ELSEVIER

NeuroImage

www.elsevier.com/locate/ynimg
NeuroImage 29 (2006) 1231 – 1243

Identification of large-scale networks in the brain using fMRI

Pierre Bellec,[a,*] Vincent Perlbarg,[a] Saâd Jbabdi,[a] Mélanie Pélégrini-Issac,[a] Jean-Luc Anton,[c] Julien Doyon,[a,b] and Habib Benali[a]

[a]*Unité 678, INSERM/UPMC, 91 Boulevard de l'Hôpital, 75634 Paris Cedex 13, France*
[b]*Department of Psychology, University of Montreal, Quebec, Canada*
[c]*Centre IRMf, 264 Rue Saint-Pierre, 13385 Marseille Cedex 05, France*

Received 24 May 2005; revised 1 August 2005; accepted 25 August 2005
Available online 24 October 2005

Cognition is thought to result from interactions within large-scale networks of brain regions. Here, we propose a method to identify these large-scale networks using functional magnetic resonance imaging (fMRI). Regions belonging to such networks are defined as sets of strongly interacting regions, each of which showing a homogeneous temporal activity. Our method of large-scale network identification (LSNI) proceeds by first detecting functionally homogeneous regions. The networks of functional interconnections are then found by comparing the correlations among these regions against a model of the correlations in the noise.

To test the LSNI method, we first evaluated its specificity and sensitivity on synthetic data sets. Then, the method was applied to four real data sets with a block-designed motor task. The LSNI method correctly recovered the regions whose temporal activity was locked to the stimulus. In addition, it detected two other main networks highly reproducible across subjects, whose activity was dominated by slow fluctuations (0–0.1 Hz). One was located in medial and dorsal regions, and mostly overlapped the "default" network of the brain at rest [Greicius, M.D., Krasnow, B., Reiss, A.L., Menon, V., 2003. Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. Proceedings of the National Academy of Sciences of the U.S.A. 100, 253–258]; the other was composed of lateral frontal and posterior parietal regions.

The LSNI method we propose allows to detect in an exploratory and systematic way all the regions and large-scale networks activated in the working brain.
© 2005 Elsevier Inc. All rights reserved.

Keywords: Functional MRI; Functional connectivity; Large-scale network; Spatial noise modeling; Robust estimation

Introduction

During the past decade, investigation of cerebral activity has put more and more emphasis on the analysis of the interactions within large-scale networks of brain areas (Horwitz et al., 1999; Varela et al., 2001). It is now widely accepted that direct, indirect and stimulus-locked interactions between spatially remote brain regions can be measured by the correlation of their fMRI time series. This correlation has been called functional connectivity (Friston, 1994).

While activation analysis allows to search for regions specifically activated during a task as compared to another, functional connectivity makes it possible to explore which networks of regions are strongly interacting for a given condition, without reference to any control condition. The seminal work of Biswal et al. (1995) has introduced (functional) connectivity maps to explore the network connected with a seed region located in the primary motor cortex on resting-state data sets. A connectivity map is a three-dimensional volume whose value at each voxel is the correlation between the time series of this voxel and that of the seed region. A suitable threshold applied to the map allows to identify the network of brain regions functionally connected to the seed. It was suggested that such a network includes mostly those regions with strong anatomical connections to the seed, either direct or indirect (Xiong et al., 1999). This technique was applied on resting-state data sets for a variety of other non-motor seed regions, located in visual (Lowe et al., 1998), language (Cordes et al., 2000) and cingulate (Greicius et al., 2003) cortices, as well as subcortical regions (Stein et al., 2000). By contrast, only few studies have investigated correlation maps for subjects steadily performing a given task (Lowe et al., 2000; Greicius et al., 2003), although many techniques concentrated on comparing patterns of connectivity between tasks (McIntosh and Gonzalez-Lima, 1994; Friston et al., 1997).

Although connectivity maps have proved to be a powerful tool, the technique is not fully satisfactory. The exploration of brain functional networks relies heavily on the choice of the seed region, which allows to get insight only into the network associated with this particular seed. In addition, for seed regions picked at random, connectivity maps do not reveal meaningful cortical areas. They are rather restricted to a few voxels close to the seed region or dominated by noise, and comprise regions located in the ventricles, blood vessels or the outline of the brain. To our knowledge, no method has been proposed yet to overcome these issues.

* Corresponding author. Fax: +33 1 53 82 84 48.
E-mail address: pbellec@imed.jussieu.fr (P. Bellec).
Available online on ScienceDirect (www.sciencedirect.com).

1053-8119/$ - see front matter © 2005 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2005.08.044

Other approaches for identifying large-scale patterns of functional connectivity exist that do not rely on a seed region. These include principal components analysis (Friston et al., 1993), independent components analysis (McKeown et al., 1998) and fuzzy clustering (Baumgartner et al., 1998). These techniques were initially developed in the general context of multivariate statistics, and they optimize certain mathematical criteria, respectively: spatial decorrelation, spatial independence and intracluster homogeneity. Unfortunately, there is no clear and systematic relationship between these criteria and functional connectivity within large-scale networks of brain regions.

In this paper, we propose a new method to identify the salient large-scale networks of the human brain in an exploratory and systematic way. Our approach is based on the main acknowledged features of large-scale neural networks. According to Varela et al. (2001), neural assemblies are defined as distributed local networks transiently linked by (large-scale) reciprocal dynamic connections. In the same paper, a local network is defined as a large patch (~1 cm) of neural tissue that synchronizes its activity through the local cytoarchitecture. This definition has its roots deep back to the concept of Hebbian cell assemblies (Hebb, 1949), which are groups of entities (neurons) that act together in a coherent fashion. The same paper also defines large-scale dynamic connections as interactions based on large fiber pathways among regions that are far apart in the brain (>1 cm).

In fMRI, dynamic connections, either local or large-scale, are thought to be reflected by high temporal correlation values (Horwitz et al., 1999). The voxels belonging to a local network (region) should therefore exhibit highly correlated time series (Zang et al., 2004). Large-scale interactions moreover imply that the time series of each region in the network exhibits strong correlation with the time series of another distant region in the network. Such strong correlations between time series may however also be related to the spatially structured noise in fMRI (Cordes et al., 2002). The spatial correlations of the noise must therefore be taken into account when dealing with the practical identification of the network.

The large-scale network identification (LSNI) method proposed in this paper consists of three steps described in Theory. Firstly, the cortex is divided into disjoint and temporally homogeneous regions (Finding homogeneous brain regions). Secondly, a procedure for estimating the spatial correlations of the noise is described (Robust estimation of the spatial correlation in the noise). Thirdly, the large-scale functional connections are identified as outlier correlations between distant regions that would not occur by chance in the distribution of noise correlations (Identification of large-scale networks). On synthetic data sets, we assess that the false-positive rate of the method is controlled, and we investigate the ability of our method to identify large-scale networks with various numbers of regions and various signal-to-noise ratios (Application to synthetic data sets). We give a description of the networks on real motor data sets with a block-designed paradigm (Application to real data sets). We finally discuss the relevance and limitations of the LSNI method, as well as its potential applications (Discussion).

Theory

Finding homogeneous brain regions

The first feature of a large-scale functional network is that each region of the network should be homogeneous, i.e. composed of voxels whose time series share some similarity. We therefore aim to segment the brain cortex into a set of disjoint regions, each being a set of voxels connected with respect to 26-connexity, and such that the resulting regions are functionally homogeneous according to a certain criterion. This is achieved by means of a competitive region growing algorithm, which is an iterative procedure for image segmentation.

This procedure starts from small regions, in our case the set of all singletons of voxels in the gray matter, that grow simultaneously at each step by merging with other neighboring voxels or regions, on the basis of a similarity criterion. As functional connectivity is measured through correlation, we chose to measure the similarity between two regions as the mean correlation between the time series of any two voxels, each belonging to a different region. As all regions compete at each step for merging, it is necessary to define a merging rule. We used the mutual nearest neighbor principle (Gowda and Krishna, 1977), which is designed to merge the most similar regions at first. At each step, if the size of a region exceeds a predefined critical value t_s, it is validated and is not involved in the merging process anymore. The algorithm stops when all regions have reached the critical size, or when no more merging is possible. The size of the validated regions ranges from to t_s to $2t_s - 2$. More details on the algorithm are given in Appendix A.

By construction, the resulting regions are connected and disjoint. Moreover, they are as homogeneous as possible, according to the similarity criterion, and have a controlled size. This step is also a process of dimension reduction, as only the mean time series $\mathbf{y}_i$ of each validated region C_i will be considered for further analysis:

$$\mathbf{y}_i = \frac{1}{\#C_i} \sum_{v \in C_i} \mathbf{y}_v,\ i = 1,\ldots,N, \tag{1}$$

where $\#C_i$ is the number of voxels in region C_i, $\mathbf{y}_v = (y_{t,v})_{t=1}^{T}$ is the time series of voxel v scaled to have zero mean and unit variance, N is the number of regions identified by the region growing method and T is the number of volumes in the acquisition. The fMRI data set is thus reduced to an $N \times T$ array $\mathbf{y}$:

$$\mathbf{y} = \left(y_{t,i}; t = 1,\ldots,T, i = 1,\ldots,N\right). \tag{2}$$

Robust estimation of the spatial correlation in the noise

To assess the existence of functional interactions between brain regions, it is mandatory to take the spatial correlation of the noise into account. Main factors contributing to this correlation include partial volume effect, non-white measurement noise, preprocessing steps (such as slice-timing correction or spatial filtering), physiology-induced fluctuations and motion-related artifacts (Woolrich et al., 2004). Yet some strategies have been proposed to reduce such structured noise (Thomas et al., 2002; Perlbarg et al., 2004), some residual correlations are still to be expected.

Statistical assumptions

The correlations in the noise can be estimated if we make a spatial stationarity assumption. More precisely, assume the following model of the fMRI time series $\mathbf{y}$:

$$\mathbf{y} = \mathbf{u} + \mathbf{e}, \tag{3}$$

where $\mathbf{u}$ is the part of $\mathbf{y}$ related to neural activity, and $\mathbf{e}$ is the noise. We will assume that the noise is Gaussian, temporally independent and identically distributed (i.i.d.). The model will be extended to

the case of time-correlated noise in Identification of large-scale networks. Formally, $\mathbf{e} = (e_{t,i})_{i=1}^{N}$, for t from to 1 to T, are some independent samples from a multivariate Gaussian random variable $\mathbf{E} = (E_i)_{i=1}^{N}$, with mean $\boldsymbol{\mu} = (\mu_i)_{i=1}^{N}$ and variance–covariance matrix $\boldsymbol{\Sigma} = (\sigma_{i,j})_{i,j=1}^{N}$. The spatial correlations in the noise are assumed to be stationary, i.e. they only depend on the spatial distance between regions:

$$\sigma_{i,j} = \sigma_i \sigma_j \rho(||i-j||), \tag{4}$$

where σ_i and σ_j are positive standard deviations, $||i-j||$ is a shortcut for the spatial distance (or lag) between regions and C_i and C_j, for example the Euclidian distance between their centroids, and ρ, the spatial correlogram, is a real-valued function that satisfies $\rho(0) = 1$ and is bounded by -1 and 1. Such a spatial model is valid if and only if the variance–covariance matrix $\boldsymbol{\Sigma}$ is positive-definite (Cressie, 1993).

Robust estimation of the spatial correlogram

Estimating the correlation structure of such a temporally i.i.d. and spatially stationary noise simplifies to estimating the spatial correlogram $\rho(h)$ for any spatial lag h. As we do not know what regions are predominantly influenced by the noise, the whole set $D_h = \{(C_i, C_j) | \; ||i-j|| = h\}$ of pairs of regions at lag h is considered for this estimation. However, the presence in D_h of functionally connected regions is expected to bias the estimation. As long as there are few such regions, we address this issue by using the following robust estimate, based on the median:

$$\hat{\boldsymbol{\rho}}(h) = \mathrm{F}^{-1}\left(\mathrm{median}\left\{\mathrm{F}(r_{i,j}), (C_i, C_j) \in D_h\right\}\right), \tag{5}$$

where F is the Fisher transform,[1] and $r_{i,j}$ is the Pearson's linear correlation between the mean time series of the two regions C_i and C_j. Under an asymptotic assumption, the estimate $\hat{\boldsymbol{\rho}}(h)$ is unbiased. This estimation is moreover "robust": it is stable in the presence of outliers, which here are the very regions of the functional network.

In practice, in order to increase the number of observations and consequently the statistical power, we do not estimate $\hat{\boldsymbol{\rho}}(h_k)$ of the spatial correlogram for these spatial lags h_k, $k = 1, \ldots, K$:

$$\hat{\boldsymbol{\rho}} = (\hat{\boldsymbol{\rho}}(h_1), \ldots, \hat{\boldsymbol{\rho}}(h_k), \ldots, \hat{\boldsymbol{\rho}}(h_K)). \tag{6}$$

Large and local scales

Once the correlation of the noise is estimated, it is possible to study its features, so as to define a local scale and a large scale in a data-driven way. However, the empirical estimator $\hat{\boldsymbol{\rho}}$ of the correlogram does not provide in general a positive-definite correlation matrix. We therefore restrict ourselves to a parametric class of valid correlograms, which was selected on the basis of the visual inspection of the empirical values $\hat{\boldsymbol{\rho}}$ estimated on real fMRI data sets. We use the following rational-quadratic model $\rho_{\boldsymbol{\theta}}(h)$ (Cressie, 1993):

$$\begin{cases} \rho_{\boldsymbol{\theta}}(0) = 1, \\ \rho_{\boldsymbol{\theta}}(h) = 1 - \theta_1 - \theta_2 \dfrac{h^2}{1 + h^2/\theta_3}, \forall h > 0, \end{cases} \tag{7}$$

where $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3)$ is a parameter vector of three non-negative real values satisfying $\theta_1 + \theta_2\theta_3 < 1$. Such a parametric expression defines a class of valid correlograms. This type of model decreases rapidly from a correlation level between neighboring voxels ρ_{0+}, which defines a local-scale correlation, towards an asymptotic level of correlation ρ_∞, which defines a large-scale correlation, and it is possible to define a critical distance h_∞^ε beyond which the correlogram is almost equal to the asymptote, with a tolerance of ε, typically 0.01 (see Fig. 1). Formally, a rational-quadratic model can be parameterized using $(\rho_{0+}, \rho_\infty, h_\infty^\varepsilon)$ as detailed in Appendix B.

In the perspective of the large-scale network identification, we need to define the large scale or, equivalently, the critical distance h_∞^ε. We thus determine the parameter set $\boldsymbol{\theta}^*$ such that $\boldsymbol{\rho}_{\boldsymbol{\theta}^*}$ fits best the empirical values $\hat{\boldsymbol{\rho}}$. The practical optimization procedure is detailed in Appendix C. On the basis of these optimal parameter values θ^*, a pair of regions is considered as being on a large scale as soon as the distance between their centroids is higher than $h_\infty^\varepsilon = \sqrt{(\theta_2^*\theta_3^*/\varepsilon - 1)\theta_3^*}$ (see Eq. (B.3)).

Identification of a large-scale network

Once the distribution of noise correlations has been estimated, it is possible to search for outlier correlations between distant regions, which are very unlikely to be due to noise. The set of regions that exhibit an outlier interaction with at least one distant region will compose the large-scale network.

Let us consider one large-scale interaction between regions C_i and C_j at a distance $h = ||i-j||$ larger than h_∞^ε. We wish to test whether the correlation $r_{i,j}$ is likely to be found only by chance in the noise distribution. Under the null hypothesis that the expected value of $r_{i,j}$ equals the expected value of the noise $\rho(h)$, the following quantity is asymptotically normally distributed with zero mean and unit variance (Anderson, 1958):

$$z_{i,j} = \frac{\mathrm{F}(r_{i,j}) - \mathrm{F}(\rho(h))}{\sigma_F}, \tag{8}$$

with $\sigma_F^2 = T - 1$ in the time i.i.d. case.

A departure from the temporally i.i.d. assumption due to a positive temporal autocorrelation will result in an increase in asymptotic variance σ_F^2 (Rodriguez, 1982). To correct such possible bias, we estimate the standard deviation of $\mathrm{F}(r_{i,j})$. This standard deviation with T observations of temporally correlated noise can be thought of as the standard deviation σ_F for an

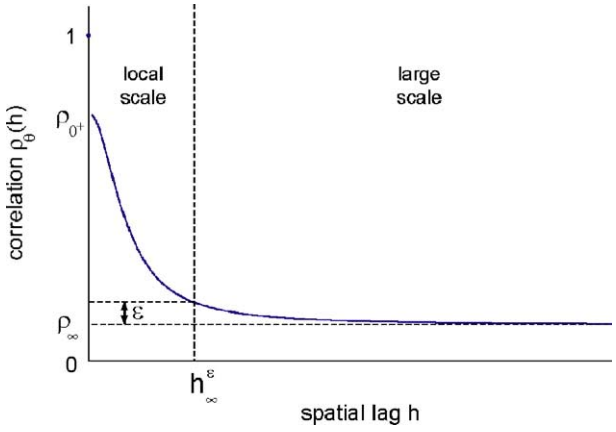

Fig. 1. Rational-quadratic correlogram: the correlation between the time series of two regions is a parametric function of the spatial distance between the regions. Graphical interpretation of the parameters ρ_{0+}, ρ_∞ and h_∞^ε is shown in the figure.

[1] $F(x) = \frac{1}{2}\log\left(\frac{1+x}{1-x}\right)$. $F(r_{i,j})$ has an asymptotic normal distribution with mean, or median, $F(\rho(h))$ and variance $T-1$ (Anderson, 1958).

equivalent number T_{eff} of effective independent observations, as was proposed in the context of the general linear model (Worsley and Friston, 1995). To estimate the standard deviation, we resort to a robust estimate, the median absolute deviation (Rousseeuw and Croux, 1993):

$$\hat{\sigma}_F = 1.48\,\text{median}\left\{|\text{F}(r_{i,j}) - \text{F}(\hat{\rho}(h))|, (C_i, C_j) \in D_h, h \geq h_\infty^\varepsilon\right\}. \quad (9)$$

We replace $\rho(h)$ and σ_F by their estimates in Eq. (8), and consider the following quantity:

$$\hat{z}_{i,j} = \frac{\text{F}(r_{i,j}) - \text{F}(\hat{\rho}(h))}{\hat{\sigma}_F}. \quad (10)$$

On simulations, the distribution of $z_{i,j}$ in the presence of both time and space correlations could be accurately approximated by the distribution of $z_{i,j}$ in the absence of both space and time correlations. This latter distribution is deduced from the known distribution of $F^{-1}(\sqrt{T-1}z_{i,j}) = r_{i,j}$ for $\rho_{i,j} = 0$ (Anderson, 1958):

$$\sqrt{T-2}\frac{r_{i,j}}{\sqrt{1-r_{i,j}^2}} \sim t_{T-2}, \quad (11)$$

where t_{T-2} is the Student's t distribution with $T-2$ degrees of freedom.

Finally, a region C_i is included in the large-scale functional network as soon as there exists a region C_i, with $\|i-j\| \geq h_\infty^\varepsilon$, such that $z_{i,j}$ is significantly different from the expected value in the noise. We have to perform a number of tests equal to the number M of pairs of regions such that the distance between them is larger than h_∞^ε. In order to correct this multiple testing problem, we apply a Bonferroni correction, which is exact for independent tests and is otherwise too conservative. If p is the type I error (false-positive rate) of the whole procedure, the Bonferroni correction consists in performing each test with a type I error p / M. We checked on simulations that the empirical false-positive rate p was indeed controlled, see Application to synthetic data sets.

Application to synthetic data sets

Synthetic data sets

All simulations used the partition of the gray matter into brain regions, namely that obtained on real data sets for subject 1 (see Data analysis). Time series with the same parameters as those of the real data sets (i.e. $T = 128$, TR = 2.33 s) were simulated in these regions with the following two different sampling procedures.

For the first simulation ($S1$), synthetic data sets met the null hypothesis of a spatially stationary noise with no large-scale interactions. The time series were composed of Gaussian noise, time- and space-correlated, generated as follows. The space correlation was based on a rational-quadratic parametric model, with a critical distance $h_\infty^{0.01} = 40$ mm and various parameters ρ_{0+} and ρ_∞. Temporally independent Gaussian data sets with this space correlation were generated as described in Cressie (1993), pp. 201–203, using a Cholesky decomposition of the space correlation matrix and Gaussian random samples generated using Matlab®.[2] The time series were then convolved with a Gaussian kernel, with various Full-Width-at-Half-Maximum (FWHM), resulting in a separable space–time correlation structure. The following parameters were used:

($S1$)a: Spatial correlations close to those estimated on real data sets (Results), small temporal correlation: $\rho_{0+} = 0.1$, $\rho_\infty = 0.001$, FWHM = 1 s.

($S1$)b: Spatial correlations close to those estimated on real data sets (Results), large temporal correlation: $\rho_{0+} = 0.1$, $\rho_\infty = 0.001$, FWHM = 5 s.

($S1$)c: Large local spatial correlation, low asymptotic spatial correlation, large temporal correlation: $\rho_{0+} = 0.3$, $\rho_\infty = 0.001$, FWHM = 5 s.

($S1$)d: Large local and asymptotic spatial correlation, large temporal correlation: $\rho_{0+} = 0.3$, $\rho_\infty = 0.1$, FWHM = 5 s.

A set of synthetic data sets was generated for each parameter set ρ_{0+}, ρ_∞, FWHM, as described in ($S1$)a–d, resulting in 6000 synthetic data sets in total.

For simulation ($S2$), the noise was both time- and space-correlated, generated as in simulation ($S1$). The parameters of the spatial correlation were set close to those estimated on real data sets (Results), with a large temporal correlation: $h_\infty^{0.01} = 40$ mm, $\rho_{0+} = 0.1$, $\rho_\infty = 0.001$, FWHM = 5 s. In addition, some randomly selected regions belonged to a functional network. In a basic attempt to model such a functional network, the time series of these selected regions were the sum of the noise and a single time series, namely a given i.i.d. Gaussian time series. As it is common to all regions of the functional network, this temporal trend induces a high correlation of the time series of the regions that belong to the network, regardless of the spatial distance between them. The empirical variances of the noise s_{n}^2 and the signal s_{s}^2 were set so as to obtain various signal-to-noise ratios (SNR), defined as follows:

$$SNR(\text{dB}) = 10\,\log_{10}(s_{\text{s}}^2/s_{\text{n}}^2).$$

We first investigated the influence of the SNR, with three different values: −1.5 dB, −0.5 dB and 0 dB (corresponding respectively to and $s_{\text{s}}^2 \simeq 0.41$, 0.47, 0.5 and $s_{\text{n}}^2 \simeq 0.59$, 0.53, 05). For each SNR, we also tested different sizes of the network: 5%, 10%, 20%, 30%, 40% and 50% of the total number of regions. We generated 250 synthetic data sets for each combination of SNR and network size, resulting in 4500 synthetic data sets in total.

Methods

We used simulations ($S1$) and ($S2$) to investigate the specificity and sensitivity of the LSNI method, respectively. For each data set of simulation ($S1$), we applied the steps described in Robust estimation of the spatial correlation in the noise and Identification of large-scale networks to find significant large-scale interactions, with various expected false-positive rates $p = 0.001$, 0.01, 0.05, 0.1. As these data sets are only composed of noise, the empirical false-positive rate $\hat{\rho}$ is the ratio between the total number of large-scale interactions detected as significant and the total number of synthetic data sets. For each synthetic data set of simulation ($S2$), the sensitivity of the method was estimated as the proportion of regions in the functional network detected by the method with $p < 0.05$. The results were averaged over the 250 synthetic data sets for each combination of SNR and network size.

[2] http://www.mathworks.com.

Results

For all sets of parameters ($S1$)a–d, Fig. 2 shows the empirical false-positive rates $\hat{\rho}$ against the expected value p. The magnitude of $\hat{\rho}$ was of the order of the expected value p for any set of parameters ($S1$)a–d. Increasing time or space correlation did not appear to bias the control of the false-positive rate on these simulations.

The sensitivity appeared to increase for networks involving from 5% up to 20% of the total number of regions, regardless of the SNR (see Fig. 3). Since a region of the network is detected as soon as one large-scale interaction is detected, the larger the network, the easier it is to detect each region, which explains such an increase in sensitivity. For networks larger than 30% of the cortex, the sensitivity sharply dropped and was zero for networks involving at least half of the regions. This is due to an increasing bias in the estimation of the structure of the noise, related to the presence of the functional network.

Moreover, we found that the method performed well with a SNR value greater than 0 dB but performed poorly, with a sensitivity lower than 50%, for a SNR value lower than −1 dB.

Application to real data sets

Real data sets

Four right-handed male healthy volunteers (age: 25 to 27 years) participated in an fMRI study of motor sequence learning. The protocol was approved by the local ethic committee. One functional run was analyzed for each subject. In this run, subjects were alternating a control task and a motor sequence task, following a conventional 30-s-long block-designed paradigm. The control task consisted of looking at a fixation cross displayed on a screen. The motor sequence task consisted of pressing the keys of an MRI-compatible keyboard with the left (non-dominant) hand, each key corresponding to one finger, following a given sequence (4,1,3,2,4, numbering the fingers from the index to the little). This sequence was performed at a fixed, comfortable rate of 2 Hz paced by an acoustic beeper. During the execution of the task,

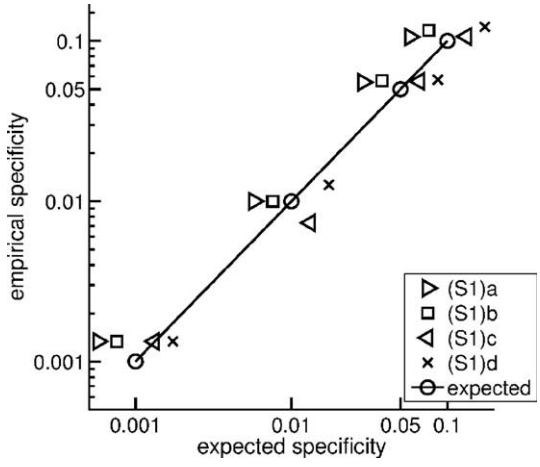

Fig. 2. False-positive rate for simulation ($S1$) meeting the null hypothesis of no large-scale interactions. The empirical false-positive rate for 1500 synthetic data sets is compared to the expected value with different parameters for the space and time correlations ($S1$)a–d.

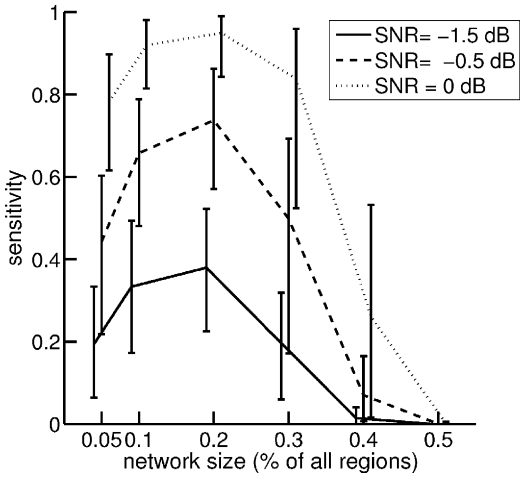

Fig. 3. Sensitivity of the LSNI procedure for functional networks involving perfectly correlated regions. The error bars indicate the 95%-percentile intervals for the estimated sensitivity over all 250 data sets.

the following text “Sequence A: 4,1,3,2,4” was displayed on a screen.

For each functional run, blood oxygenation level dependent signals were recorded in 136 volumes, which were acquired using a single-shot echo-planar imaging sequence (TR/TE = 2333/30 ms, 64 × 64 matrix, 42 contiguous slices, FOV = 192 mm × 192 mm, slice thickness 3 mm and flip angle = 81°) on a Bruker 3.0 T MRI scanner at the fMRI Center in Marseille.[3] A high-resolution T_1-weighted scan was also acquired using the following MPRAGE sequence: TR = 11.6 ms, TE = 5.67 ms, TI = 800 ms, 256 × 192 × 104 matrix, FOV = 256 mm × 230 mm × 182 mm and flip angle = 30°.

Data analysis

Preprocessing

On the one hand, the functional data were corrected of slice timing and checked for motion using SPM99,[4] and then corrected of second-order polynomial time drifts. No correction of small movements was applied as such correction could itself induce some correlation in the time series to an unknown degree. To improve the spatial stationarity of the noise, we reduced the noise with a highly preponderant influence in well-localized brain areas using a dedicated statistical procedure (Thomas et al., 2002; Perlbarg et al., 2004). This procedure started with a spatial independent component analysis (ICA) of the fMRI data using the infomax algorithm[5] (Bell and Sejnowski, 1995), initialized using all components of a principal component analysis. The lateral ventricles and basilar artery were then manually segmented on the structural image of each subject. This segmentation was used to extract time series predominantly influenced by cardiac, respiratory and movement-related noise, and clearly unrelated to neural activity. These time series were used to identify the independent components related to physiological noise, whose contribution was removed from the fMRI data set.

On the other hand, we segmented the gray matter from the anatomical image using SPM99. This segmentation was considered

[3] http://irmfmrs.free.fr/.
[4] http://www.fil.ion.ucl.ac.uk/spmspm99.html.
[5] http://www.sccn.ucsd.edu/fmrlab/.

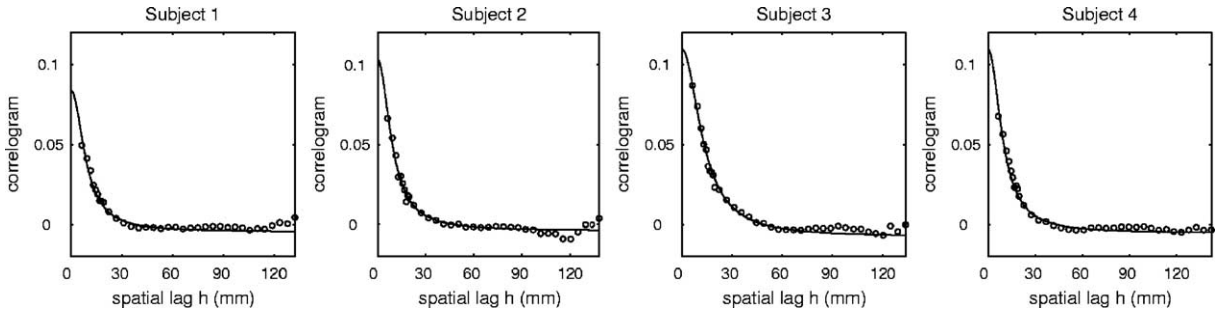

Fig. 4. Empirical (circles) and fitted (solid line) correlograms for the four subjects.

as a mask of interest for further analysis. The anatomical images were also spatially normalized in the standard stereotaxic space of the Montreal Neurological Institute (MNI) using SPM99, and the same transformation was applied to the regions identified by the LSNI procedure.

Identification of large-scale networks

All procedures were implemented using Matlab® 7.0. The computation times indicated in Results were obtained for a workstation with an Intel™Xeon® 2.4 GHz processor and 1 gigabyte of memory. We applied the LSNI method as described in Theory: firstly, we ran the region growing algorithm on the mask of interest with a critical size of $t_s = 10$ voxels; we then computed the empirical correlograms for each subject, and identified significant large-scale interactions with a false-positive rate $p < 0.05$.

Post-processing

Instead of describing individually all the regions detected by the LSNI method, we grouped these regions into distinct functional networks to simplify their visualization and interpretation. This was achieved using a hierarchical clustering of the time series of the regions with the Ward criterion as implemented in Matlab®. The hierarchy was thresholded in an arbitrary way, in order to obtain a minimal number of clusters while roughly preserving functional coherence. To clarify which phenomenon gave rise to the connectivity within a given cluster, we computed the first principal component of the time series associated with the regions of the cluster, as well as the Gaussian window Fourier transform of this component.[6]

Results

The region growing procedure assigned most of the voxels (range: 94% to 95% of the gray matter) to a homogeneous region whose size varied from 10 to 18 voxels. The number of regions was comparable from one subject to another (range: 1566 to 1767). This first step required the highest computation time (22 to 29 min).

The shape of the empirical correlograms was very reproducible, see Fig. 4, suggesting that a rational-quadratic spatial model was well adapted. For spatial lags smaller than one half of the maximum lag, the fit of the spatial model was very accurate. For the largest lags, small departures from the fitted model were observed, mainly due to the decreasing number of pairs of regions used in the estimation. When computing the correlograms, the 10% largest lags were not considered, as there were few pairs of such distant regions, which made an accurate estimation of the correlogram difficult. Estimated local and asymptotic correlations only varied slightly across subjects: local correlation ρ_{0+} ranged from 0.08 to 0.11; asymptotic correlation ρ_∞ was found to be almost zero for all subjects (range -0.01 to 0). The critical distance h_∞^ε ranged from 25 mm to 43 mm. Estimation and fit of the correlogram took from 21 s to 29 s by subject.

Once the correlation of the noise had been estimated, the large-scale network was identified. This step required a little computation time (7 to 9 s). There was an important intersubject variability concerning the size of the large-scale networks: 73, 134, 205 and 171 regions for subjects 1 to 4, respectively. Fig. 5 shows, for each subject, the correlation matrix between the time series of all the regions belonging to the large-scale network. Reordering the correlation matrix using hierarchical clustering evidenced a clear organization into functional clusters (Fig. 5a). The diagonal squares associated with each cluster are shown in Fig. 5b. The spatial location of each cluster numbered on Fig. 5 is described hereafter, as well as the first principal component associated with each cluster, which explained from 30% to 55% of the variance of the time series in the cluster.

Cluster of activated regions

For all subjects, the first principal component associated with one of the clusters was strongly correlated with the expected response to the block-designed stimulus, as modeled with SPM99 using a box-car stimulus convolved by the standard hemodynamic response function. This correlation was 0.85, 0.84, 0.74 and 0.76 for subjects 1 to 4, respectively. The cluster was the largest in size for subjects 1 to 3 (number 1, see Fig. 5b), and the third largest in size (number 3, see Fig. 5b) for subject 4. Some regions were systematically found for all subjects (see Fig. 6a), such as the right primary motor and sensorimotor cortices, bilateral premotor cortex ($z = 50$), bilateral occipital cortex ($z = -8$) and cerebellar cortex ($z = -20$). Other regions could be found for at least 2 subjects, such as bilateral supplementary motor area ($z = 50$), and bilateral insula ($z = 16, 20$). Fig. 6b shows that the first principal component of each cluster was indeed clearly related to the block-designed paradigm. The time-frequency distribution of this component revealed a steady power at the fundamental frequency of the stimulation ($1/60 \simeq 0.016$ Hz) (Fig. 6c).

Cluster of parieto-frontal regions

For all subjects, one cluster was predominantly located in the lateral frontal and posterior parietal cortices. This cluster was the second largest in size for all subjects (number 2, see Fig. 5b). The spatial location of this cluster is displayed in Fig. 7a. Regions were

[6] http://www-stat.stanford.edu/~wavelab/.

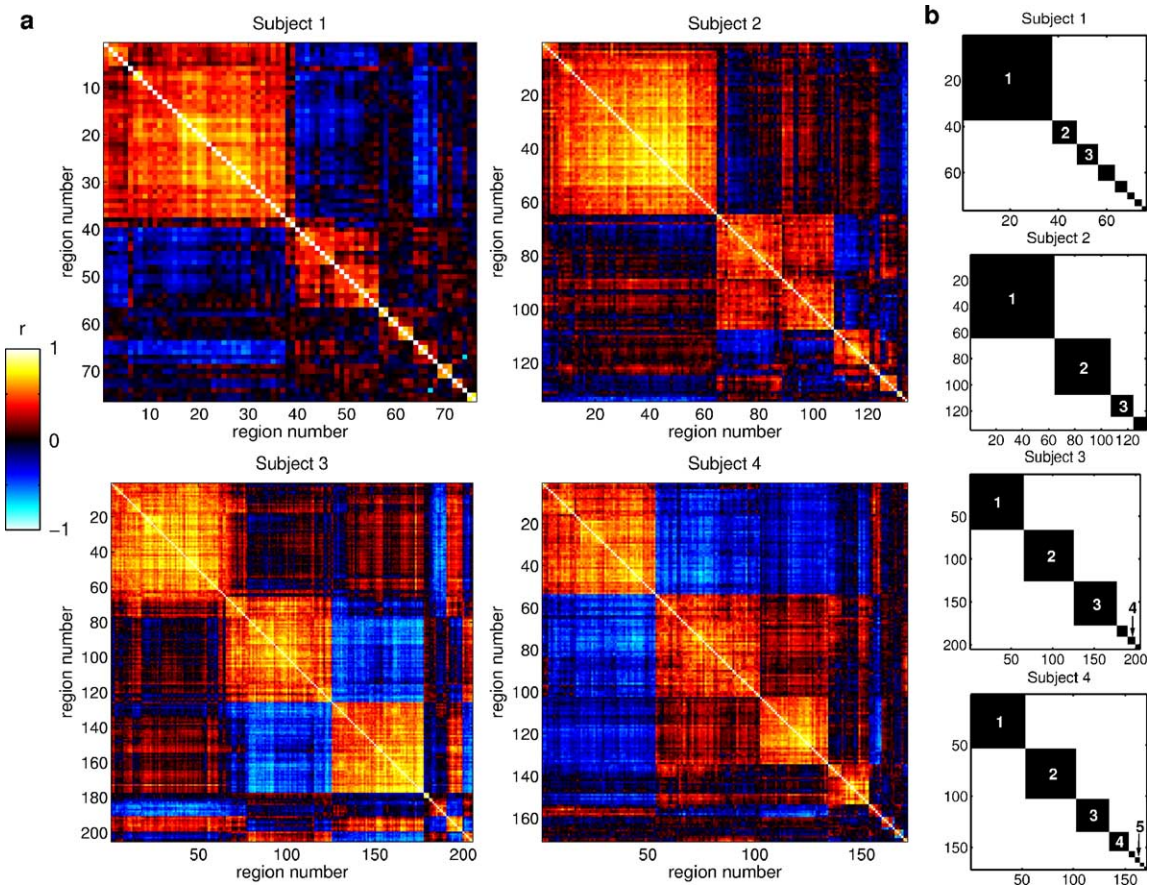

Fig. 5. (a) Correlation matrix of the time series associated with the regions of the large-scale network of each subject. (b) Regions were grouped into distinct functional clusters. Black squares show the within-cluster correlations inside the matrix of panel a.

systematically found bilaterally in the superior parietal cortex (z = 45, 50), and in the lateral frontal cortex (z = −7, 25, 50). There were additional regions in the cerebellar cortex (z = −20) and bilaterally in the temporal cortex for three out of four subjects. The first principal component of each cluster had a low correlation with the expected response to the stimulation: −0.22, 0.1, −0.02, 0.35 for each subject, respectively (Fig. 7b). The time-frequency distribution of this component did not reveal any steady power at a given frequency, but the power spectrum was steadily located in the [0,0.1] Hz band, and highly predominant in the [0,0.05] Hz band (Fig. 7c).

Cluster of medio-posterior regions

For all subjects, one cluster was predominantly located medially in the cingulate/precuneus and frontal cortices, as well as in posterior regions. This cluster was the third largest in size for subjects 1 to 3 (number 3, see Fig. 5b), and the largest in size for subject 4 (number 1, see Fig. 5b). Fig. 8a shows the spatial location of this cluster. Some regions were repeatedly found in the posterior and medial regions, in the precuneus (z = 43, 50) and in the posterior cingulate cortex (z = 24, 28) for subjects 2, 3, 4. Moreover, some regions were also found in the medial part of the frontal cortex (all slices) both ventrally and dorsally except for subject 1, who only had dorsal regions. There were additional regions in the superior and inferior parietal cortices (z = 28, 43, 50), and in the occipital cortex (z = 24). The first principal component had low or moderate correlations with the expected response to the stimulation: −0.04, 0.13, 0.33, −0.52 for each subject, respectively (Fig. 8b). The time-frequency distribution patterns of this component were very similar to those of the cluster of parieto-frontal regions, with a power spectrum steadily predominant in the [0,0.05] Hz band (Fig. 8c).

Other clusters

The remaining clusters were small and shared no clear commonalities between subjects. A few typical examples of these clusters are displayed in Fig. 9. Cluster 4 of subject 4 (Fig. 5b) had a very bilateral spatial location (Fig. 9a) involving frontal, ventral premotor and anterior cingulate cortices (z = 25, 34), as well as left anterior cerebellar cortex (z = −8) and cingulate cortex (z = 6) (Fig. 9a, first row). The first principal component had a time-frequency distribution similar to that of the large non-activated clusters (Fig. 9c, first row). Moreover, there was no clear strong correlation between this cluster and any other one. Such a small and isolated cluster might have been missed by the LSNI method for the other subjects. Cluster 4 of subject 3 (Fig. 5) is another small cluster with a medial spatial location (Fig. 9a, second row). The power spectrum of the first principal component was steadily located in the 0–0.1 Hz band (Fig. 9c, second row). This cluster had strong correlations with both clusters 1 and 3 (Fig. 5b), which shows the limits of functional clustering: as clusters 1 and 3 were only weakly correlated, it was therefore impossible to define a large homogeneous cluster comprising both clusters 1, 3 and 4. Cluster 5 of subject 4 was an example of a very small cluster

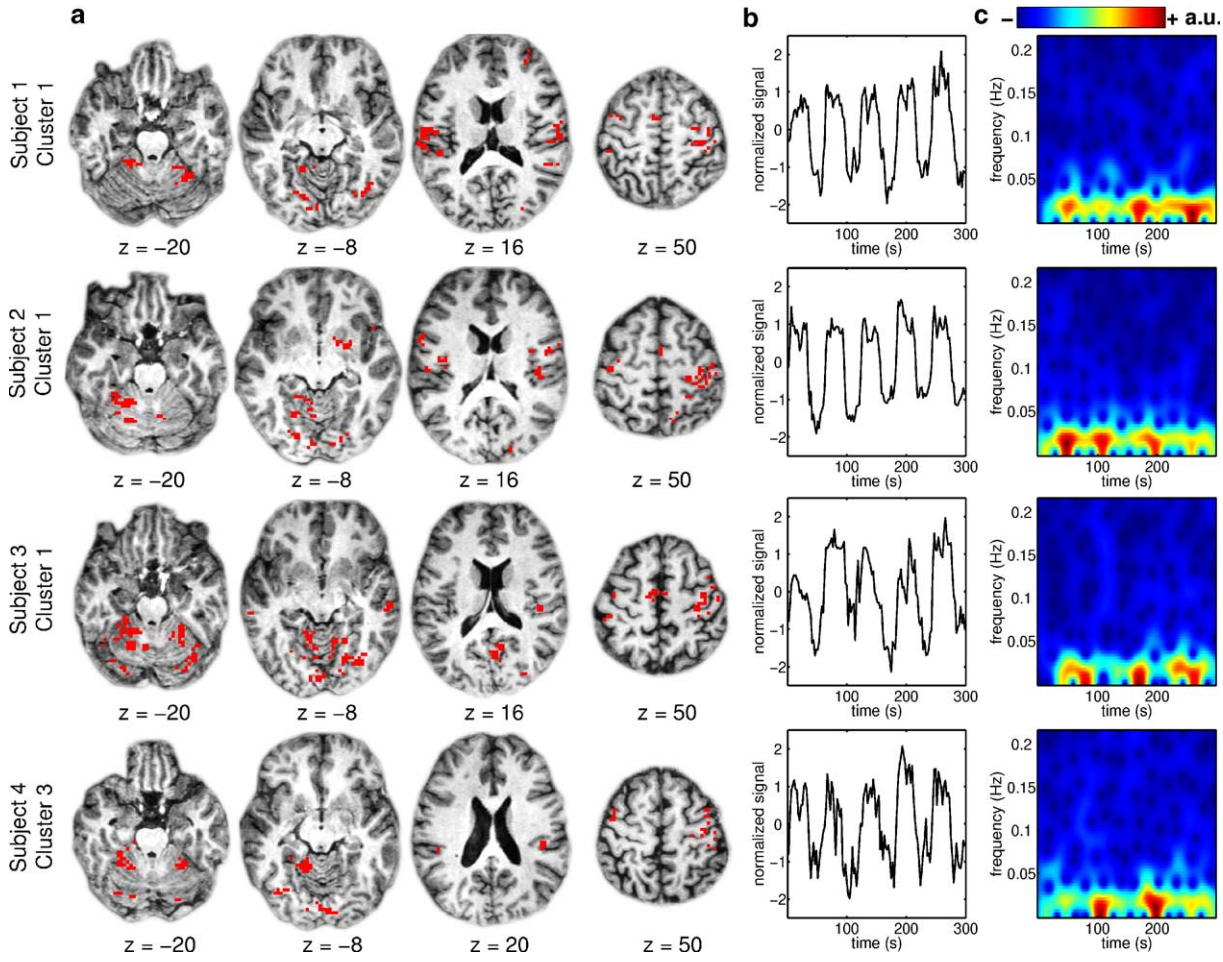

Fig. 6. Cluster of activated regions: regions of this cluster have a stimulus-locked activity, and are mainly located in the motor system. (a) Spatial location for each subject, superimposed on the corresponding axial slice of the anatomical image (left on the image is left of the brain). (b) First principal component of the time series associated with the cluster. (c) Gaussian window Fourier transform of the first principal component (arbitrary units).

located near large blood vessels (Fig. 9a, third row). The power spectrum of the first principal component was spread over all the frequencies (Fig. 9c, third row). The time-frequency distribution associated with this cluster is a strong indication that it is dominated by residual physiological noise. It did not show any strong correlation with other clusters of the large-scale network.

Discussion

The LSNI method proposed in this paper consists of three steps: definition of homogeneous regions, estimation of the correlation in the noise and identification of large-scale networks. These steps were designed under the following hypotheses regarding the statistical properties of the noise:

(1) The noise is predominant, i.e. the time series in most brain regions are only noise. We proved on fully synthetic data sets that it was possible to achieve a sensitive detection of a network with 30% of the cortex included in the network. While such performance cannot strictly be extended to real data sets, we can notice that, in most fMRI studies, effects of interest are reported in less than 30% of the brain, which is quite compatible with this result.

(2) The noise is Gaussian. Although we did not investigate the robustness of our method to departures from a Gaussian distribution, this hypothesis is thought to be roughly exact for fMRI data sets. Even if a Gaussian distribution may not be strictly valid for the time series, it may still hold asymptotically for the Fisher transform of the correlations.

(3) The noise is spatially stationary. More sophisticated models of the noise in fMRI have already been proposed, including separable stationary (Kruggel et al., 2002), non-separable stationary (Benali et al., 2001) and non-separable non-stationary (Woolrich et al., 2004) models. However, these works did not concentrate on large-scale correlations in the noise, that were expected to be reduced or even suppressed using a noise reduction strategy prior to the modeling step (Woolrich et al., 2004). We proposed that a first step towards modeling the residual correlations in the noise is an hypothesis of stationarity on the gray matter. We also took the temporal autocorrelation of the noise into account by estimating the variance of the Fisher transform of the correlations (Identification of large-scale networks), and by further neglecting the influence of temporal correlation. This correction is analogous to the estimation of effective degrees of freedom applied in the context of the general linear model (Worsley and Friston, 1995). We did not investigate

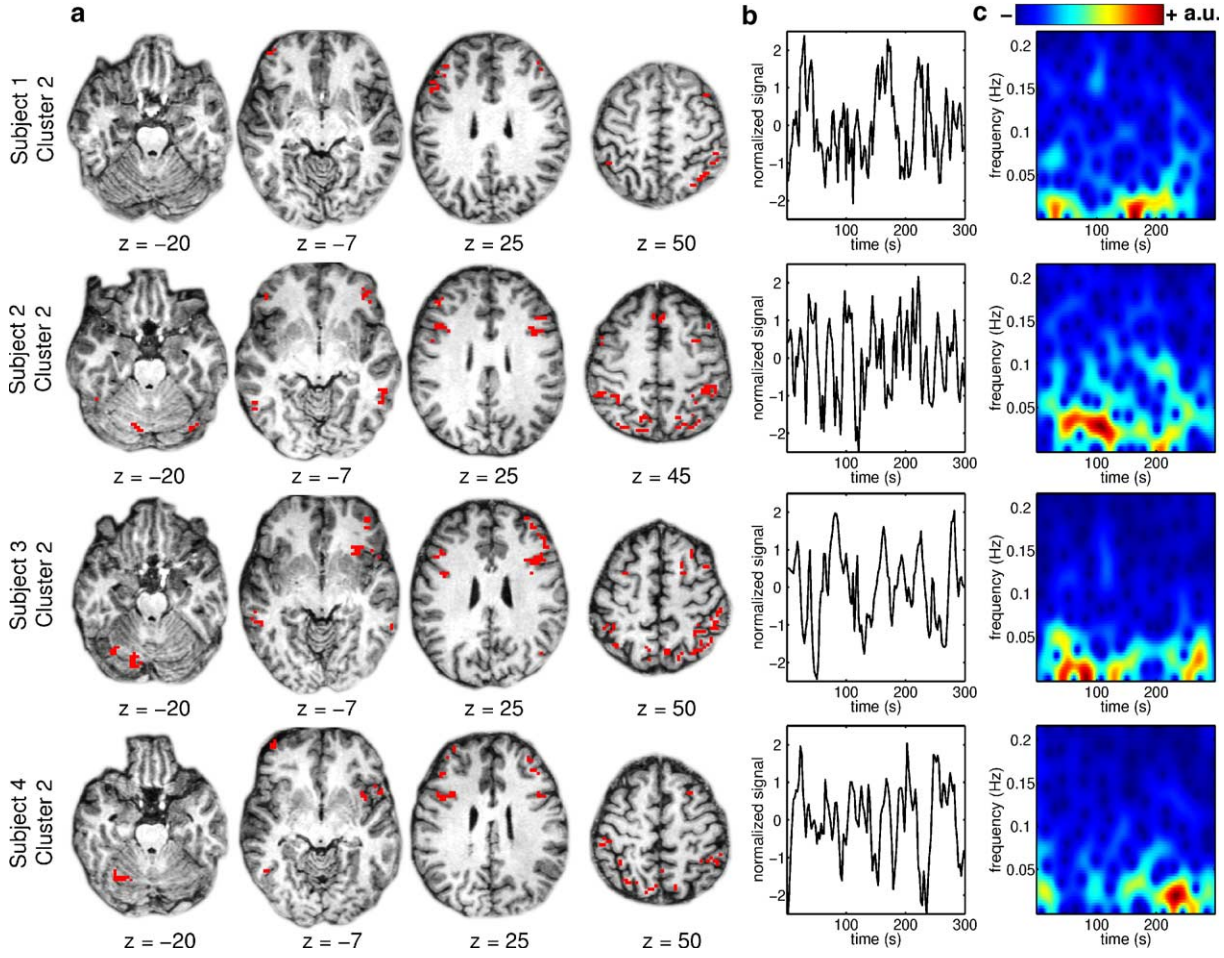

Fig. 7. Cluster of parieto-frontal regions: regions of this cluster have a low-frequency temporal activity, and are mainly located in the posterior parietal and in the lateral frontal cortices. (a) Spatial location for each subject, superimposed on the corresponding axial slice of the anatomical image (left on the image is left of the brain). (b) First principal component of the time series associated with the cluster. (c) Gaussian window Fourier transform of the first principal component (intensity in arbitrary units).

theoretically to which extent this estimation could be biased by spatial correlation. However, it proved sufficient on simulations of temporally and spatially correlated Gaussian noise.

(4) The spatial correlation of the noise is a rational-quadratic function of the spatial (Euclidean) distance. The reproducibility of the shape of the empirical correlograms over subjects, as was observed on real data sets, supports the use of this parametric model. Besides, the choice of the distance, e.g. Euclidean or geodesic, is not crucial. The large-scale effects were indeed almost constant beyond some critical distance, so the only expected difference between the Euclidian and the geodesic distance is the local shape of the correlogram, but not its asymptotic behavior. The Euclidean distance was the easier to handle.

An important parameter of the LSNI method is the critical size of the brain regions we used in the region growing algorithm. We used a critical size of 10 voxels with an in-plane resolution of 3×3 mm^2 because it roughly corresponds to a 1 cm^2 patch of cortex, compatible with the expected size of the local networks (Varela et al., 2001). This critical size moreover leads to less than 2000 regions on the cortex. This appeared as a trade-off between spatial resolution and numerical tractability when handling the correlation matrix. We did not systematically investigate the influence of the critical size parameter on the whole procedure.

We assessed the false-positive rate of the LSNI method under the null-hypothesis (pure noise) using fully synthetic data sets, with both simulated noise and simulated functional networks. In many studies, the noise comes from real resting-state data sets, e.g. Lu et al. (2003), and only the functional activity is simulated. In resting-state time series, the noise is indeed truly realistic, although there remains a spontaneous neural activity in the brain. For methods based on the general linear model, this activity can fairly be considered as independent of the simulated activity, and does not interfere with the detection procedure. Such an assumption does not hold in our case because spontaneous interactions within brain regions at rest cannot be neglected when large-scale functional interactions are concerned. We therefore used fully synthetic data sets made as similar as possible to real data sets: same location of regions, same number of time samples and same parameters of the spatial correlation as estimated on real data sets.

Before applying the LSNI method on real data sets, some preprocessing was necessary. Among these, a gray matter

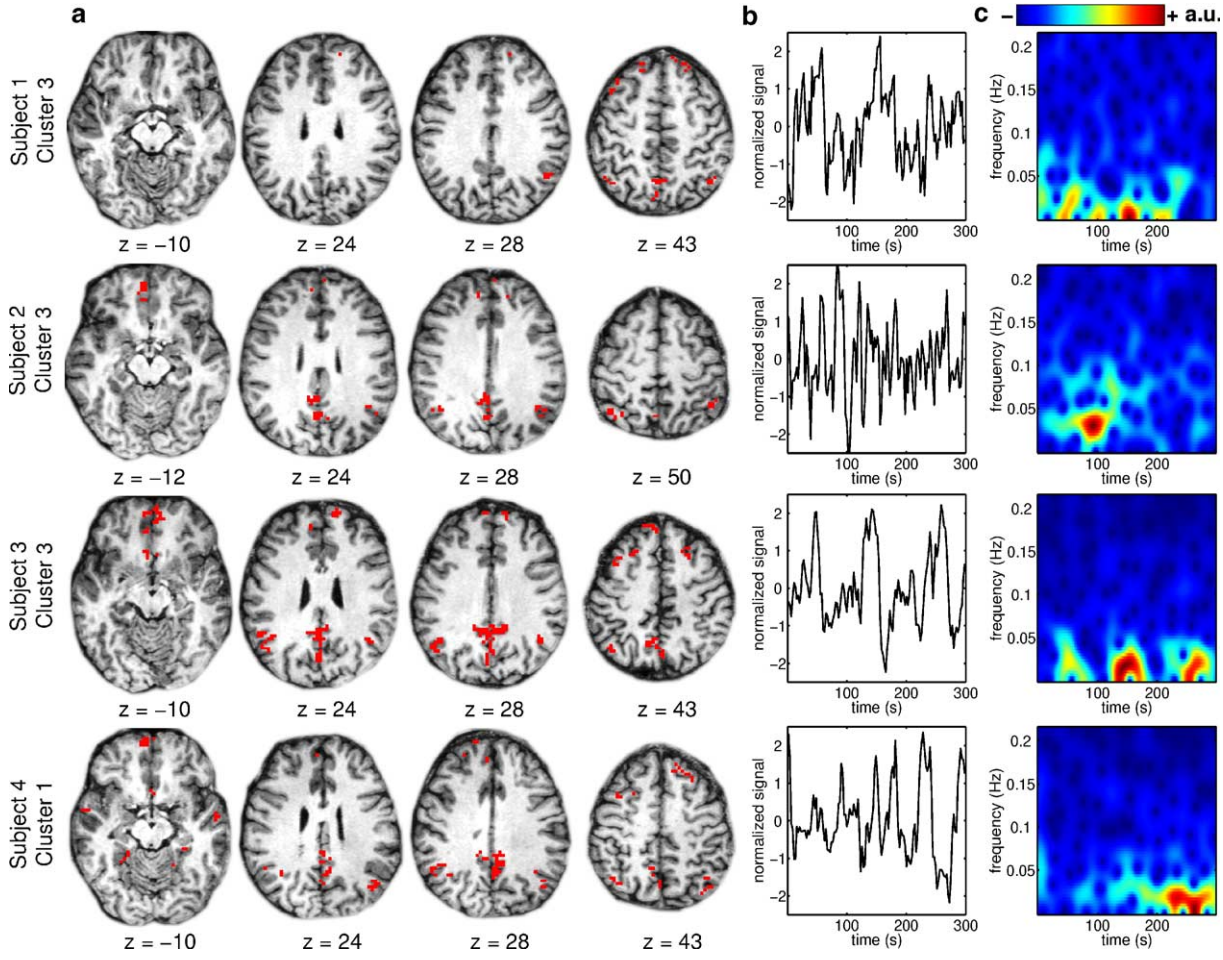

Fig. 8. Cluster of medio-posterior regions: regions of this cluster have a low-frequency temporal activity, and are mainly located in the posterior cingulate, the medial frontal cortex and in posterior regions. (a) Spatial location for each subject, superimposed on the corresponding axial slice of the anatomical image (left on the image is left of the brain). (b) First principal component of the time series associated with the cluster. (c) Gaussian window Fourier transform of the first principal component (intensity in arbitrary units).

segmentation was conducted, which makes the hypothesis of spatial stationarity of the noise more plausible than for a whole brain analysis. The noise reduction step was of considerable importance, because spatially localized noise is clearly a confound for our method, as it is for any functional connectivity-based method. We used one of the few available procedures for long TR acquisition, while it is still not fully automated and requires a manual segmentation of the ventricles and basal artery. This step will profit of future works in the field of noise reduction.

Once the large-scale networks were detected using the LSNI method, a post-processing was applied in order to facilitate the interpretation of the results. It consisted of grouping the identified regions into distinct subnetworks using hierarchical clustering. Our aim was to describe as simply as possible the networks that we found, in order to discuss whether or not they were relevant from a functional point of view. Hierarchical clustering appeared well-suited for this purpose, even if the number of clusters had to be specified arbitrarily. Many other ways to analyze the structure of large-scale networks proposed so far in the fMRI literature do not require to specify such arbitrary parameters, e.g. multi-dimensional scaling (Salvador et al., 2005) or functional integration (Tononi et al., 1998), and could highlight other aspects of the connectivity patterns.

Results on real data sets showed that the main features of the correlogram were reproducible across subjects ($h_\infty^\varepsilon \sim 3$ cm and $\rho_\infty \sim 0$). Given these parameters, it is interesting to notice that the identification of the large-scale network roughly simplifies to finding non-zero correlation between regions more than 3 cm apart. However, it is not fully elucidated yet whether this result is general or whether it depends on field strength, sequence type or preprocessing strategies.

Some small networks were not reproducibly found across subjects. As simulations showed that our method was less sensitive for small networks than for large ones, this effectively implies a large inter- and intrasubjects variability for small networks. Moreover, the time-frequency analysis showed that some clusters were physiology-induced false-positives. Such issue could be overcome by improving the noise reduction step, or by identifying a posteriori regions whose power spectrum contains high frequencies.

The large functional clusters shared many characteristics across subjects, despite an important variability concerning their sizes. First, all subjects had a cluster of regions with stimulus-locked activity. Although not surprising, finding this cluster is a

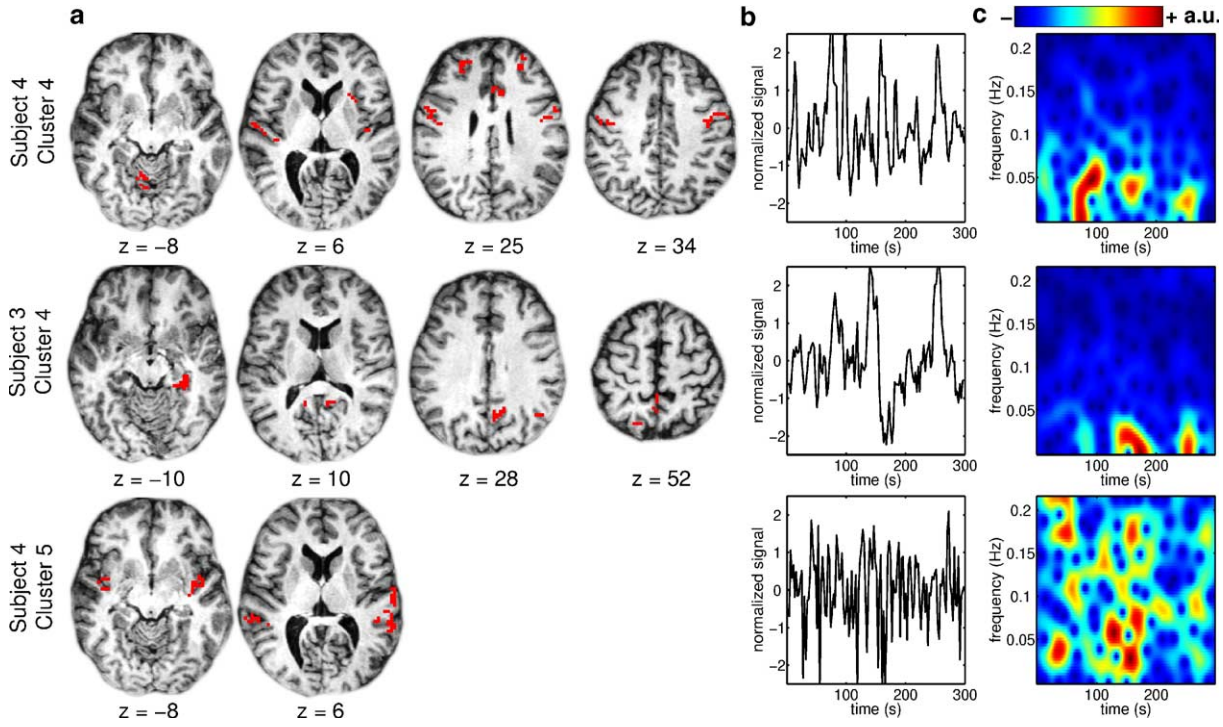

Fig. 9. Other clusters: typical examples of clusters that did not share clear commonalities across subjects. (a) Spatial location of the cluster, superimposed on the corresponding axial slice of the anatomical image (left on the image is left of the brain). (b) First principal component of the time series associated with the cluster. (c) Gaussian window Fourier transform of the first principal component (intensity in arbitrary units).

first-step validation of the LSNI method. For all subjects, there were also two large clusters whose temporal activities were dominated by low-frequency fluctuations (0–0.1 Hz). Such low-frequency fluctuations have been consistently reported for functional connectivity at rest (Biswal et al., 1995; Cordes et al., 2001). The spatial distribution of the medio-posterior cluster involving posterior cingulate was remarkably similar to the connectivity map reported by Greicius et al. (2003), obtained using a seed region in the posterior cingulate at rest, and which overlaps the so-called "default-mode network". The same work reported that the map was minimally disrupted during a simple visual task, hence this network could also be expected in our motor paradigm. The default-mode network is thought to contribute to the "deactivation" phenomenon, i.e. regions exhibiting a stronger activity during rest than during a task (Gusnard and Raichle, 2001), and moreover includes most regions of the parieto-frontal cluster.

Because it is an exploratory method that does not rely on an explicit model of the acquisition, the LSNI method allows to explore new types of data sets. For instance, as far as resting-state data sets are concerned, only isolated individual subnetworks have been studied in the literature, partly because selecting the regions was a hard task. A systematic study of the relations between subnetworks in the resting-state could be handled using the LSNI procedure. Also, an extensive study of the brain large-scale networks would have important clinical applications, for example it would help to better distinguish between Alzheimer's disease and healthy aging (Greicius et al., 2004). The LSNI method would also allow to select subnetworks for which it would be possible to apply more hypothesis-driven methods, such as structural equation modeling (McIntosh and Gonzalez-Lima, 1994) or dynamical causal modeling (Friston et al., 2003), without excluding important nodes from a functional connectivity point of view.

Conclusion

In this paper, we have proposed a method for large-scale network identification (LSNI) in fMRI on the basis of two main spatial features, namely the local temporal homogeneity of the regions and the existence of strong large-scale interactions. To achieve this identification, the spatial structure of the noise was robustly estimated. A simulation study demonstrated that the procedure correctly controlled the false-positive rate, and was indeed able to identify functional networks involving up to 30% of the cortex. On real data sets with a block-designed motor task, both temporal and anatomical features of the large-scale networks strongly support that the regions identified by the method indeed have large-scale functional connections. This work provides a first statistical formulation of the large-scale network identification problem, and enables the systematic identification of the salient networks of the working brain.

Acknowledgments

The authors are thankful to Roberto Toro and Jean Daunizeau for helpful comments.

Appendix A. Competitive region growing algorithm

The goal of the algorithm is to partition the set of all voxels on the brain cortex into a set F of regions C that are connected,

disjoint, of a given size $t_s \geq 1$, and that are optimally homogeneous, according to a measure that is the sum of all regional homogeneities defined hereafter:

$$\hom(C) = \frac{2}{\#C(\#C-1)} \sum_{v,w \in C, v \neq w} \mathrm{corr}(\mathbf{y}_v, \mathbf{y}_w),$$

where corr is Pearson's linear correlation, $\#C$ is the number of voxels in region C, and $\mathbf{y}_u = (y_v^t)_{t=1}^T$ is the time series of voxel v. The associated similarity measure between two regions is:

$$s(C,D) = \frac{1}{\#C\#D} \sum_{(v,w) \in C \times D} \mathrm{corr}(\mathbf{y}_v, \mathbf{y}_w).$$

Finding the best set F is computationally hard, and could even be impossible if the number of voxels was not a multiple of the critical size t_s. The proposed method to find a "good" set F is an iterative procedure derived from Clustering Algorithm for Medical Image Sequences, CAMIS (Benali et al., 1994). At step n, let E_n denote all regions that are candidates for merging, and denote by F_n the set of regions that have already been validated.

At the initial step ($n = 0$), regions are singletons of voxels, all are candidates for merging and none is validated yet:

$$E_0 = \{\{v\}, v \in M\},\ F_0 = \varnothing.$$

At step n, the first condition for a pair of candidate regions to merge is that those regions are neighbors. Let $\mathcal{N}(v)$ denote the standard 26-connexity neighborhood for voxel v, and let us define the neighborhood of a region C in E_n as follows:

$$\mathcal{N}_n(C) = \{D \in E_n, D \neq C | \exists (v,w) \in C \times D, w \in \mathcal{N}(v)\}.$$

A pair of neighboring regions (C, D) in $E_n \times E_n$ will merge if it fulfills the following *mutual nearest neighbors* condition (Gowda and Krishna, 1977):

$$mnn(C,D) \Leftrightarrow \begin{cases} C &= \arg\min_{K \in \mathcal{N}_n(D)} s(K,D) \\ \text{and} & \\ D &= \arg\min_{L \in \mathcal{N}_n(C)} s(C,L) \end{cases}.$$

A region of E_n is validated if its number of voxels is greater or equal to the critical size t_s:

$$F_{n+1} = F_n \cup \{C \cup D | C, D \in E_n, mnn(C,D), \#(C \cup D) \geq t_s\}.$$

At next step, regions candidate for merging are those that have not merged at this step, or those that are still too small to be validated. We thus define E_{n+1} as follows:

$$E_{n+1} = \{C \cup D | C, D \in E_n, mnn(C,D), \#(C \cup D) < t_s\} \cup \{C \in E_n | \exists \neq D \in E_n, mnn(C,D)\}.$$

The algorithm stops if $E_n = \varnothing$ or if there remains no pair of regions in E_n that verifies the mutual nearest neighbors criterion.

As the number of regions in E_n is an integer and strictly decreases at each step, as long as pairs of mutual nearest neighbors still exist, the algorithm necessarily stops at some point. The final set of homogeneous regions F is the set F_n once convergence of the algorithm is achieved. Note that the size of a validated region ranges from t_s to $(2t_s - 2)$, because it is possible that two regions of size $(t_s - 1)$ merge during the growing process. Moreover, if the set E_n is not empty once convergence is achieved, the final set F_n is not a partition of all voxels because some voxels remain that do not belong to any validated region. These voxels are discarded from further analysis.

Appendix B. Parametric model for spatial correlogram of fMRI data sets

The rational-quadratic model of the spatial correlogram as defined in Eq. (7) can be parameterized by the following quantities, which have a simple graphical interpretation as illustrated in Fig. 1:

$$\begin{cases} \rho_{0+} = \lim_{h \to 0+} (\rho_\theta(h)), & 0 < \rho_{0+} \leq 1, \\ \rho_\infty = \lim_{h \to +\infty} (\rho_\theta(h)), & 0 < \rho_\infty < \rho_{0+}, \\ h_\infty^\varepsilon = \min\{h > 0 : \rho_\theta(h) - \rho_\infty < \varepsilon\}, & 0 < \varepsilon < \rho_{0+} - \rho_\infty \end{cases} \quad \text{(B.1)}$$

The following relations link the parameters $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3)$ of Eq. (7), and the parameters of interest ρ_{0+}, ρ_∞ and h_∞^ε:

$$\begin{cases} \theta_1 = 1 - \rho_{0+}, \\ \theta_3 = \frac{\varepsilon}{\rho_{0+} - \rho_\infty - \varepsilon} \left(h_\infty^\varepsilon\right)^2, \\ \theta_2 = \left(\frac{\rho_{0+} - \rho_\infty}{h_\infty^\varepsilon}\right)\left(\frac{\rho_{0+} - \rho_\infty - \varepsilon}{\varepsilon}\right), \end{cases} \quad \text{(B.2)}$$

and

$$\begin{cases} \rho_{0+} = 1 - \theta_1, \\ \rho_\infty = 1 - \theta_1 - \theta_2\theta_2, \\ \left(h_\infty^\varepsilon\right)^2 = \left(\frac{\theta_2\theta_3}{\varepsilon} - 1\right)\theta_3. \end{cases} \quad \text{(B.3)}$$

Appendix C. Optimization procedure for the parametric spatial model

Let $\rho_{\boldsymbol{\theta}}$ be a valid correlogram model, and let us consider the Fisher transform of both the expected and empirical values of the correlogram at lags $(h_1, \ldots, h_K)$:

$$\boldsymbol{\zeta}_{\boldsymbol{\theta}} = (F(\rho_{\boldsymbol{\theta}}(h_1)), .., F(\rho_{\boldsymbol{\theta}}(h_K))), \hat{\boldsymbol{\zeta}} = (F(\hat{\boldsymbol{\rho}}(h_1)), .., F(\hat{\boldsymbol{\rho}}(h_K))).$$

In the asymptotic assumption of a Gaussian distribution for $F(\rho(h_k))$, with no space–time correlation, maximizing the probability of observing $\hat{\boldsymbol{\zeta}}$ for a given $\boldsymbol{\zeta}_{\boldsymbol{\theta}}$ is equivalent to minimizing the following cost function:

$$\sum_{k=1}^{K} \#D_{h_k} \{F(\hat{\boldsymbol{\rho}}(h_k)) - F(\rho_{\boldsymbol{\theta}}(h_k))\}^2. \quad \text{(C.1)}$$

where $\#D_{h_k}$ is the total number of pairs of regions at spatial lag h_k. The general least-squares fitting consists in finding the parameters $\boldsymbol{\theta}^*$ that minimize the cost function (C.1). Although the effect of space and time correlations is neglected, such a cost function still weights differently each spatial lag depending on the number of observations $\#D_{h_k}$ (Cressie, 1993).

In the optimization procedure of the parameters $\boldsymbol{\theta}$, only the spatial lags less than half the largest lag were considered, following the recommendations of Journel and Huijbregts (1978). Optimization of the cost function (C.1) was achieved using the Nelder–Mead method implemented in Matlab®, starting from a random initialization. This step was repeated until the parameters could not be improved for 50 iterations, in order to avoid local minima.

References

Anderson, T.W., 1958. An introduction to multivariate statistical analysis. Wiley Publications in StatisticsJohn Wiley and Sons, New York.

Baumgartner, R., Windischberger, C., Moser, E., 1998. Quantification in functional magnetic resonance imaging: fuzzy clustering vs. correlation analysis. Magn. Reson. Imaging 16, 115–125.

Bell, A.J., Sejnowski, T.J., 1995. An information maximization approach to blind separation and blind deconvolution. Neural Comput. 7, 1129–1159.

Benali, H., Buvat, I., Frouin, F., Bazin, J.P., Di Paola, R., 1994. CAMIS: clustering algorithm for medical image sequences using a mutual nearest neighbour criterion. In: Bookstein, F.L., Duncan, J.S., Lange, N., Wilson, D.C. (Eds.), Mathematical Methods in Medical Imaging III. Proc. SPIE 2299, pp. 336–347.

Benali, H., Pélégrini-Issac, M., Kruggel, F., 2001. Spatio-temporal covariance model for medical image sequences: application to functional MRI data. In: Insana, M.F., Leahy, R.M. (Eds.), Information Processing in Medical Imaging, Lecture Notes in Computer Science, vol. 2082. Springer, Berlin, pp. 197–203.

Biswal, B.B., Yetkin, F.Z., Haughton, V.M., Hyde, J.S., 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. Magn. Reson. Med. 34, 537–541.

Cordes, D., Haughton, V.M., Arfanakis, K., Wendt, G.J., Turski, P.A., Moritz, C.H., Quigley, M.A., Meyerand, M.E., 2000. Mapping functionally related regions of brain with functional connectivity MR imaging. Am. J. Neuroradiol. 21, 1636–1644.

Cordes, D., Haughton, V.M., Arfanakis, K., Carew, J.D., Turski, P.A., Moritz, C.H., Quigley, M.A., Meyerand, E., 2001. Frequencies contributing to functional connectivity in the cerebral cortex in "resting-state" data. Am. J. Neuroradiol. 22, 1326–1333.

Cordes, D., Haughton, V., Carew, J.D., Arfanakis, K., Maravilla, K., 2002. Hierarchical clustering to measure connectivity in fMRI resting-state data. Magn. Reson. Imaging 20, 305–317.

Cressie, N., 1993. Statistics for Spatial Data. John Wiley and Sons, New York.

Friston, K.J., 1994. Functional and effective connectivity in neuroimaging: a synthesis. Hum. Brain Mapp. 2, 56–78.

Friston, K.J., Frith, C.D., Liddle, P.F., Frackowiak, R.S.J., 1993. Functional connectivity: the principal component analysis of large (PET) data sets. J. Cereb. Blood Flow Metab. 1, 153–171.

Friston, K., Buechel, C., Fink, G.R., Morris, J., Rolls., E., Dolan, R.J., 1997. Psychophysiological and modulatory interactions in neuroimaging. NeuroImage 6, 218–229.

Friston, K.J., Harrison, L., Penny, W., 2003. Dynamic causal modeling. NeuroImage 19, 1273–1302.

Gowda, K.C., Krishna, G., 1977. Agglomerative clustering using the concept of mutual nearest neighbourhood. Pattern Recogn. 10, 105–112.

Greicius, M.D., Krasnow, B., Reiss, A.L., Menon, V., 2003. Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. Proc. Natl. Acad. Sci. U. S. A. 100, 253–258.

Greicius, M.D., Srivastava, G., Reiss, A.L., Menon, V., 2004. Default-mode network activity distinguishes Alzheimer's desease from healthy aging: evidence from functional MRI. Proc. Natl. Acad. Sci. U. S. A. 101, 4637–4642.

Gusnard, D.A., Raichle, M.E., 2001. Searching for a baseline: functional imaging and the resting human brain. Nature 2, 685–694.

Hebb, D., 1949. The Organization of Behavior. Wiley, New York.

Horwitz, B., Tagamets, M.-A., McIntosh, A.R., 1999. Neural modeling, functional brain imaging, and cognition. Trends Cogn. Sci. 3, 91–98.

Journel, A., Huijbregts, C., 1978. Mining Geostatistics. Academic Press, London.

Kruggel, F., Pélégrini-Issac, M., Benali, H., 2002. Estimating the effective degrees of freedom in univariate multiple regression analysis. Med. Image Anal. 6, 63–75.

Lowe, M.J., Mock, B.J., Sorenson, J.A., 1998. Functional connectivity in single and multislice echoplanar imaging using resting-state fluctuations. NeuroImage 7, 119–132.

Lowe, M.J., Dzemidzic, M., Lurito, J.T., Mathews, V.P., Phillips, M.D., 2000. Correlations in low-frequency BOLD fluctuations reflect cortico-cortical connections. NeuroImage 12, 582–587.

Lu, Y., Jiang, T., Zang, Y., 2003. Region growing method for the analysis of functional MRI data. NeuroImage 20, 455–465.

McIntosh, A.R., Gonzalez-Lima, F., 1994. Structural equation modeling and its application to network analysis of functional brain imaging. Hum. Brain Mapp. 2, 2–22.

McKeown, M., Makeig, S., Brown, G., Jung, T., Kindermann, S., Bell, A., Sejnowski, T., 1998. Analysis of fMRI data by blind separation into independent spatial components. Hum. Brain Mapp. 6, 160–188.

Perlbarg, V., Bellec, P., Marrelec, G., Jbabdi, S., Pélégrini-Issac, M., Lehéricy, S., Doyon, J., Benali, H., 2004. Selection of spatially independent components to reduce physiological noise in fMRI data. Proceedings of the 10th International Conference on Functional Mapping of the Human Brain: CDROM, Neuroimage, vol. 22.

Rodriguez, R.N., 1982. Correlation. In: Kotz, S., Johnson, N.L. (Eds.), Encyclopedia of Statistical Sciences, vol. 1. Wiley, New York, pp. 193–204.

Rousseeuw, J., Croux, C., 1993. Alternatives to the median absolute deviation. J. Am. Stat. Assoc. 88, 1273–1283.

Salvador, R., Suckling, J., Coleman, M.R., Pickard, J.D., Menon, D., Bullmore, E., 2005. Neurophysiological architecture of functional magnetic resonance images of human brain. Cereb. Cortex 34, 387–413.

Stein, T., Moritz, C., Quigley, M., Haugton, D., Meyerand, E., 2000. Functional connectivity in the thalamus and hippocampus studied with functional MR imaging. Am. J. Neuroradiol. 21, 1397–1401.

Thomas, C.G., Harshman, R.A., Menon, R.S., 2002. Noise reduction in BOLD-based fMRI using component analysis. NeuroImage 17, 1521–1537.

Tononi, G., McIntosh, A.R., Russel, D.P., Edelman, G.M., 1998. Functional clustering: identifying strongly interactive brain regions in neuroimaging data. NeuroImage 7, 133–149.

Varela, F., Lachaux, J.-P., Rodriguez, E., Martinerie, J., 2001. The brainweb: phase synchronization and large-scale integration. Nat. Rev., Neurosci. 2, 229–239.

Woolrich, M., Jenkinson, M., Brady, M., Smith, S.M., 2004. Fully Bayesian spatio-temporal modeling of fMRI data. IEEE Trans. Med. Imaging 23, 213–230.

Worsley, K.J., Friston, K.J., 1995. Analysis of fMRI time-series revisited-again. NeuroImage 2, 173–181.

Xiong, J., Parsons, L.M., Gao, J.-H., Fox, P.T., 1999. Interregional connectivity to primary motor cortex revealed using MRI resting state images. Hum. Brain Mapp. 8, 151–156.

Zang, Y., Jiang, T., Lu, Y., He, Y., Tian, L., 2004. Regional homogeneity approach to fMRI data analysis. NeuroImage 22, 394–400.