardware imperfections, oscillatory noise due to respiration and cardiac pulsation, residual movement artefacts not accounted for by (non)-rigid body registration, and so on. These noise contributions give rise to temporal autocorrelations in the fMRI signals. Furthermore, it is also known [33] that the nonwhite noise is structured in space. It is clear that the proposed CCA-solution will be sensitive to this spatial correlation in the noise, and that we have to take these noise correlation into account if we want to measure the connectivity between ROIs.

Let $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{y}}$ be 2 unobservable fMRI signals. Due to the noise, we only observe $\mathbf{x}$ and $\mathbf{y}$ with

$$\mathbf{x} = \tilde{\mathbf{x}} + \mathbf{u}_x \tag{28}$$

$$\mathbf{y} = \tilde{\mathbf{y}} + \mathbf{u}_y \tag{29}$$

in which $\mathbf{u}_x$ and $\mathbf{u}_y$ are the noise components. In the CCA-approach, we determine the weight vectors $\mathbf{w}_x$ and $\mathbf{w}_y$ so that the correlation between $\mathbf{x}$ and $\mathbf{y}$ is maximized. We now change the goal function to maximize $(\mathbf{x'y} - \mathbf{u}_x'\mathbf{u}_y)$, which is the part of correlation that is not due to the correlated noise. More in detail, this goal function is an approximation of $\tilde{\mathbf{x}}'\tilde{\mathbf{y}}$ in which we assume zero-correlation between $\tilde{\mathbf{x}}$ and $\mathbf{u}_y$, and between $\tilde{\mathbf{y}}$ and $\mathbf{u}_x$. We keep the constraints $\mathbf{w}_x'\mathbf{x'x}\mathbf{w}_x = 1$ and $\mathbf{w}_y'\mathbf{y'y}\mathbf{w}_y = 1$ unchanged. For the CCA with multiple regions, the matrix $\mathbf{Q}$ in (22) has been changed to (30), shown at the bottom of the page. The matrices $\mathbf{U}_i$ contain the noise signals. If the fMRI data is obtained in a block design, the following procedure can be followed: first, take for each signal the mean over all runs, next calculate in each run the noise as the difference between the signal in that run and the mean signal over runs. However, any noise estimation method (e.g., [34] and [35]) can be used to fill in the matrices $\mathbf{U}_i$.

III. REDEFINING THE ROIs

The analysis in Section II occurs *within* the given sets of voxels. In order to incorporate voxels in the analysis which are not included in the initial ROIs, we developed a procedure to redefine the ROIs. This procedure works as follows. First, we identify local maxima in the weight vectors of the CCA, and

$$
\mathbf{Q} = \begin{bmatrix}
\mathbf{w}_1'\mathbf{X}_1'\mathbf{X}_1\mathbf{w}_1 & \cdots & \mathbf{w}_1'(\mathbf{X}_1'\mathbf{X}_p - \mathbf{U}_1'\mathbf{U}_p)\mathbf{w}_p \\
\vdots & \ddots & \vdots \\
\mathbf{w}_p'(\mathbf{X}_p'\mathbf{X}_1 - \mathbf{U}_p'\mathbf{U}_1)\mathbf{w}_1 & \cdots & \mathbf{w}_p'\mathbf{X}_p'\mathbf{X}_p\mathbf{w}_p
\end{bmatrix} \tag{30}
$$

take a large cube of voxels around each local maximum. If the weight vector $\mathbf{w}_x$ of a ROI X has a single (global) maximum, then the region can still move. If $\mathbf{w}_x$ has multiple local maxima then the region will be splitted into separate regions. In order to avoid an explosion of new regions, we test whether nearby regions (i.e., if some of their voxels overlap) can be merged into one region. Finally, overlapping voxels between two or more regions are assigned to one local maximum in order to end up with a new set of nonoverlapping ROIs. These new ROIs can then be used as inputs to a new CCA-stage. The algorithm stops when it has converged to a stable solution i.e., when the position of the local maxima and the assignment of voxels to these maxima is the same in two consecutive iterations. In the next two sections, we first describe the test for region merging, and then the assignment of voxels to regions.

A. Merging Local Maxima

Let $\mathbf{y}_1$ and $\mathbf{y}_2$ be the time signals of two neigbouring local maxima in the weight vectors of the canonical correlation analysis. These two weight vectors can come from one weight vector (from one ROI) or from different weight vectors (from different ROIs). If CCA had been performed on p regions, let $\mathbf{X}$ be the data matrix that contains the representative time signals of the ROIs that were not the source for $\mathbf{y}_1$ or $\mathbf{y}_2$. Thus, the dimension of $\mathbf{X}$ is $[T \times (p-1)]$ or $[T \times (p-2)]$, depending on whether $\mathbf{y}_1$ and $\mathbf{y}_2$ come from the same or different ROIs, respectively. The connectivity patterns from $\mathbf{y}_1$ and $\mathbf{y}_2$ to the other regions can be modelled as regression coefficients $\boldsymbol{\beta}_1$ and $\boldsymbol{\beta}_2$ from the two regression models

$$\mathbf{y}_1 = \mathbf{X}\boldsymbol{\beta}_1 + \boldsymbol{\epsilon}_1 \tag{31}$$

$$\mathbf{y}_2 = \mathbf{X}\boldsymbol{\beta}_2 + \boldsymbol{\epsilon}_2. \tag{32}$$

These two regression models can be summarized by the model

$$\begin{bmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{X} & \mathbf{0} \\ \mathbf{0} & \mathbf{X} \end{bmatrix} \cdot \begin{bmatrix} \boldsymbol{\beta}_1 \\ \boldsymbol{\beta}_2 \end{bmatrix} + \begin{bmatrix} \boldsymbol{\epsilon}_1 \\ \boldsymbol{\epsilon}_2 \end{bmatrix}. \tag{33}$$

In order to test whether the connectivity patterns $\boldsymbol{\beta}_1$ and $\boldsymbol{\beta}_2$ are similar, we impose equality constraints $\boldsymbol{\beta}_1 = \boldsymbol{\beta}_2 = \boldsymbol{\beta}$ on (33) and compare the sum of the squared residuals of the constrained model with that of the unconstrained model. One can easily verify that the solution of the constrained model can be obtained as

$$\boldsymbol{\beta} = (2\mathbf{X}'\mathbf{X})^{-1}(\mathbf{X}'\mathbf{y}_1 + \mathbf{X}'\mathbf{y}_2) \tag{34}$$

and the solution of the unconstrained problem is identical to that of the two separate regression models (31) and (32)

$$\boldsymbol{\beta}_1 = (\mathbf{X}'\mathbf{X})^{-1} \cdot (\mathbf{X}'\mathbf{y}_1) \tag{35}$$

$$\boldsymbol{\beta}_2 = (\mathbf{X}'\mathbf{X})^{-1} \cdot (\mathbf{X}'\mathbf{y}_2) \tag{36}$$

which is the standard OLS solution. Let eer and eeu be the sums of the squared residuals in the constrained and the unconstrained models, respectively, the test for equality of the two models is then

$$F = \frac{(\text{eer} - \text{eeu})/d_x}{\text{eeu}/(T - d_x)} = F(d_x, T - d_x) \tag{37}$$
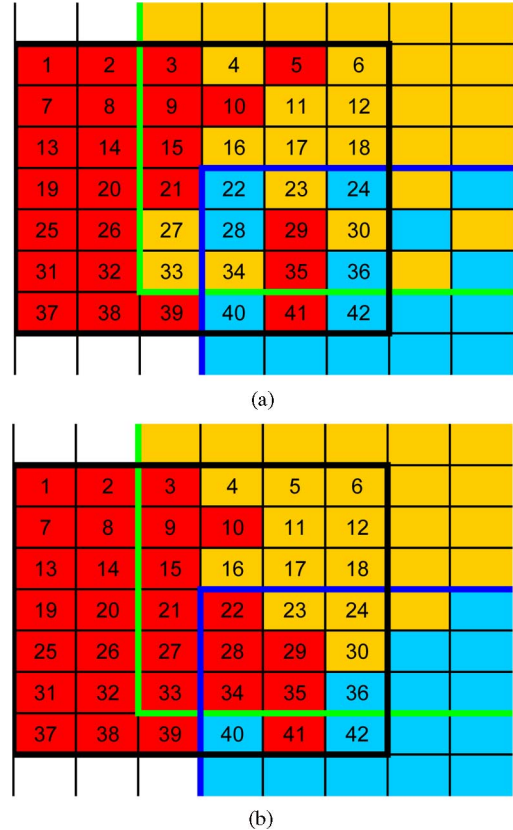

Fig. 2. Illustration of the voxel assignment problem. Three overlapping regions are defined (black border: X, green border: Y, dark blue border: Z) and labels are assigned to the voxels (red: X, yellow: Y, blue: Z). The left panel shows an initial labeling with isolated voxels, the right panel shows a desired solution consisting of contiguous parts.

where d_x is equal to $(p-1)$ or $(p-2)$, i.e., the number of columns in $\mathbf{X}$.

B. Assigning Voxels to ROIs Using Graph-Cuts Optimization

The retained local mimima are further used as seeds around which a (large) spherical or cubical set of voxels is defined. However, since these large regions may consist of many overlapping voxels, we are now concerned with the problem of assigning each voxel to exactly one seed (and, hence, to one ROI). As a first labeling, we could assign a voxel to the seed with which it has the largest correlation. However, since no spatial regularization has been included yet, this labeling will yield regions consisting of isolated voxels or unconnected groups of them. This is conceptually shown in Fig. 2(a). In this figure, three overlapping regions are defined, say region X defined by the black border, region Y defined by the green border and region Z defined by the dark blue border. The labels that are currently assigned to the voxels are shown by a color code: red means that a voxel has been assigned to region X, yellow that is has been assigned to region Y, and light blue that is has been assigned to region Z. Fig. 2(b) shows a desired solution with regions consisting of contiguous parts without isolated or unconnected voxels. In order to obtain such a solution, we include a spatial smoothness term in the labeling. The assignment

problem can then be formulated as the minimization of the following well-known energy function

$$\min_f E(f) = \sum_v D_v(f_v) + \sum_{\{v,q\}\in N} V(f_v, f_q) \quad (38)$$

in which f_v is the label of voxel v, or in other words, the label of the local maximum to which voxel v has been assigned. The first term, $\sum_v D_v(f_v)$ measures the cost of labeling voxel v as f_v, or the cost of assigning voxel v to the seed f_v. The second term, $\sum_{\{v,q\}\in N} V(f_v, f_q)$ controls the spatial regularization, where N is the set of all pairs of neighboring voxels. This energy function can efficiently be minimized using graph-cuts.

Graph-cuts is basically an energy minimization technique based on combinatorial optimization. It has many interesting characteristics such as global minimization, solid theoretical background and flexibility in the energy function. Many energy functions can be minimized via graph-cuts as soon as the problem can be formalized as a network flow problem. The construction of such networks or graphs is fully described in [36]. In our method, we used the $\alpha\beta$-swap algorithm described in [37]. In our application of the graph-cut algorithm, we define $D_v(f_v)$ as $(1 - \rho(v, f_v))$. For the regularization term, we adopt the Potts model $V(\alpha, \beta) = K.T(\alpha \neq \beta)$ where $T(.)$ equals 1 if its argument is true, and otherwise 0. The degree of spatial regularization has been controlled by the parameter K. The $\alpha\beta$-swap algorithm works iteratively on all pairs of labels. For a given pair of labels, say α and β, the algorithm can switch labels from α to β or vice-versa. Voxels that are not currently labeled as α or β remain unchanged in that iteration. In principle, the algorithm works on all pairs of labels, however, since in the voxel assignment problem only the overlapping voxels have to be (re-)assigned, the iteration must not go over all pairs of ROIs. For each $\alpha\beta$-swap, a graph has been constructed consisting of nodes and links between them. Let $P_{\alpha\beta}$ be the set of voxels, the labels of which may be switched during the current iteration. The nodes in the constructed graph are then the voxels in $P_{\alpha\beta}$, together with a source node α, and a sink node β. As an illustration, consider again Fig. 2. If an $\alpha\beta$-swap has been performed between region X and Y, the voxels in $P_{\alpha\beta}$ are [3, 4, 5, 6, 9, 10, 11, 12, 15, 16, 17, 18, 21, 23, 27, 29, 30, 33, 34, 35]. Two types of links are introduced in the graph, first the t_v^α and t_v^β links which connect the voxels $v \in P_{\alpha\beta}$ with the source and sink nodes, and, second the $e_{\{v,q\}}$ links which connect neighbouring voxels v and q ($v, q \in P_{\alpha\beta}$). According to [37], the weights assigned to these links are

$$t_v^\alpha = D_v(\alpha) + \sum_{\substack{q\in N_v \\ q\notin P_{\alpha\beta}}} V(\alpha, f_q) \quad \text{for } v\in P_{\alpha\beta} \quad (39)$$

$$t_v^\beta = D_v(\beta) + \sum_{\substack{q\in N_v \\ q\notin P_{\alpha\beta}}} V(\beta, f_q) \quad \text{for } v \in P_{\alpha\beta} \quad (40)$$

$$e_{\{v,q\}} = V(\alpha, \beta) \quad \text{for } \begin{bmatrix} \{v,q\} \in N \\ v, q \in P_{\alpha\beta} \end{bmatrix} \quad (41)$$

where N_v is the neighbourhood of voxel v, and N the set of all pairs of neighbouring voxels. A minimal cut on this graph yields an optimal labeling for the $\alpha\beta$ swap, given the current labeling of the other voxels which are not in $P_{\alpha\beta}$. A minimal

cut on the graph can be computed in polynomial time via one of the maximum-flow algorithms as proposed by [38]. We used a more recent and efficient maximum-flow algorithm proposed by [39].

IV. RESULTS

A. Synthetic Data Set

First, we demonstrate our method on a synthetic data set. In order to make synthetic data compliant with the properties of real fMRI data, we added simulated connected ROIs to a real fMRI data set (which will be described in the next section). As a ground truth we added 6 connected ROIs to the fMRI data. Signals are generated simulating 10 runs of 3 conditions with 12 data points per condition (signals have length 360). The length of the simulated runs is different from the length of the runs in the real fMRI data, which is in this case considered as noise data. We defined two underlying signals X_A and X_B and generated signals for the 6 ground truth ROIs using the generator model defined in Fig. 3 (where ϵ represents noise, coming from the real fMRI data). Let $X_1, X_2, \ldots, X_6$ be the ground truth ROIs (shown in Fig. 3) and let the ROIs defined in the algorithm be denoted by Y, whereby $Y_{(i,k)}$ denotes the kth ROI in iteration i. The Y-ROIs for the different iterations are shown in Fig. 4. For example, $Y_{(1,4)}$ is represented by the yellow blob in Fig. 4(a). The right panels show the weight vectors of the CCA algorithm. Since the eigenvector $\mathbf{v}$ [see (22)] gives an indication of the importance of the different ROIs, we multiplied the weight vectors of the ROIs with this eigenvector for a more complete visualization.

We started our algorithm with 7 cubical ROIs. The ROI X_4 has been covered by two Y-ROIs: $Y_{(1,4)}$ and $Y_{(1,9)}$ and the CCA weight vectors yield local maxima which are close to each other [see Fig. 4(b)]. After testing the similarity of their connectivity patterns (significance level = 0.05) these 2 local maxima are merged into one ROI for the next iteration [see Fig. 4(c)]. Another interesting case is the weight vector of ROI $Y_{(1,10)}$, which shows two local maxima. These caused a splitting of the ROI in two new ROIs: $Y_{(2,5)}$ and $Y_{(2,6)}$. A similar case happened in ROI $Y_{(2,1)}$ which also has been splitted in two new ROIs: $Y_{(3,1)}$ and $Y_{(3,2)}$. In the boundary between $Y_{(3,1)}$ and $Y_{(3,2)}$ we see the effect of the graph-cut minimization algorithm for assigning voxels to neighbouring ROIs. After the third iteration, there were no further splittings or mergings. Only small movements of the ROIs still happened, and at the fifth iteration the solution became stable. Finally, notice that the ROIs $Y_{(1,7)}$ and $Y_{(1,8)}$, which were not in the ground truth, had no influence on the algorithm.

B. Real fMRI Data Set

We used contrast agent-enhanced fMRI in awake monkeys, from a motion localizer study described in [40]. In the experiment, random textured patterns (50% white 0.075° dots, 50% black 0.075° dots) were presented in a circular aperture (14° diameter) either stationary or translating at five different speeds (1, 2, 4, 8, and 16°/s). The monkeys were scanned in a horizontal bore, 1.5T MR scanner (Sonata; Siemens, Erlangen, Germany) equipped with echoplanar imaging. Each functional time series
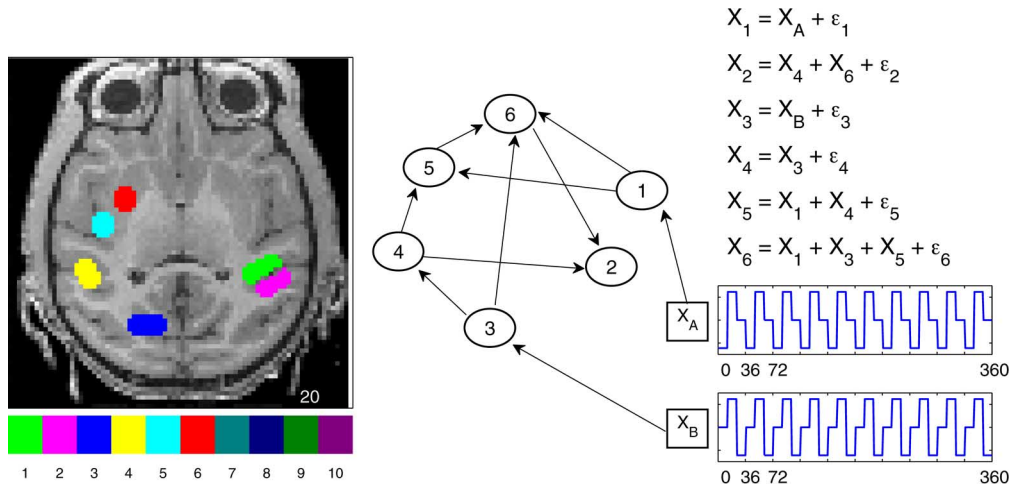
Fig. 3.   Ground truth for the synthetic data set: underlying signals and generator model (right); connected ROIs are identified on the brain slice (left).
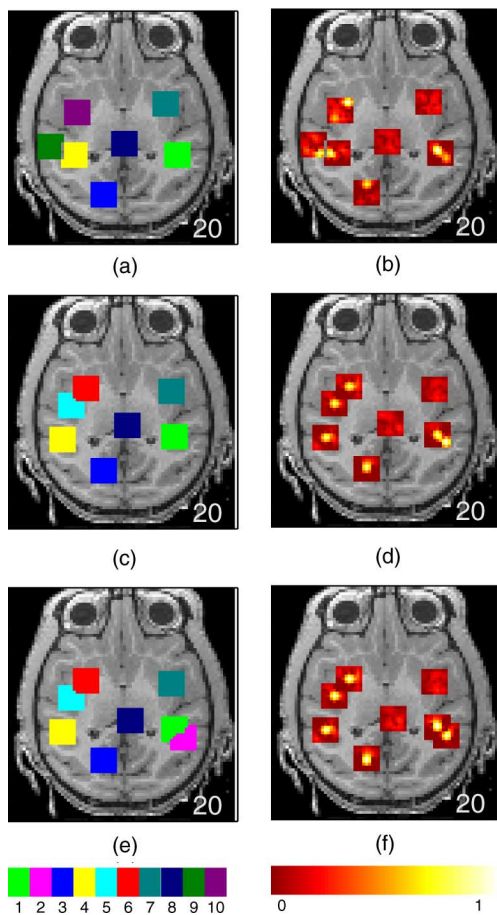

Fig. 4.   Evolution of the algorithm on a synthetic data set. The left panels show the selected voxels in each ROI, the right panels show the CCA-weight vectors. Each row represents one iteration, with the top row being the first iteration.

consisted of gradient-echo echoplanar whole brain images: repetition time (TR), 2.4 s; echo time (TE), 27 ms; 64×64 matrix; $2 \times 2 \times 2$ mm voxels (32 sagittal slices). A block design was used, with 7 conditions per run (5 speeds, 1 stationary and 1 fixation condition). Within the same run, the presentation order of conditions was repeated two times. Between the runs, the order of the conditions was randomized. 10 runs were selected (only

runs in which the monkey fixated for more than 85% of the time), which correspond to $1400 \, (= 7 \times 2 \times 10 \times 10)$ volumes. As an example, 4 time signals (obtained after convergence, see later) are shown in Fig. 8. The time signals are reordered so that that the order of the conditions is the same in all runs: (1, 2, 4, 8, and 16°/s, static, fixation). The monkey functional volumes were realigned and nonrigidly co-registered with the anatomy of a template. The brain of the template has been scanned at a spatial resolution of 1 mm^3. The monkey functional volumes were subsampled to 1 mm^3, but not smoothed.

1) Start From a Small Set of Given ROIs: We start from 5 given ROIs which are known to be involved in motion observation [40]: the dorsal and the ventral part of region V2 in the occipital lobe, namely V2d and V2v; two ROIs in the temporal lobe: MT/V5 in the lower bank of the caudal superior temporal sulcus (STS), and FST, the floor of the superior temporal sulcus; and finally, the frontal eye field (FEF) in the frontal lobe. These 5 ROIs are shown in Fig. 5(a). The algorithm converged in nine iterations to 11 new ROIs, which are shown in Fig. 5(b). Finally, Fig. 5(c) shows the CCA weight vectors in these 11 ROIs multiplied with the eigenvector **v**.

2) Start From a Full Parcellation: The final solution of the algorithm may depend on the inital set of starting seeds. Possibly interesting ROIs that are too far away from the inital ROIs may not be reached by the region-growing phase. In order to get rid of this initial ROI selection, we divided the full data on a blind manner into a lot of cubical ROIs that were only a few millimeters apart from each other, so that the voxels that were not covered by the initial ROIs are certainly reachable by the region growing phase. In this manner, 104 cubical ROIs are defined, as shown in Fig. 6(a) (color code represents labels). From these 104 initial ROIs the algorithm converged in 16 iterations to 40 final ROIs which are shown in Fig. 6(b). The evolution of the algorithm is plotted in Fig. 7. Subplot (a) shows the number of ROIs in 20 subsequent iterations. Subplot (b) shows for each iteration the value of the maximum eigenvalue on the noise-corrected correlation matrix from the representative signals, divided by the number of ROIs. This value lies between 0 and 1, i.e., 0 if all the representative signals were totally uncorrelated, and 1 if they were all the same. In subplots (c) and (d),
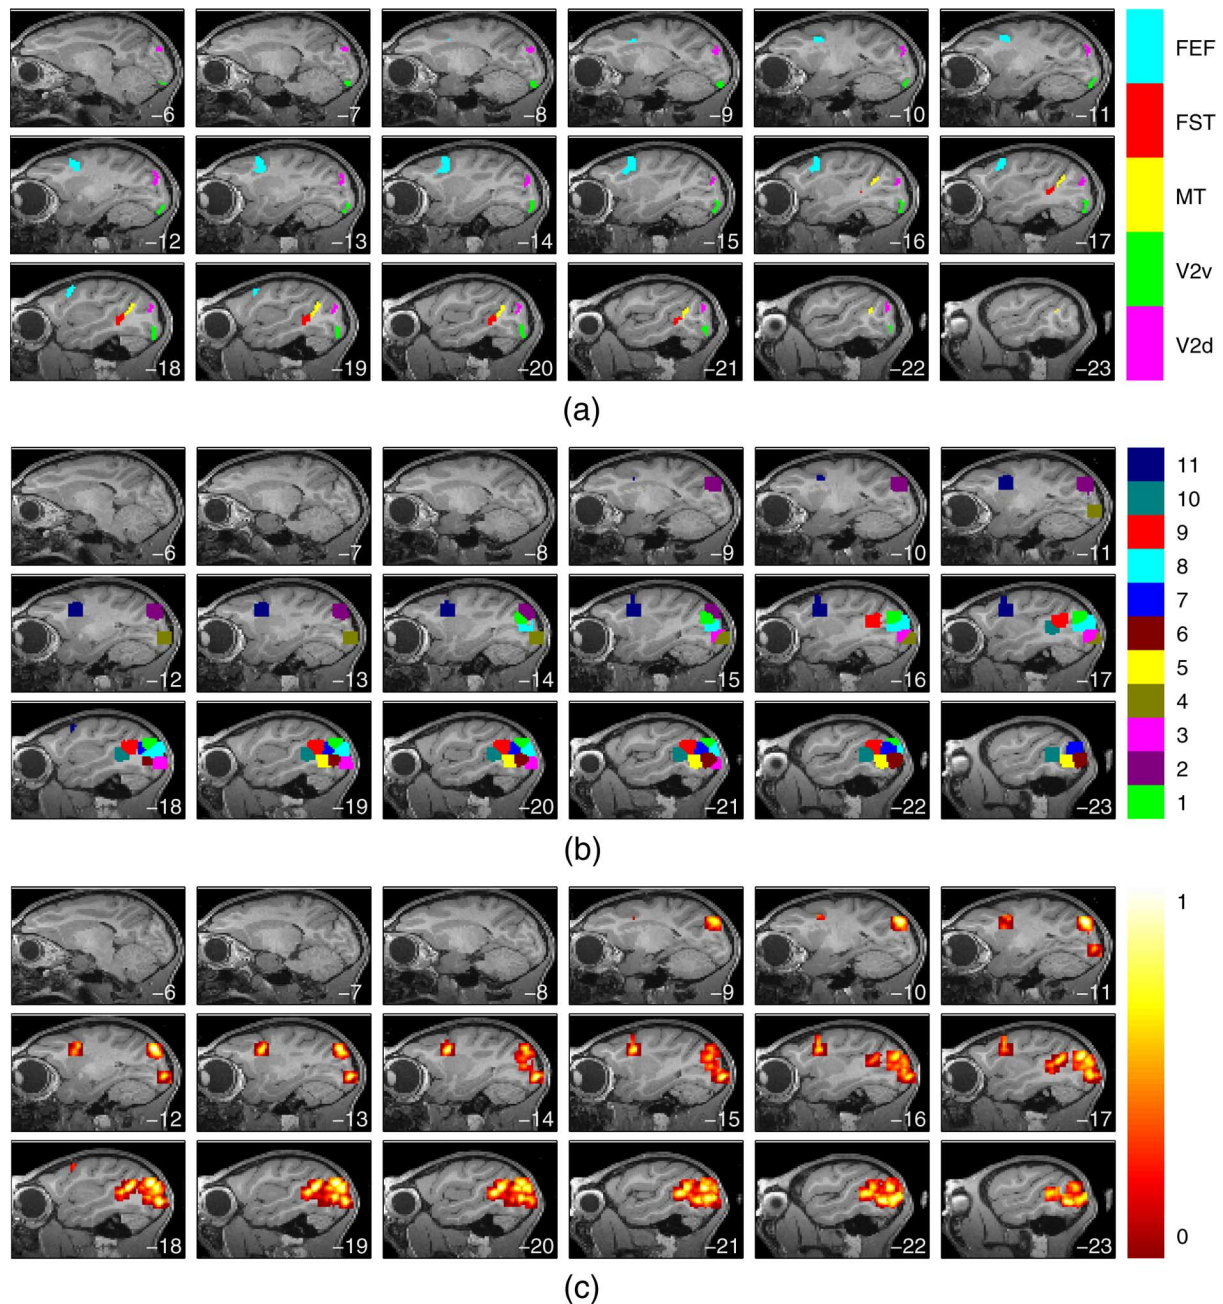
Fig. 5. (a) 5 starting ROIs; (b) 11 final ROIs; (c) CCA weight vectors in the 11 ROIs.

we show the eigenvector $\mathbf{v}$ obtained on the 104 initial ROIs in the first iteration, (c), and on the final ROIs in the last iteration, (d). The componens of $\mathbf{v}$ give a measure of the importance (the contribution) of each ROI in obtaining a maximum eigenvalue. The horizontal red line represents the value that all the components of $\mathbf{v}$ would have if they were all the same, i.e., if each ROI had the same importance. This value is equal to $\sqrt{1/n}$, with n the number of ROIs. In order to get rid of the ROIs with a low importance (low value in the eigenvector), we performed a surrogate-based test to test the hypothesis of null-importance for each ROI.

As an example, we show 4 representative time signals in Fig. 8. The top panels (a,b) represent signals that are taken from the inital set of 104 ROIs. Subplot (a) comes from ROI 41,

which is the ROI with the highest value in the eigenvector $\mathbf{v}$ [see Fig. 7(c)]. This signal shows already a clear repeating block paradigm, while the signal in subplot (b), which comes from ROI 50 and has a low value in $\mathbf{v}$, seems to represent only noise. The bottom panels (c,d) in Fig. 8 are taken from the final set of 40 ROIs. The signal in (c) comes from the ROI with the highest value in $\mathbf{v}$, while the signal in (d) comes from the ROI with the smallest value in $\mathbf{v}$. However, also this latter still seems to show a repeating block paradigm (hence, it is still in the final set of ROIs).

3) Comparison With k-Means Clustering: In this last section we compare our method with k-means clustering. In order to have a noise correction phase, as we have in our method, to correct for the spatially correlated noise, we take for each
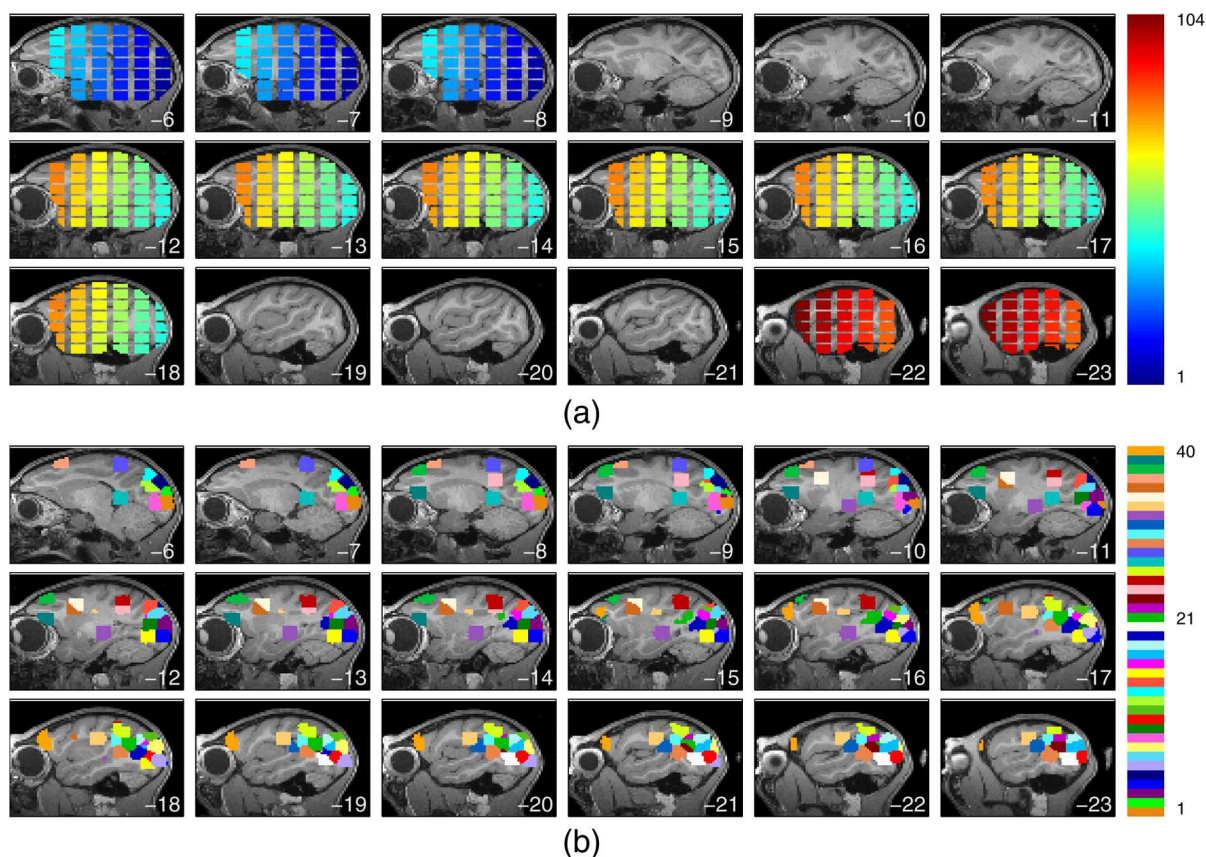
Fig. 6. Initial parcellation of the brain in 104 ROIs (a), and final parcellation in 40 ROIs (b).
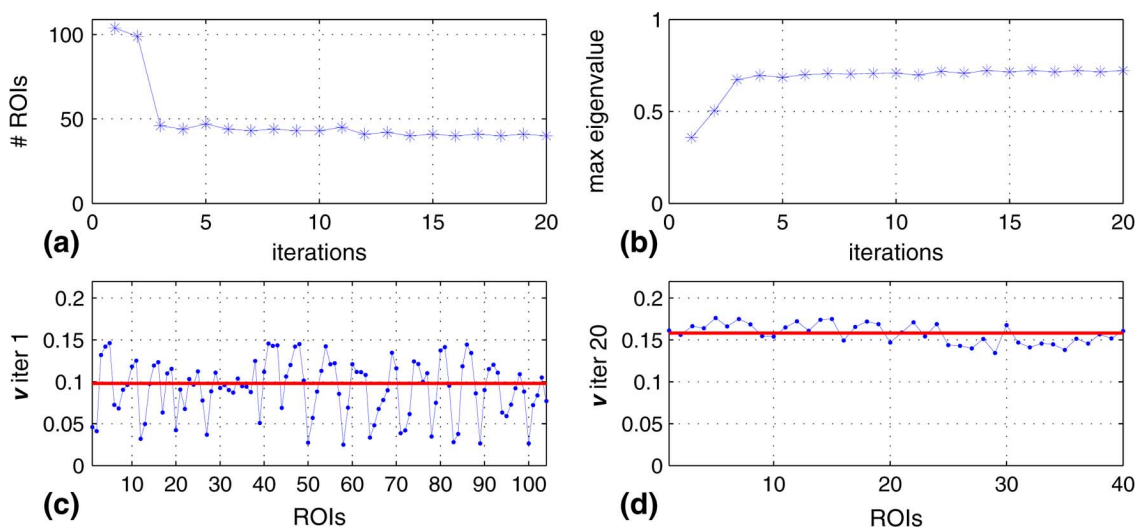

Fig. 7. Evolution of the algorithm with (a) number of ROIs, (b) the maximum eigenvalue divided by the number of ROIs, (c) the eigenvector in the first iteration, and (d the eigenvector in the last iteration.

time signal the mean time signal over the runs. Thus, the clustered time signals have length $140 \ (= 1400/10)$. In order to account for the spatial information in the data, as we also do in our method, namely in the graph-cuts optimization phase, we extend each signal with its 3 spatial coordinates, multiplied with a parameter s which controls for the weight that is given to this spatial information. Each vector (data point) for the k-means clustering has then 143 dimensions, and k-means will be per-

formed for different values of the spatial parameter s. However, due to the ill-balanced data problem in fMRI, i.e., the activated regions represent a small proportion of the brain and can be embedded in the large amount of nonactivated voxels [41], every fMRI cluster analysis usually performs an initial voxel selection procedure. As a first voxel selection procedure, we take the set of voxels that are in our final set of 40 ROIs and that have a strictly positive CCA-weight [see Fig. 6(b)]. This set consists of 8314
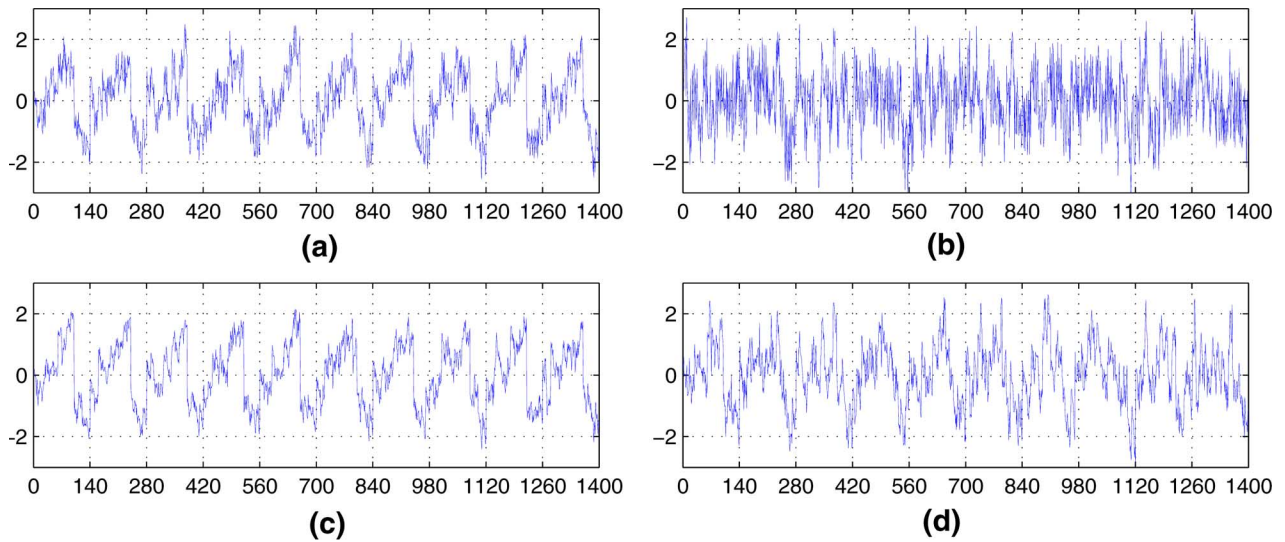
Fig. 8. Representative time signals: top panels (a), (b) are time signals obtained from the initial set of 104 ROIs, bottom panels (c), (d) are taken from the final set of 40 ROIs, left panels (a), (c) are "good" time signals (the ROI has a high value in the eigenvector), right panels (b), (d) are "bad" time signals (the ROI has a low value in the eigenvector). (a) Signal 41 from the 104 initial ROIs, (b) Signal 50 from the 104 initial ROIs, (c) Signal 14 from the 40 final ROIs, (a) Signal 29 from the 40 final ROIs.

voxels. For the number of clusters in the k-means clustering, we take the number of ROIs on convergence of our method, i.e., 40. In order to have another voxel selection which is not biased by our method, we calculate Kendall's concordance coefficient W [42] on each voxel with its 26 neighbours. Kendall's W is a measure of the intragroup homogeneity on a group of signals (here 27 signals) and ranges between 0 and 1, with $W = 0$ indicates no concordance among the signals, and $W = 1$ holds for a perfect match. We calculated Kendall's W on the original time signals of length 1400, and on the time signals that were averaged over the runs (signals of length 140) and we selected each time the 8314 voxels with the highest W-value. K-means clustering has been performed on the 3 sets of voxels, for different values of the spatial parameter s and with the number of clusters equal to 40. The results are shown in Fig. 9. The blue curves ($\diamond$) are the results on the set of voxels obtained by our method, the green curves ($\star$) are the results on the set of voxels obtained with Kendall's W on signals of length 1400, the red curves ($\square$) are the results on the set of voxels obtained with Kendall's W on signals of length 140. The horizontal axis represents the spatial parameter s. Subplot (c) shows the total intracluster error on the time signals (140 dimensions), subplot (b) shows the total intracluster error on the spatial coordinates (3 dimensions). We calculated a representative signal for each cluster as the weighted mean of its signals, where a voxel's weight is the inverse of its intracluster error to the cluster's centroid. Next, the noise-corrected correlation matrix is formed with these representative signals and the maximum eigenvalue of this correlation matrix is plotted in subplot (a) for the 3 different methods. The dotted horizontal lines in the three panels show the results when we apply the parcellation of 40 ROIs that was obtained with our method. We see that our results are comparable with those of k-means clustering with a high influence on the spatial information, and on a set of voxels obtained with a noise-corrected voxel selection procedure (Kendall's W on signals of length 140).
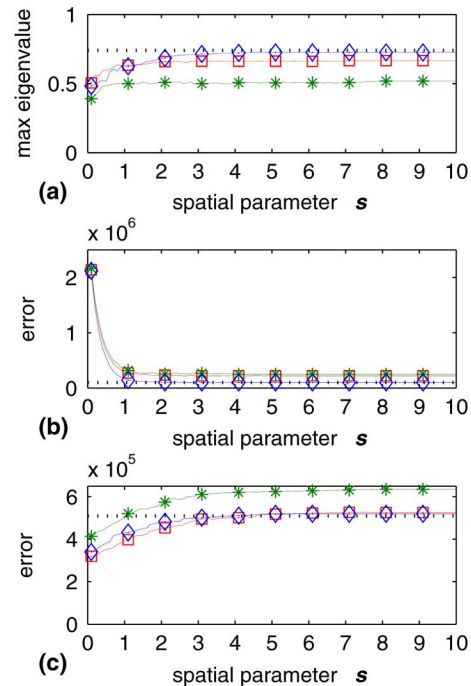

Fig. 9. Results of k-means clustering for different values of the spatial parameter s (horizontal axes) and with the number of clusters equal to 40. Three different sets of voxels are clustered: a set of voxels obtained by our method (blue curves, $\diamond$), a set of voxels obtained by Kendall's concordance without noise correction (green curves, $\star$), a set of voxels obtained by Kendall's concordance with noise correction (red curves, $\square$). Subplot (a) shows the maximum eigenvalue on the noise corrected correlation matrix of the representative signals from the different clusterings, (b) shows the total intracluster error on the spatial coordinates (3 dimensions), (c) shows the total intracluster error on the time signals (140 dimensions).

V. DISCUSSION AND CONCLUSION

We introduced a novel method for defining ROIs in fMRI data, based on their functional connectivity to other ROIs. In

our method we used first spatially regularized canonical correlation analysis for identifying maximally correlated signals, hereby accounting for spatially correlated noise, second a test for merging ROIs based on the similarity of their connectivities to other ROIs, third a graph-cuts optimization for assigning voxels to ROIs.

Several parameters are included in the method. First, the parameter κ, which controls the spatial smoothness of the CCA weight vectors. This parameter ranges between 0 and 1. With $\kappa = 0$ there is no spatial smoothing, while with $\kappa = 1$ all the weight values in one weight vector are equal and the representative signal is the unweighted mean of the original signals. In our application we used $\kappa = 0.7$. Second, the parameter K, which controls the spatial regularization in the graph-cuts optimization for assigning voxels to ROIs. In our application we test at each step a number of different values for K and determine in each ROI the voxels that are spatially connected with the local maximum (thus, we disregard the isolated voxels or islands). Next, we take in each ROI the summation of the connected signals and calculate the variance of this new signal. Finally, we select the smallest K for which the sum of the variances is maximum. The intuition behind this heuristic is that we try to get as much connected voxels as possible (in general, the more variables the higher the total variance), but without assigning bad signals to ROIs (which do not help to increase the variance). The last important parameter in the method is the threshold for region merging. This threshold is in terms of a p-value and, hence, has the classical statistical meaning. In our application, we set this p-value to 0.05.

Our method is based on functional connectivity in the sense that a ROI will only arise if its representative signal is correlated with the representative signals of other ROIs. This functional connectivity-based character occurs in two stages in the method: first, in the CCA step which seeks to find maximally (noise-correted) correlated signals; and second, in the test for region merging, which tests the similarity of connection patterns to other ROIs. This functional connectivity between ROIs is the main difference between our method and other data-driven methods for defining ROIs, i.e., cluster analysis and region growing. These 2 latter methods only take the *within-cluster* or *intracluster* performance into account, but not the *relatedness to other* ROIs. Hence, a set of neighboring voxels with highly correlated signals, but with the high correlations only due to some artefact on that specific place, will not form a ROI in our approach, since the representative signal will not be correlated with other representative signals, while it is likely to form a ROI in the other methods that look only to the intracluster performance. Our method is also related to PCA. However, in PCA on fMRI, the correlation matrix of (a subset of) the voxels' time signals is calculated and the first eigenvector of this correlation matrix is the principal component. In our method we calculate the correlation matrix of the representative signals of the ROIs (not of the individual voxels) and calculate its first eigenvector. The representative signals are generated by our method in such a way that the subsequent PCA can be performed in the most optimal circumstances, i.e., so that a maximum value of the first eigenvalue can be obtained.

Since our method for defining ROIs with corresponding representative signals is based on functional connectivity, one could argue that the results of a subsequent effective connectivity analysis (like SEM, DCM, BN, etc.) are biased. However, we can decrease this possible bias if we split the data in two parts and take one part to obtain the ROIs and the CCA weights, and next, apply these weights on the other half of the data to generate the representative signals for the further analysis.

We demonstrated our method on a synthetic and on a real fMRI data set. We showed that our method can start from a set of given ROIs, for example those that are known to be involved in a particular task, or from a fully blind parcellation of the brain to have a totally data-driven method for defining functionally connected ROIs.

REFERENCES

[1] K. Friston, C. Frith, P. Liddle, and R. Frackowiak, "Functional connectivity: The principal component analysis of large (PET) data sets," *J. Cerebr. Blood Flow Metab.*, vol. 1, pp. 153–171, 1993.

[2] M. McKeown, S. Makeig, G. Brown, T. Jung, S. Kindermann, A. Bell, and T. Sejnowski, "Analysis of fMRI data by blind separation into independent spatial components," *Human Brain Map.*, vol. 6, pp. 160–188, 1998.

[3] A. McIntosh and F. Gonzalez-Lima, "Structural equation modeling and its application to network analysis in functional brain imaging," *Human Brain Map.*, vol. 2, pp. 2–22, 1994.

[4] K. Friston, L. Harrison, and W. Penny, "Dynamic causal modelling," *Neuroimage*, vol. 19, pp. 1273–1302, 2003.

[5] X. Zheng and J. Rajapakse, "Learning functional structure from fmr images," *Neuroimage*, vol. 31, no. 4, 2006.

[6] J. Rajapakse and J. Zhou, "Learning effective brain connectivity with dynamic bayesian networks," *Neuroimage*, vol. 37, pp. 749–760, 2007.

[7] L. Harrison, W. Penny, and K. Friston, "Multivariate autoregressive modeling of fmri time series," *Neuroimage*, vol. 19, no. 4, pp. 1477–1491, 2003.

[8] A. Roebroeck, E. Formisano, and R. Goebel, "Mapping directed influence over the brain using granger causality and fmri," *Neuroimage*, vol. 25, no. 1, pp. 230–242, 2005.

[9] O. Yamashita, N. Sadato, T. Okada, and T. Ozaki, "Evaluating frequency-wise directed connectivity of bold signals applying relative power contribution with the linear multivariate time-series models," *Neuroimage*, vol. 25, no. 2, pp. 478–490, 2005.

[10] J. Sato, E. Junior, D. Takahashi, M. de Maria Felix, M. Brammer, and P. Morettin, "A method to produce evolving functional connectivity maps during the course of an fMRI experiment using wavelet-based time-varying granger causality," *Neuroimage*, vol. 31, no. 1, pp. 187–196, 2006.

[11] J. Rademacher, V. Caviness, H. Steinmetz, and A. Galaburda, "Topographical variation of the human primary cortices: Implications for neuroimaging, brain mapping, and neurobiology," *Cerebral Cortex*, vol. 3, no. 4, pp. 313–329, 1993.

[12] V. Caviness, J. Meyer, N. Makris, and D. Kennedy, "Mri-based topographic parcellation of human neocortex: An anatomically specified method with estimate of reliability," *J. Cogn. Neurosci.*, vol. 8, no. 6, pp. 566–587, 1996.

[13] G. Flandin, F. Kherif, X. Pennec, D. Riviere, N. Ayache, and J.-B. Poline, "Parcellation of brain images with anatomical and functional constraints for fMRI data analysis," presented at the IEEE Int. Symp. Biomed. Imaging, Washington, DC, 2002.

[14] R. Baumgartner, G. Scarth, C. Teichtmeister, R. Somorjai, and E. Moser, "Fuzzy clustering of gradient-echo functional MRI in the human visual cortex: Part I. Reproducibility," *J. Magn. Res. Imag.*, vol. 7, pp. 1094–1101, 1997.

[15] C. Goutte, P. Toft, E. Rostrup, F. Nielsen, and L. Hansen, "On clustering fMRI time series," *Neuroimage*, vol. 9, pp. 298–310, 1999.

[16] A. Baune, F. Sommer, M. Erb, D. Wildgruber, B. Kardatzki, G. Palm, and W. Grodd, "Dynamical cluster analysis of cortical fMRI activation," *Neuroimage*, vol. 9, pp. 477–489, 1999.

[17] P. Filzmoser, R. Baumgartner, and E. Moser, "A hierarchical clustering method for analyzing functional MR images," *Magn. Res. Imag.*, vol. 17, no. 6, pp. 817–826, 1999.

[18] Y. Lu, T. Jiang, and Y. Zang, "A split-merge-based region-growing method for fMRI activation detection," *Human Brain Map.*, vol. 22, no. 4, pp. 271–279, 2004.

[19] P. Bellec, V. Perlbarg, S. Jbabdi, M. Pelegrini-Issac, J.-L. Anton, J. Doyon, and H. Benali, "Identification of large-scale networks in the brain using fMRI," *Neuroimage*, vol. 29, no. 4, pp. 1231–1243, 2006.

[20] Y. Zang, T. Jiang, Y. Lu, Y. He, and L. Tian, "Regional homogeneity approach to fmri data analysis," *Neuroimage*, vol. 22, pp. 394–400, 2004.

[21] R. Baumgartner, R. Somorjai, R. Summers, and W. Richter, "Assesment of cluster homogeneity in fMRI data using Kendall's coefficient of concordance," *Magn. Res. Imag.*, vol. 17, no. 10, pp. 1525–1532, 1999.

[22] H. Chen, H. Yuan, L. Yao, D. ans Chen, and W. Chen, "An integrated neighborhood correlation and hierarchical clustering approach for functional MRI," *IEEE Trans. Biomed. Eng.*, vol. 53, no. 3, pp. 452–458, Mar. 2006.

[23] H. Hotelling, "Relations between two sets of variates," *Biometrika*, vol. 28, pp. 321–377, 1936.

[24] O. Friman, J. Cedefamn, P. Lundberg, M. Borga, and H. Knutsson, "Detection of neural activity in functional mri using canonical correlation analysis," *Magn. Reson. Med.*, vol. 45, pp. 323–330, 2001.

[25] O. Friman, M. Borga, P. Lundberg, and H. Knutsson, "Detection of neural activity in fmri using maximum correlation modeling," *Neuroimage*, vol. 15, no. 2, pp. 386–395, 2002.

[26] O. Friman, M. Borga, P. Lundberg, and H. Knutsson, "Adaptive analysis of fmri data," *Neuroimage*, vol. 19, no. 3, pp. 837–845, 2003.

[27] J. Rydell, H. Knutsson, and M. Borga, "On rotational invariance in adaptive spatial filtering of fMRI data," *Neuroimage*, vol. 30, pp. 144–150, 2006.

[28] K. Katanoda, Y. Matsuda, and M. Sugishita, "A spatial-temporal regression model for the analysis of functional MRI data," *Neuroimage*, vol. 17, pp. 1415–1428, 2002.

[29] G. Tononi, A. McIntosh, D. Russell, and G. Edelman, "Functional clustering: Identifying strongly interactive brain regions in neuroimaging data," *Neuroimage*, vol. 7, pp. 133–149, 1998.

[30] P. Gill, W. Murray, and M. Wright, *Practical Optimization*. London, U.K.: Academic, 1981.

[31] J. Kettenring, "Canonical analysis of several sets of variables," *Biometrika*, vol. 58, no. 3, pp. 433–451, 1971.

[32] J. Van de Geer, "Linear relations among k sets of variables," *Psychometrika*, vol. 49, no. 1, pp. 79–94, 1984.

[33] M. Dagli, J. Ingeholm, and J. Haxby, "Localization of cardiac-induced signal change in fMRI," *Neuroimage*, vol. 9, no. 4, pp. 407–415, 1999.

[34] J. Diedrichsen and R. Shadmehr, "Detecting and adjusting for artifacts in fMRI time series data," *Neuroimage*, vol. 27, pp. 624–634, 2005.

[35] T. Lund, K. Madsen, K. Sidaros, W.-L. Luo, and T. Nichols, "Non-white noise in fMRI: Does modelling have an impact?," *Neuroimage*, vol. 29, pp. 54–66, 2006.

[36] V. Kolomogorov and R. Zabih, "What energy functions can be minimized via graph-cuts?," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 26, pp. 147–159, 2004.

[37] Y. Boykov, O. Veksler, and R. Zabih, "Fast approximate energy minimization via graph cuts," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 11, pp. 1222–1238, 2001.

[38] L. Ford and D. Fulkerson, *Flows in Networks*. Princeton , NJ: Princeton Univ. Press, 1962.

[39] Y. Boykov and V. Kolomogorov, "An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 26, no. 9, pp. 1124–1137, Sep. 2004.

[40] K. Nelissen, W. Vanduffel, and G. Orban, "Charting the lower superior temporal region, a new motion-sensitive region in monkey superior temporal sulcus," *J. Neurosci.*, vol. 26, no. 22, pp. 5929–5947, 2006.

[41] M. Fadili, S. Ruan, D. Bloyet, and B. Mazoyer, "A multistep unsupervised fuzzy clustering analysis of fmri time series," *Human Brain Map.*, vol. 10, pp. 160–178, 2000.

[42] M. Kendall and J. Gibbons, *Rank Correlation Methods*. Oxford, U.K.: Oxford Univ. Press, 1990.

Filip Deleus received the M.Sc. degree in commercial engineering (business informatics) and the M.Sc. degree in artificial intelligence from the K. U. Leuven, Leuven, Belgium, in 1997 and 1999, respectively, where he is currently pursuing the Ph.D. degree in medical sciences.

Since 2000, he has been with the K. U. Leuven, Laboratorium voor Neuro- en Psychofysiologie. His research interests include computational neuroscience, neural networks, data mining, and signal processing.

Marc M. Van Hulle (SM'00) received the M.Sc. degree in electrotechnical engineering (electronics) and the Ph.D. degree in applied sciences from the K. U. Leuven, Leuven, Belgium, in 1985 and 1990, respectively, and the B.Sc.Econ. and M.B.A. degrees.

In 1992, he joined the Brain and Cognitive Sciences Department, Massachusetts Institute of Technology (MIT), Boston, as a postdoctoral scientist. From 1985–2000, he has been with the Fund for Scientific Research—Flanders, first as doctoral researcher and later as a postdoctorate and research leader. Since 2000, he has been with the K. U. Leuven, first as an Associate Professor and, since 2006, as a Full Professor. He has authored the monograph *Faithful representations and topographic maps: From distortion- to information-based self-organization* (Wiley, 2000; also translated in Japanese), and more than 175 scientific publications (http://simone.neuro.kuleuven.be). His research interests include computational neuroscience, neural networks, computer vision, data mining, and signal processing.

Prof. Van Hulle heads the Computational Neuroscience Group of the Laboratorium voor Neuro- en Psychofysiologie at the K. U. Leuven. He is co-editor of the *International Journal of Neural Systems*, associate editor of the *Computational Intelligence and Neuroscience* journal, associate editor of the IEEE TRANSACTIONS ON NEURAL NETWORKS, and editor of several special issues in signal processing and neural network journals. He is also member of the Program Committees of the Machine Learning in Signal Processing (MLSP) workshops, the International Conference on Intelligent Data Engineering and Automated Learning (IDEAL), International Symposium on Neural Networks (ISNN), International Conference on Pattern Recognition (ICPR), and the Workshop on Self-Organizing Maps (WSOM). In 2003, he received from Queen Margrethe II of Denmark the Doctor Technices degree for his work.