ELSEVIER

NeuroImage

www.elsevier.com/locate/ynimg
NeuroImage 32 (2006) 228 – 237

Partial correlation for functional brain interactivity investigation in functional MRI

Guillaume Marrelec,[a,b,c,*] Alexandre Krainik,[d,e] Hugues Duffau,[a,b,f] Mélanie Pélégrini-Issac,[a,b] Stéphane Lehéricy,[g] Julien Doyon,[a,b,c] and Habib Benali[a,b,c]

[a]*Inserm, U678, Paris F-75013, France*
[b]*Université Pierre et Marie Curie, Faculté de médecine Pitié-Salpêtrière, Paris F-75013, France*
[c]*Université de Montréal, MIC/UNF, Montreal, Canada H3W 1W5*
[d]*Inserm, U594, Grenoble F-38000, France*
[e]*CHU La Tronche, Department of Neuroradiology, Grenoble F-38000, France*
[f]*AP-HP, Hôpital Pitié-Salpêtrière, Department of Neurosurgery, Paris F-75013, France*
[g]*AP-HP, Hôpital Pitié-Salpêtrière, Department of Neuroradiology, Paris F-75013, France*

Received 8 August 2005; revised 16 December 2005; accepted 29 December 2005
Available online 13 June 2006

Examination of functional interactions through effective connectivity requires the determination of three distinct levels of information: (1) the regions involved in the process and forming the spatial support of the network, (2) the presence or absence of interactions between each pair of regions, and (3) the directionality of the existing interactions. While many methods exist to select regions (Step 1), very little is available to complete Step 2. The two main methods developed so far, structural equation modeling (SEM) and dynamical causal modeling (DCM), usually require precise prior information to be used, while such information is sometimes lacking. Assuming that Step 1 was successfully completed, we here propose a data-driven method to deal with Step 2 and extract functional interactions from fMRI datasets through partial correlations. Partial correlation is more closely related to effective connectivity than marginal correlation and provides a convenient graphical representation for functional interactions. As an instance of brain interactivity investigation, we consider how simple hand movements are processed by the bihemispheric cortical motor network. In the proposed framework, Bayesian analysis makes it possible to estimate and test the partial statistical dependencies between regions without any prior model on the underlying functional interactions. We demonstrate the interest of this approach on real data.

Keywords: Functional MRI; Functional brain interactivity; Motor network; Partial correlation; Functional connectivity; Effective connectivity; Bayesian analysis

Introduction

Blood oxygen level-dependent (BOLD) functional magnetic resonance imaging (fMRI), which allows to dynamically follow metabolic and hemodynamic consequences of brain activity (Chen and Ogawa, 1999; Huettel et al., 2004), has deeply modified our knowledge about the brain (Frackowiak et al., 2004). Since its appearance, it has primarily been concerned with locating brain processes in determinate regions, thus proposing a topography of brain activity (Friston et al., 1994; Worsley and Friston, 1995). However, this approach conveys a rather static idea of brain processes. By contrast, it is now believed that processing of a functional task by the brain can only be performed through interaction of segregated regions within a complex network (Hebb, 1949; Tononi et al., 1998; Frackowiak et al., 2004; Sporns et al., 2004). Consequently, various methods have been proposed in fMRI data analysis to extract information of interaction from datasets, most of which rely on either functional or effective connectivity (for reviews and discussions, see, e.g., Stone and Kötter, 2002; Horwitz, 2003; Lee et al., 2003).

Effective connectivity aims at examining the influence that regions exert on each other (Friston et al., 1993a). Its investigation requires three distinct levels of information: (1) the regions involved in the process and forming the spatial support of the network, (2) the presence or absence of interactions between each pair of regions, and (3) the directionality of the existing interactions. Step 1, though far from being trivial, has already been examined in detail by the existing literature. As a result, many methods are available to this end, including activation maps (Huettel et al., 2004), as well as methods deriving from functional connectivity, such as principal or independent component analysis (Friston et al., 1993b; Arfanakis et al., 2000), or correlation maps (Biswal et al., 1995, 1997; Xiong et al., 1999). In our study, we

* Corresponding author. MIC/UNF, CRIUGM, 4565 chemin Queen-Mary, Montréal, QC, Canada H3W 1W5.
E-mail address: guillaume.marrelec@umontreal.ca (G. Marrelec).
Available online on ScienceDirect (www.sciencedirect.com).

1053-8119/$ - see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2005.12.057

assume that this step has been successfully completed, providing us with a set of regions and corresponding signals.

Step 2, on which we will henceforth concentrate, then focuses on functional interactions per se. Structural equation modeling (SEM) is the most widespread way to model effective connectivity (Gonzalez-Lima and McIntosh, 1995; Bullmore et al., 2000), even though dynamical causal modeling (DCM) has recently been developed to the same goal (Friston et al., 2003; Penny et al., 2004b). Both methods rely on the definition of a model in the form of a directed graph prior to the analysis. The graph links are usually thought of as anatomical connections that are functionally relevant for the experiment under consideration. In this setting, there is hence an anatomical connection underlying each functional interaction. Unfortunately, several drawbacks still hamper the use of SEM and DCM, most of which originate from the difficulty to propose a prior model of effective connectivity.

First, it is virtually impossible for current methods to perform extensive model comparison. Indeed, the complexity of interactivity analysis increases exponentially with the number of regions. For instance, working with D regions in SEM brings up to $4^{D(D-1)/2}$ potential graphs. For a mere $D = 6$ regions, this amounts to $4^{15} \approx 10^9$ graphs, which is already humongous and computationally out of reach. The complexity of DCM is at least as high. Methods have been developed to cope with this issue (Bullmore et al., 2000; Mechelli et al., 2002; Penny et al., 2004a). Yet, they are only able to perform model comparison within a limited subset of potential models given a priori.

Accordingly, it lies on the expert's hands to sift through the set of possible graphs and provide only a few potentially relevant models on which the method will focus. This process requires information that usually comes from literature on brain anatomy. A large amount of information is available from previous studies regarding the underlying neuroanatomy of some networks. Such evidence may come from lesion studies, electrophysiological investigations in monkeys or imaging analyses. However, the results of those experiments are usually incommensurate with one another for various reasons. The relevance of each finding to a different or more general group of subjects cannot always easily be assessed. Information originating from monkey studies must be used with care to constrain human structural connections. In what measure patient case studies can be used for healthy subjects remains an open issue; conversely, results on healthy subjects cannot usually be applied to just any patient. Results from different modalities suffer from different kinds of artifacts that make them hard to compare. Last, and most importantly, only a very limited amount of information is available for many other networks (e.g., language, executive control).

Further hypotheses have hence to be made, such as concentrating on anatomical connections that are assumed to be functionally relevant to the task under investigation. As a matter of fact, even though the brain is strongly interconnected, the principle of segregation states that only a few regions should participate to a given functional task (Tononi et al., 1998). Consequently, it could be hypothesized that only few of the anatomical connections are actually used during the processing of a task, particularly if the given task is simple. Unfortunately, relating anatomical connections to functional interactions remains an open issue. A functional interaction requires, yet is not entailed by, the presence of an anatomical connection. The cortical motor network is a remarkable illustration of this fact. From a purely anatomical perspective, this network is very likely to be fully connected, preventing the use of SEM due to a lack of degrees of freedom. Nonetheless, it still seems meaningful to study how regions of this network interact when performing a simple hand movement.

These issues clearly show that existing methods of effective connectivity require much prior information relative to the structure of functional interactions. Since the goal of the study is precisely to investigate this structure, that information is lacking in most cases. One must then resort to information relative to anatomical connectivity, which is often diffuse and cannot be easily translated in terms of functional interactions.

Use of directed graphs as structural models finally raises conceptual questions inherent to directed graphs, such as observational equivalence or identifiability (Pearl, 2001), which have not fully been taken into account in fMRI data analysis yet. These problems, which are critical for directed acyclic graphs, are even more crucial for cyclic graphs, which are commonly used in effective connectivity modeling, while many theoretical questions remain unsolved (Spirtes, 1995; Pearl and Dechter, 1996; Neal, 2000). These points remain to be tackled before interpretation of SEM and DCM analyses can be fully trusted.

For all these reasons, many situations occur where existing model-based approaches of effective connectivity cannot be applied. The issue to be faced is then the following: assuming that we are provided with a network of brain regions and have no other information of any kind relative to their functional interactivity, how can the global interaction structure of the network be inferred from fMRI data? In other words, is it possible to propose data-driven measures of direct functional interactions?

Although computationally convenient, data-driven functional connectivity (e.g., temporal correlation) cannot be used to this end, for it is well known that it can emerge from various configurations that are not related to direct interactions. Unlike effective connectivity, functional connectivity hence does not ensure that each functional interaction will be supported by an anatomical connection. Indeed, we showed that one of the reasons why functional connectivity was unable to extract information that could be interpreted in terms of effective connectivity was that it could not quantify interaction mediation (Marrelec et al., 2005a,b). We introduced conditional correlation as a way to circumvent this flaw and hypothesized that a mediated interaction should translate into a zero conditional correlation.

On this premise, we here propose a novel procedure that extracts partial correlation from the data. Partial correlation has the interesting features of providing a convenient summary of conditional independences (Whittaker, 1990) and, hence, of being more closely related to effective connectivity than marginal correlation. Furthermore, its use provides a convenient framework for graphical representation. Methods using partial correlation have already been proposed (e.g., McIntosh et al., 1996; Sun et al., 2004). In these previous approaches, though, partial correlation was used to eliminate the effect of the experimental design. In our approach, its use is rather to subtract and remove mutual dependencies on common influences from other brain areas. By conditioning the dependencies between two areas on other areas, the ensuing functional connectivity (i.e., partial correlation) reflects interactions between the two areas in question. Therefore, the use of partial correlation, allowing access to a quantity that is more closely related to direct interaction, takes the analysis of functional connectivity closer to the characterization of functional interactions in terms of effective connectivity. It is data-driven in the sense that, unlike

existing methods such as SEM and DCM, it does not require any prior information regarding functional interactions to proceed. The partial correlation approach is also unique in that it provides a first insight into the effective connectivity of the network. As such, it is not meant to replace SEM or DCM but may act as a preliminary step and possibly allow for a more efficient use of those methods.

In this paper, we examine a simple instance of brain interactivity investigation. More precisely, we concentrate on the bihemispheric cortical motor network involved in simple hand movements. Much is known about the regions involved in motor processing and their anatomical connections, mostly through investigation in primates and studies in patients (Kunzle, 1978; Leichnetz, 1986; Zigmond et al., 1999; Gazzaniga, 2000). Within the hierarchical organization of the motor system, the cortex stands as the highest level of motor control. Cortical regions known to be involved in the motor network include at least the two sensorimotor cortices, the two premotor areas, and the two supplementary motor areas. The main cortical inputs to the motor areas stem from the premotor and the supplementary motor cortices. Cortico-cortical inputs are also present, originating from the opposite hemisphere. Movement can elicit complex brain processes that are still under investigation (Kandel et al., 2000; Swinnen, 2002). How the available anatomical information can be translated in terms of models of effective connectivity therefore often remains unknown. Few structural models have been proposed so far (Calautti and Baron, 2003; Rogers et al., 2004), none of which is based on the six aforementioned regions and considers the interactions between hemispheres.

The objective of this article is therefore to investigate functional brain interactivity within the aforementioned network during simple hand movements. Since the regions defining the functional network are already known, correlation analysis (and, hence, functional connectivity) can only confirm that the regions selected are indeed relevant for the analysis. As to effective connectivity, it cannot be applied without the definition of a directed graph. Since the anatomical information available regarding this network does not put any constraint on the anatomical connectivity, no such graph can be proposed without strong hypothesis. We hence examine what kind of information can be extracted from the data using a partial correlation analysis.

The outline of this paper is the following. In the next section, we present the partial correlation model. A Bayesian scheme is then proposed to infer the interaction structure from the data. The following section is devoted to demonstrating the relevance of this approach through the analysis of real data from simple hand movements. Further issues are addressed in the discussion.

Partial correlations

Our objective is to investigate the functional interactions occurring between the $D = 6$ following cortical regions: the two supplementary motor areas, RSMA and LSMA (R standing for "Right", L for "Left"), the two sensorimotor cortices, RSMC and LSMC, and the two premotor cortices, RPMC and LPMC. This set of regions is denoted by $\mathcal{R}$, i.e.,

$$\mathcal{R} = \{\mathrm{RSMA}, \mathrm{LSMA}, \mathrm{RSMC}, \mathrm{LSMC}, \mathrm{RPMC}, \mathrm{LPMC}\}.$$

Let $\boldsymbol{y} = (\boldsymbol{y}_t)_{t = 1,\ldots,T}$ be the BOLD fMRI time courses of these six regions. Each $\boldsymbol{y}_t$ is further assumed to be a realization of a D-dimensional Gaussian variable $\boldsymbol{x} = (x_i)_{i = 1,\ldots,D}$ of mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. A measure of direct interaction strength between regions is the partial correlation matrix (Whittaker, 1990; Lauritzen, 1996). A partial correlation coefficient between two regions i and j, denoted by Π_{ij}, is a particular case of a conditional coefficient: it is the correlation between regions i and j conditioned on the set of remaining regions, i.e.,

$$\Pi_{ij} = \mathrm{Corr}\left[x_i, x_j | \boldsymbol{x}_{\mathcal{R}\setminus\{i,j\}}\right]. \quad (1)$$

There are hence $D(D-1)/2 = 15$ partial correlation coefficients that form the D-by-D partial correlation matrix $\boldsymbol{\Pi} = (\Pi_{ij})$. Each partial correlation coefficient is a measure of the interaction between the time courses of two regions once these signals have been projected on the subspace orthogonal to the time courses of all other regions. It hence only considers the "direct correlation" between both regions, i.e., the correlation that cannot be accounted for by the influence of any other area in the network.

Set $\boldsymbol{\Upsilon} = (\Upsilon_{ij}) = \boldsymbol{\Sigma}^{-1}$ the inverse covariance matrix of $\boldsymbol{x}$, also called concentration, or precision, matrix. Given $\boldsymbol{\Upsilon}$, the partial correlation coefficients can readily be calculated through the following relationship (Whittaker, 1990):

$$\Pi_{ij} = -\frac{\Upsilon_{ij}}{\sqrt{\Upsilon_{ii}\Upsilon_{jj}}} \quad (2)$$

for two distinct regions i and j, and $\Pi_{ii} = 1$.

Bayesian inference

This model being set, it is now necessary to infer the true partial correlation matrix underlying the group data. If we knew exactly the model parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$, the exact partial correlation coefficients would be unambiguously determined by Eq. (2). However, since the true values of $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are unknown and only partly accessible through the data, so is the value of $\boldsymbol{\Pi}$, which must hence be inferred.

Model

Since we want to draw inference at a group level, we have to take both intra- and intersubject variability into account. This can readily be performed in a Bayesian framework by introduction of a hierarchical model, where $\boldsymbol{\Sigma}_0$ stands for the group covariance structure and $\boldsymbol{\Sigma}_s$ for that of each subject s, $s = 1,\ldots, \boldsymbol{S}$.

Using standard Bayesian theory, it can first be shown that the likelihood function of a covariance matrix $\boldsymbol{\Sigma}_s$ for a given dataset $\boldsymbol{y}_s$ is proportional to an inverse Wishart distribution with T-D-2 degrees of freedom and scale matrix $\boldsymbol{S}_s = \sum_{t=1}^{T}\left(\boldsymbol{y}_{s,t} - \overline{\boldsymbol{y}_{s,t}}\right)\left(\boldsymbol{y}_{s,t} - \overline{\boldsymbol{y}_{s,t}}\right)^{\mathrm{t}}$, proportional to the sample covariance matrix of subject s (see Appendix A for a proof of this assertion):

$$\boldsymbol{y}_s | \boldsymbol{\Sigma}_s \sim \mathrm{Inv-Wishart}\left(T - D - 2, \boldsymbol{S}^{-1}\right).$$

For the sake of convenience, we assume that the covariance matrix of each subject s originates from an inverse Wishart distribution with scale matrix $\boldsymbol{\Sigma}_0$ and ν_0 degrees of freedom:

$$\boldsymbol{\Sigma}_s | \boldsymbol{\Sigma}_0, \nu_0 \sim \mathrm{Inv-Wishart}\left(\nu_0, \boldsymbol{\Sigma}_0^{-1}\right).$$

While the use of a conjugate prior for the $\boldsymbol{\Sigma}_s$'s greatly simplifies calculations, the proposed model can still efficiently capture the

intersubject variability through the tuning of parameter v_0 (Gelman et al., 1998).

Numerical sampling

Our objective is to obtain $\Pr(\boldsymbol{\Pi}_0|\boldsymbol{y})$, the posterior distribution for the group partial correlation matrix $\boldsymbol{\Pi}_0 = (\Pi_{0,ij})$, given the data. In the aforementioned model, this quantity cannot be calculated in close form. Nonetheless, we can resort to the following sampling scheme.

We first run Gibbs sampling on the model to propose a numerical approximation of $\Pr(\boldsymbol{\Sigma}_0|\boldsymbol{y})$ (Ruanaidh and Fitzgerald, 1996; Gelman et al., 1998) and successively sample each variable given the set of remaining variables. To allow for burn-in effect, we discard the first half of the samples and only keep the second half for consideration, that we note $(\boldsymbol{\Sigma}_0^{[l]})$, $l = 1,\ldots, L$ (e.g., $L = 5000/2$).

Given a sample $(\boldsymbol{\Sigma}_0^{[l]})$ of group covariance matrices, it is then possible to use Eq. (2) to obtain an approximation of the sought distribution. For sample l, calculate $\boldsymbol{\Upsilon}_0^{[l]} = (\boldsymbol{\Sigma}_0^{[l]})^{-1}$ and then $\boldsymbol{\Pi}_0^{[l]}$ from $\boldsymbol{\Upsilon}_0^{[l]}$ using Eq. (2). The marginal distribution of a given partial correlation can be approximated by the frequency histogram obtained from the sample. Likewise, all statistics and estimators can be approximated by their sample counterparts. For instance

$$\mathrm{E}\left[\Pi_{0,ij}|\mathbf{y}\right] \approx M_{ij} = \frac{1}{L}\sum_{l=1}^{L}\Pi_{0,ij}^{[l]} \tag{3}$$

$$\mathrm{Var}\left[\Pi_{0,ij}|\boldsymbol{y}\right] \approx V_{ij} = \frac{1}{L}\sum_{l=1}^{L}\left(\Pi_{0,ij}^{[l]} - M_{ij}\right)^2, \tag{4}$$

where "E" and "Var" respectively stand for the posterior expectation and variance. Having the mean and variance of the posterior distribution makes it possible to use Laplace (i.e., normal) approximation to $\Pr(\Pi_{0,ij}|\boldsymbol{y})$:

$$\Pr\left(\Pi_{0,ij}|_{\mathbf{y}}\right) \approx \mathcal{N}\left(M_{ij}, V_{ij}; \Pi_{0,ij}\right),$$

where $\mathcal{N}(m, v; x)$ is the normal distribution with mean m and variance v calculated at point x. Significance tests can also be approximated the same way. For instance, testing against the hypothesis $\Pi_{0,ij} = 0$ can be associated to the following significance: $\Pr(\Pi_{0,ij} > 0)$ if $M_{ij} > 0$ or $\Pr(\Pi_{0,ij} < 0)$ if $M_{ij} < 0$. These quantities can be approximated by

$$\Pr\left(\Pi_{0,ij} > 0\right) \approx \frac{1}{L}\#\left\{\Pi_{0,ij}^{[l]} > 0\right\}$$

and

$$\Pr\left(\Pi_{0,ij} < 0\right) \approx \frac{1}{L}\#\left\{\Pi_{0,ij}^{[l]} < 0\right\},$$

respectively, where "#" stands for the cardinal of a set. Comparing a given correlation coefficient $\Pi_{0,ij}$ between two conditions, e.g., right-hand movement and left-hand movement, can be achieved by testing against the hypothesis $\Pi_{\text{left hand},0,ij} = \Pi_{\text{right hand},0,ij}$ in a similar fashion.

The same procedure can incidentally also be used to generate an approximation of the posterior distribution of the marginal correlation coefficients, $\Pr(R_{0,ij}|\boldsymbol{y})$.

Real data

Imaging

The MR protocol was carried out with a General Electric 1.5 T Sigma system. Functional MRI using BOLD contrast was performed. The protocol included (1) two runs comprising 42 T_2*-weighted functional volumes each, each volume covering the whole frontal lobes (TR/TE/flip angle: 3000 ms/60 ms/90°, 20 contiguous slices per volume, 5 mm slice thickness, in-plane pixel size: 3.75 mm × 3.75 mm); and (2) one axial inversion recovery three-dimensional T_1-weighted image for anatomical localization.

Seven healthy right-handed male volunteers were scanned after giving informed consent set by the local ethic committee. The experimental design protocol consisted of two different blocked-trial tasks: self-paced flexion/extension of the fingers of the right or left hand, depending on the session. Before the experiment started, all subjects practiced each movement to keep amplitude, acceleration, and strength constant. The paradigm was block-designed, alternating rest (R) and activation (A), and consisted of seven epochs of 18 s each for either activation or rest (total duration of each run: 2 min 06 s in this order: R-A-R-A-R-A-R). Rest consisted of lying eyes closed in the magnet. The task instructions were auditory-cued using a digital audio tape and presented using standard headphones customized for fMRI experiments and inserted in a noise-protecting helmet that provided isolation from scanner noise. Direct observation of the tasks was performed by an investigator during the fMRI acquisitions.

Preprocessing

Preprocessing was performed in MATLAB®[1] with SPM99 software[2]. The first six images of each run were discarded for signal stabilization. For each subject, images were corrected for rigid subject motion with the first volume of each run used as a reference and transformed stereotactically to common spatial coordinates using the standard template of the Montreal Neurological Institute (MNI). The resulting images were smoothed with a Gaussian isotropic spatial filter (FWHM = 5 × 5 × 5 mm).

Region and signal selection

The six regions composing the network (RSMA, LSMA, RSMC, LSMC, RPMC, and LPMC) were manually drawn by an expert onto normalized T_1-weighted anatomical images without reference to the activation patterns, using a standard sulcal atlas (Talairach and Tournoux, 1988; Naidich et al., 2001). Coregistration across anatomical and functional images and across subjects was assessed on anatomical landmarks located in the vicinity of the ROIs such as interhemispheric fissure, "hand knob", and the crossing between precentral and frontal superior sulci. We also used standardized ROIs to avoid an effect of ROIs volume across subjects.

The signal characteristic of each region was then selected as the spatial average of the time course of all voxels within the region. This signal was then translated and scaled to be of mean 0 and

[1] The Mathworks Inc., Natick, MA, USA.
[2] Wellcome Department of Cognitive Neurology, UCL, London, UK.

variance 1. We finally obtained seven times twelve time courses of 36 times samples each.

Results

We first estimated the correlation and the partial correlation matrices corresponding to the right- and left-hand movements, respectively. This was performed by use of the sampling scheme detailed earlier. We used the samples to obtain approximations of the mean and standard deviation for both Π_{ij} and R_{ij}. We plotted these results in Fig. 1. For an easier visualization, the average interaction structure has also been represented in graphical form in Fig. 2.

Using the same sampling scheme, it was also possible to determine which marginal and partial correlations were significantly different from zero for a given hand movement. The data confirmed that all marginal correlations were strongly positive and significantly different from 0 ($P < 0.025$). This confirms the fact that the six regions considered are indeed part of a same network and relevant for the analysis. On the other hand, the significance of partial correlations varied depending on the coefficient considered. Thresholding to a level $P = 0.025$ led to the same structure for left- and right-hand movements, depicted in Fig. 3. Almost all intrahemispheric partial correlations remain (five out of six), while none of the non-homologous interhemispheric partial correlations significantly differed from 0.

It was also possible to test for significant changes from a left to a right-hand movement. The results, summarized in Figs. 2 and 4, clearly indicate that, when changing from a right to a left-hand movement, the RPMC-RSMC partial correlation significantly increased ($P = 0.05$), while the LPMC-LSMC partial correlation decreased. Conversely, when changing from a left to a right-hand movement, the RPMC-RSMC partial correlation significantly decreased, while the LPMC-LSMC increased (same threshold). The same pattern was observed with marginal correlation only with a slightly less conservative threshold.

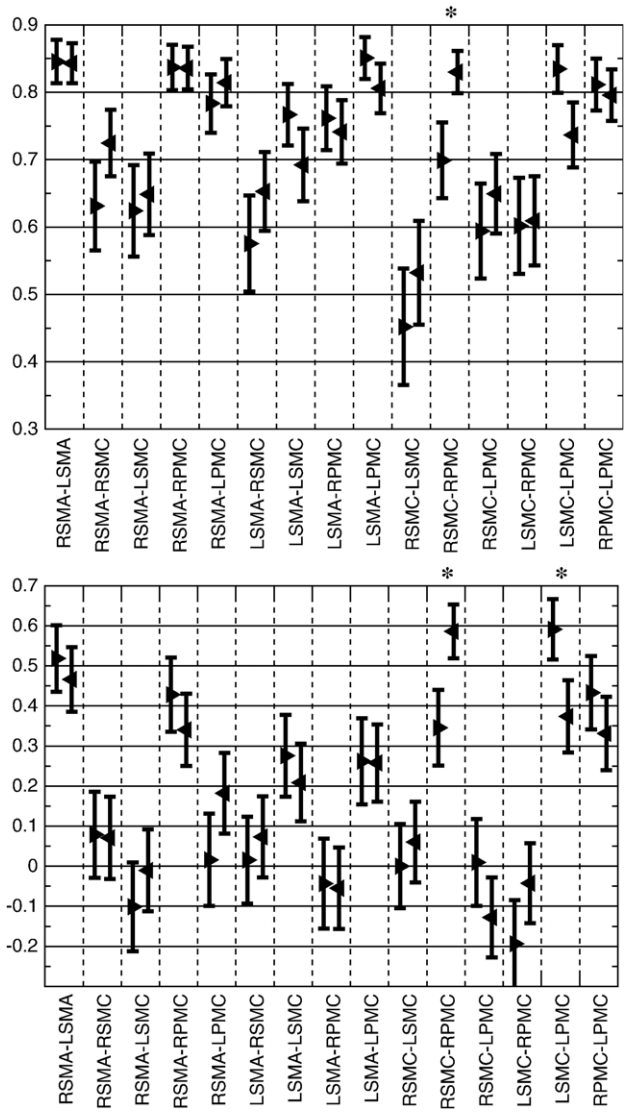

Fig. 1. Estimated marginal (top) and partial (bottom) correlations corresponding to a right-hand (triangle right) or a left-hand (triangle left) movement. A star indicates a significant difference in the level of correlation between a right- and a left-hand movement ($P < 0.05$).

Discussion

We introduced partial correlation analysis as a way to measure the statistical dependencies between two regions after removing the confounding effects of all other regions, hence providing data-driven measures that are closer to effective connectivity than marginal correlation. Given a set of regions and their corresponding time courses, we developed a Bayesian scheme that allowed to infer the underlying interaction structure from fMRI data. The proposed numerical sampling scheme allowed to approximate the posterior distributions of both the group marginal and partial correlation coefficients, $\Pr(R_{0,ij}|\boldsymbol{y})$ and $\Pr(\Pi_{0,ij}|\boldsymbol{y})$.

As expected, the time series of the six cortical regions selected – the two sensorimotor cortices (RSMC and LSMC), the two supplementary motor areas (RSMA and LSMA), and the two premotor areas (RPMC and LPMC) – were strongly correlated, confirming that all six regions did indeed interact, even during simple hand movements. Left motor areas do not usually appear on right-handed healthy subjects' activation maps during right-handed movements. Nevertheless, movements of the non-dominant hand have consistently been reported to induce a much more bilateral activation pattern (Kobayashi et al., 2003; Verstynen et al., 2005). Similarly, participation of these regions for left-hand movements to the recovery of motor skills has amply been documented in stroke patients (Calautti and Baron, 2003; Rossini et al., 2003; Ward, 2005) or patients with a tumor or undergoing surgery (Duffau, 2000; Duffau and Capelle, 2001a,b; Krainik et al., 2004), or with multiple sclerosis (Rocca et al., 2005). More recently, studies of functional connectivity, particularly on healthy subjects in resting state (Biswal et al., 1995; Lowe et al., 1998; Xiong et al., 1999) have confirmed that these regions of the motor network are strongly correlated to one another.

The most striking result of the partial correlation analysis is the seemingly central role of the premotor cortices and, to a slightly lesser extent, the supplementary motor areas. Particularly noticeable is the switch induced by a change in hand: the dependency between SMC and PMC contralateral to the hand movement is significantly higher than that of the other hemisphere. These results, that have been obtained by examination of functional interactions within the cortical motor network, are nonetheless in

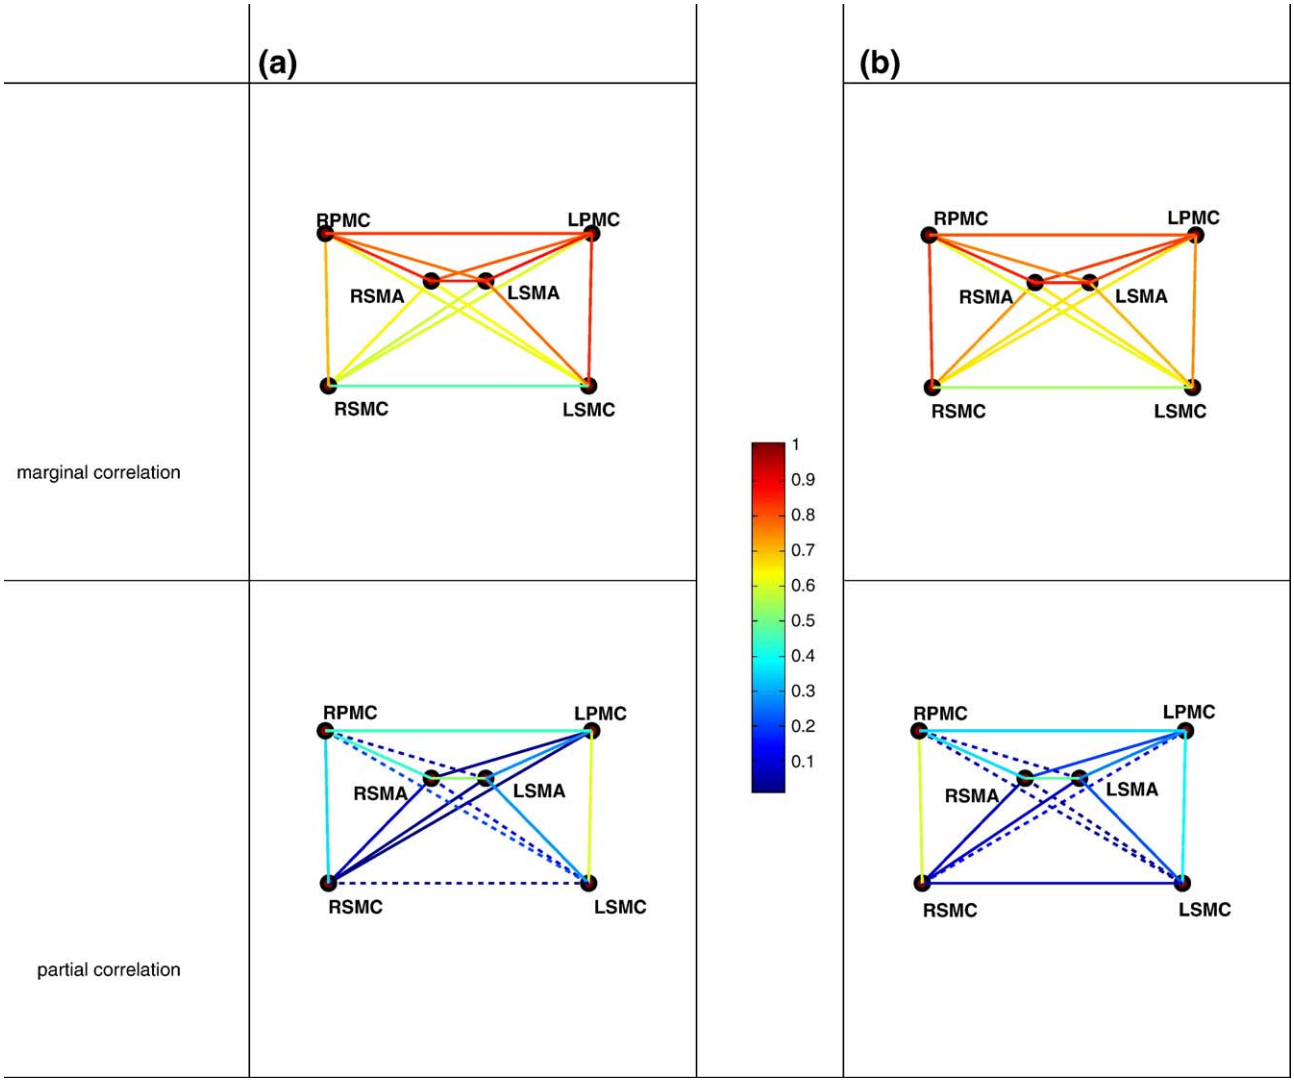

Fig. 2. Estimated mean marginal (top) and partial (bottom) correlations for left and right-hand movements. Positive values between regions are represented by solid links, while negative values are represented by dashed links. The interaction strength is coded according to the color bar.

agreement with the literature on activation of this network. Globally, the results emphasized the central role of the premotor cortex in simple hand movements.

The results also show a preponderance of intrahemispheric and homologous connections, at the expense of non-homologous interhemispheric connections. For instance, strong interhemispheric dependencies between homologous premotor and supplementary areas were observed. This effect has not been observed between sensorimotor cortices and is hence unlikely to be a general, spurious effect from the data. Dependencies between both SMCs remain at a low level, hinting that their connections could be mediated by premotor and secondary motor areas.

Note that partial correlations, by construction, should be more robust than marginal correlation to effects that could affect the

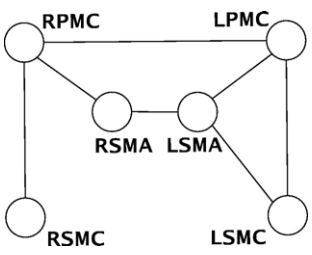

Fig. 3. Partial correlations that are significantly different from zero during left-hand, or, equivalently, right-hand movement ($P < 0.025$).

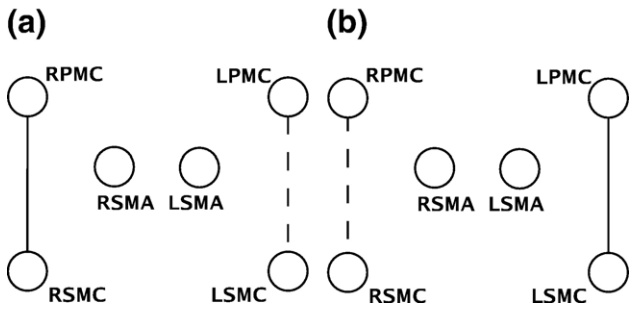

Fig. 4. Partial correlations that significantly change from (a) a right-to a left-hand movement, or (b) a left to a right-hand movement ($P < 0.05$). Correlations that significantly increase are represented in solid lines; those that significantly decrease are represented with dashed lines.

brain over a large scale and hence influence the time course of several brain areas. As a matter of fact, if such an effect is present in the time courses of several regions, then the conditioning property of partial correlation will project the data on a space that is orthogonal to this global, spurious effect, hence removing it, at least partly, from the time series considered. To what extent this property is true with fMRI data, and whether it could allow to produce results that are not affected by, for instance, motion artifacts or stimulus-induced signal, remains to be investigated.

Of course, there still remains the possibility that a dependency detected by partial correlation as significant was actually induced by the effect of a brain area that has not been considered in the analysis. For instance, basal ganglia and the cerebellum are notoriously known, among others, to participate in the motor network, particularly for motor skill learning (Kandel et al., 2000; Doyon et al., 2003). Since these regions have not been considered for this investigation, nothing can be said regarding their influence on the cortical motor network under investigation in this paper. This uncertainty is strongly related to that of the effect of latent variables in graphical models (Whittaker, 1990) and remains an open theoretical issue, that even structural equation modeling has to face. What can be said, however, in the case of partial correlations, is that a significant dependency between two areas translates an interaction that is not mediated by any region that has been included in the analysis. On the other hand, a non-significant dependency would rather hint a lack of direct interaction for, if there were no significant dependency between two areas, there would be no reason to expect any dependency when other areas would be added to the analysis.

Partial correlation analysis allows for a graphical representation of the interaction structure. The connection between these graphs and SEM models remains to be investigated from a practical perspective. Providing such a connection would make it possible to use the data to provide insight into potential structural models underlying the functional network under investigation. The method expounded here should actually be envisioned in a broader procedure for investigation of functional brain interactivity. Indeed, it performs a preliminary, exploratory analysis of interactivity that could conveniently be taken into account when devising a structural model for effective connectivity investigation. Its use would be particularly relevant in cases where model-driven approaches are problematic due to a lack of experimentally designed manipulations, e.g., studies on sleep or hallucinations. Here, SEM or DCM cannot be easily applied because the inputs necessary to drive the system are unknown. For instance, DCM, as it has commonly been applied, cannot be meaningfully performed without an initial statistical parametric map (SPM), since it "reexplains" how the local effects in an SPM result from the connectivity between the regions and its modulation.

One major gap that remains to be bridged for partial correlation to provide data-driven investigation of structural models is related to Step 3. i.e., interaction directionality. Interestingly, determining direct functional interactions and setting their direction are two separate issues. In fMRI data analysis, both concepts appear jointly as required prior information in SEM and DCM, giving the impression that both issues are strongly tied. But we advocate that they can be investigated separately—at least to a certain extent. For instance, Granger causality, as applied in Chávez et al. (2003) in electroencephalography (EEG) to extract causal relationships, cannot decide whether such effects are direct or mediated. Partial correlation does not provide direct information regarding the direction of an effective connectivity. Nonetheless, a given structure of effective connectivity implies a particular temporal behavior that can be characterized to some extent in terms of partial correlation (Whittaker, 1990). We hence believe that the partial correlation approach could eventually been used to retrieve some information of directionality. It would be unable to fully extract the flow of information without any prior information, e.g., in systems such as "A→B→C". It should nonetheless infer from the data that areas A and C do not directly interact. Hence, what would result from: our method would be a (undirected) pathway instead of a flow: "A–B–C", instead of "A→B→C". This pattern of connectivity would rule out "A→B←C" without being able to discriminate between "A→B→C", "A←B←C", and "A←B→C" (Pearl, 2001).

As is the case for SEM, the method proposed here can only calculate one partial correlation structure per run and then compare run structures to one another. When conditions are intertwined, like it is the case for event-related designs, it cannot be applied as easily. For epoch-related designs, a way to circumvent the problem is to separate "on" and "off" blocks and compare the partial correlations calculated for both cases.

The relationship between psycho-physiological interaction (PPI) as proposed by Friston et al. (1997) and partial correlation analysis also remains to be clarified. The goal of PPI is to investigate how a factor (regional activation or experimental factor) modulates the influence of another factor on a regional response. As such, we can see two major differences. The first one is that PPI explicitly takes the paradigm into account, while it is implicit for partial correlation. The other one is that PPI analysis requires a very precise hypothesis to be tested regarding a set of factors. This is hence clearly a model-based method, while our method is data driven.

Despite the fact that partial correlation coefficient extraction can blindly exhibit some patterns of effective connectivity, it still relies on prior anatomical and functional background. For instance, there exists various ways to a priori select the set of regions on which partial correlation analysis, as well as other methods (e.g., SEM and DCM), will take place, e.g., according to group-based or subject-specific anatomical and/or functional criteria. The effects of the selection method onto these methods remain unclear. Another question is what the signal that "represents" the region should be: the time course of the most significantly activated voxel in the region or the spatial mean of time courses over the whole region (Gavrilescu et al., 2004)? Once this has been addressed, it remains to decide on which part connectivity analysis should be performed: the raw signal, or the filtered signal? If filtered, what components should be kept? It is not yet obvious, what part of the signal carries the connectivity information (Arfanakis et al., 2000; Cordes et al., 2000). Nonetheless, these issues are not distinctive of the method expounded here, since they must be considered in SEM or DCM analysis as well. On the other hand, partial correlation provides a way to circumvent the issue of prior selection of a structural model, that SEM or DCM has to face.

Conclusion

In this paper, we proposed to use partial correlations as data-driven measures of functional dependencies that are more closely related to effective connectivity than marginal correlations. In this framework, we measured the interaction strengths between six

cortical regions of the motor network. Once the regions and the corresponding time courses were selected, the Bayesian scheme developed allowed for a fully data-driven procedure that led investigation of dependencies closer to effective connectivity. This technique provided relevant insight into the functional relationships between regions of the motor network and, particularly, confirmed the central role of the premotor cortices during simple-hand movements. Further research includes investigating procedures to integrate this new kind of information to construct structural models for effective connectivity.

Acknowledgments

The authors are in debt to the two anonymous referees for significantly improving the quality of the original manuscript. They are also grateful to Vincent Perlbarg, Odile Jolivet, Saâd Jbabdi, and Pierre Bellec for insightful discussions. Guillaume Marrelec is supported partly by the Ministère de la Recherche, de la Science et de la Technologie du Québec and partly by the Fondation Fyssen (Paris, France). The work was supported by Grant PHRC AOR01109.

Appendix A. Likelihood function of $\boldsymbol{\Sigma}$

$\Pr(\boldsymbol{y}|\boldsymbol{\mu}, \boldsymbol{\Sigma})$, the likelihood of $\boldsymbol{\Sigma}$ and $\boldsymbol{\mu}$, reads:

$$\begin{aligned}\Pr(\boldsymbol{y} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}) &= \prod_{t=1}^{T} \Pr(\boldsymbol{y}_t|\boldsymbol{\mu},\boldsymbol{\Sigma}) \\ &\propto |\boldsymbol{\Sigma}|^{T/2} \exp\left[-\frac{1}{2} \sum_{t=1}^{T} (\boldsymbol{y}_t - \boldsymbol{\mu})^{\mathrm{t}} \boldsymbol{\Sigma}^{-1} (\boldsymbol{y}_t - \boldsymbol{\mu}) \right].\end{aligned}$$

Each term of the sum in the exponential can be expanded as

$$(\boldsymbol{y}_t - \boldsymbol{\mu})^{\mathrm{t}} \boldsymbol{\Sigma}^{-1} (\boldsymbol{y}_t - \boldsymbol{\mu}) = \boldsymbol{\mu}^{\mathrm{t}} \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu} - 2\boldsymbol{\mu}^{\mathrm{t}} \boldsymbol{\Sigma}^{-1} \boldsymbol{y}_t + \boldsymbol{y}_t^{\mathrm{t}} \boldsymbol{\Sigma}^{-1} \boldsymbol{y}_t.$$

The sum hence rereads

$$\sum_{t=1}^{T} (\boldsymbol{y}_t - \boldsymbol{\mu})^{\mathrm{t}} \boldsymbol{\Sigma}^{-1} (\boldsymbol{y}_t - \boldsymbol{\mu}) = T\boldsymbol{\mu}^{\mathrm{t}} \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu} - 2\boldsymbol{\mu}^{\mathrm{t}} \boldsymbol{\Sigma}^{-1} \left[\sum_{t=1}^{T} \boldsymbol{y}_t \right] + \sum_{t=1}^{T} \boldsymbol{y}_t^{\mathrm{t}} \boldsymbol{\Sigma}^{-1} \boldsymbol{y}_t,$$

or, setting $\hat{\boldsymbol{\mu}} = T^{-1} \boldsymbol{\Sigma}_{t=1}^{T} \boldsymbol{y}_t$,

$$\begin{aligned}\sum_{t=1}^{T} (\boldsymbol{y}_t - \boldsymbol{\mu})^{\mathrm{t}} \boldsymbol{\Sigma}^{-1} (\boldsymbol{y}_t - \boldsymbol{\mu}) &= T(\boldsymbol{\mu} - \hat{\boldsymbol{\mu}})^{\mathrm{t}} \boldsymbol{\Sigma}^{-1} (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}) + \sum_{t=1}^{T} \boldsymbol{y}_t^{\mathrm{t}} \boldsymbol{\Sigma}^{-1} \boldsymbol{y}_t - T\hat{\boldsymbol{\mu}}^{\mathrm{t}} \boldsymbol{\Sigma}^{-1} \hat{\boldsymbol{\mu}} \\ &= T(\boldsymbol{\mu} - \hat{\boldsymbol{\mu}})^{\mathrm{t}} \boldsymbol{\Sigma}^{-1} (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}) + \sum_{t=1}^{T} (\boldsymbol{y}_t - \hat{\boldsymbol{\mu}})^{\mathrm{t}} \boldsymbol{\Sigma}^{-1} (\boldsymbol{y}_t - \hat{\boldsymbol{\mu}}).\end{aligned}$$

Setting $\boldsymbol{S} = \boldsymbol{\Sigma}_{t=1}^{T} (\boldsymbol{y}_t - \hat{\boldsymbol{\mu}})(\boldsymbol{y}_t - \hat{\boldsymbol{\mu}})^{\mathrm{t}}$, the sum in the right-hand side of the above equation rereads $\mathrm{tr}(\boldsymbol{S}\boldsymbol{\Sigma}^{-1})$, where tr(.) denotes the trace operator. The likelihood hence reads

$$\Pr(\boldsymbol{y}|\mu,\boldsymbol{\Sigma}) \propto |\boldsymbol{\Sigma}|^{-T/2} \exp\left[-\frac{1}{2} \left(T(\boldsymbol{\mu} - \hat{\boldsymbol{\mu}})^{\mathrm{t}} \boldsymbol{\Sigma}^{-1} (\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}) + \mathrm{tr}\left(\boldsymbol{S}\boldsymbol{\Sigma}^{-1}\right) \right) \right].$$

To obtain the likelihood as a function of Σ only, one has to integrate over $\boldsymbol{\mu}$ as follows:

$$\begin{aligned}\Pr(\boldsymbol{y}|\boldsymbol{\Sigma}) &= \int \Pr(\boldsymbol{y}, \boldsymbol{\mu}|\boldsymbol{\Sigma}) \mathrm{d}\boldsymbol{\mu} \\ &= \int \Pr(\boldsymbol{y}|\boldsymbol{\mu},\boldsymbol{\Sigma}) \cdot \Pr(\boldsymbol{\mu}|\boldsymbol{\Sigma}) \mathrm{d}\boldsymbol{\mu}.\end{aligned}$$

Assuming no dependence a priori between $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$, $\Pr(\boldsymbol{\mu}|\boldsymbol{\Sigma})$ is classically simplified as $\Pr(\boldsymbol{\mu})$ and set to a uniform, non-informative distribution. Integration with respect to $\boldsymbol{\mu}$ finally yields the desired likelihood:

$$\Pr(\boldsymbol{y}|\boldsymbol{\Sigma}) \propto |\boldsymbol{\Sigma}|^{-\frac{T-1}{2}} \exp\left[-\frac{1}{2} \mathrm{tr}(\boldsymbol{S}\boldsymbol{\Sigma})^{-1} \right],$$

which is an inverse Wishart distribution with $T-D-2$ degrees of freedom and scale matrix $\boldsymbol{U} = \boldsymbol{S}^{-1}$ (Gelman et al., 1998, Appendix A).

References

Arfanakis, K., Cordes, D., Haughton, V.M., Moritz, C.H., Quigley, M.A., Meyerand, M.E., 2000. Combining independent component analysis and correlation analysis to probe interregional connectivity in fMRI task activation datasets. Magn. Reson. Imaging 18, 921–930.

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

Biswal, B.B., Kylen, J.V., Hyde, J.S., 1997. Simultaneous assessment of flow and BOLD signals in resting-state functional connectivity maps. NMR Biomed. 10, 165–170.

Bullmore, E., Horwitz, B., Honey, G., Brammer, M., Williams, S., Sharma, T., 2000. How good is good enough in path analysis of fMRI data? NeuroImage 11, 289–301.

Calautti, C., Baron, J.-C., 2003. Functional neuroimaging studies of motor recovery after stroke in adults. Stroke 34, 1553–1566.

Chávez, M., Martinerie, J., Le Van Quyen, M., 2003. Statistical assessment of nonlinear causality: application to epileptic EEG signals. J. Neurosci. Methods 124, 113–128.

Chen, W., Ogawa, S., 1999. Principles of BOLD functional MRI. In: Moonen, C., Bandettini, P. (Eds.), Functional MRI. Springer, Berlin, pp. 103–113.

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.

Doyon, J., Penhune, V., Ungerleider, L.G., 2003. Distinct contribution of the cortico-striatal and cortico-cerebellar systems to motor skill learning. Neuropsychologia 41, 252–262.

Duffau, H., 2000. Intraoperative direct subcortical stimulation for identification of the internal capsule, combined with an image-guided stereotactic system during surgery for basal ganglia lesions. Surg. Neurol. 53, 250–254.

Duffau, H., Capelle, L., 2001a. Récupération fonctionnelle à la suite de lésions de l'aire somato-sensorielle primaire: étude des mécanismes de compensation. Neuro-Chirurgie 47, 557–563.

Duffau, H., Capelle, L., 2001b. Récupération fonctionnelle après résection de gliomes infiltrant l'aire somato-sensorielle primaire (SI): étude par stimulations électriques per-opératoires. Neuro-Chirurgie 47, 534–541.

Frackowiak, R.S.J., Friston, K.J., Frith, C.D., Dolan, R.J., Price, C.J., Zeki, S., Ashburner, J., Penny, W. (Eds.), 2004. Human Brain Function, 2nd ed. Academic Press.

Friston, K.J., Frith, C.D., Frackowiak, R.S.J., 1993a. Time-dependent changes in effective connectivity measured with PET. Hum. Brain Mapp. 1, 69–79.

Friston, K.J., Frith, C.D., Liddle, P.F., Frackowiak, R.S.J., 1993b. Functional connectivity: the principal component analysis of large (PET) data sets. J. Cereb. Blood Flow Metab. 13, 5–14.

Friston, K.J., Jezzard, P., Turner, R., 1994. Analysis of functional MRI time-series. Hum. Brain Mapp. 1, 153–171.

Friston, K.J., 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 modelling. NeuroImage 19, 1273–1302.

Gavrilescu, M., Stuart, G.W., Waites, A., Jackson, G., Svalbe, I.D., Egan, G.F., 2004. Changes in effective connectivity models in the presence of task-correlated motion: an fMRI study. Hum. Brain Mapp. 21, 49–63.

Gazzaniga, M.S. (Ed.), 2000. The New Cognitive Neurosciences, 2nd ed. The MIT Press, Cambridge, MA.

Gelman, A., Carlin, J.B., Stern, H.S., Rubin, D.B., 1998. Bayesian Data Analysis. Chapman and Hall, London.

Gonzalez-Lima, F., McIntosh, A.R., 1995. Analysis of neural interactions related to associative learning using structural equation modeling. Math. Comput. Simul. 40, 115–140.

Hebb, D.O., 1949. The Organization of Behavior: A Neurophysiological Theory. Wiley, New York.

Horwitz, B., 2003. The elusive concept of brain connectivity. NeuroImage 19, 466–470.

Huettel, S.A., Song, A.W., McCarthy, G., 2004. Functional Magnetic Resonance Imaging. Sinauer, Sunderland.

Kandel, E.R., Schwartz, J.H., Jessell, T.M., 2000. Principles of Neural Science, 4th ed. McGraw-Hill, New York.

Kobayashi, M., Hutchinson, S., Schlaug, G., Pascual-Leone, A., 2003. Ipsilateral motor cortex activation on functional magnetic resonance imaging during unilateral hand movements is related to interhemispheric interactions. NeuroImage 20, 2259–2270.

Krainik, A., Duffau, H., Capelle, L., Cornu, P., Boch, A.-L., Mangin, J.-F., Bihan, D.L., Marsault, C., Chiras, J., Lehéricy, S., 2004. Role of the healthy hemisphere in recovery after resection of the supplementary motor area. Neurology 62, 1323–1332.

Kunzle, H., 1978. An autoradiographic analysis of the efferent connections from premotor and adjacent prefrontal regions (areas 6 and 9) in *Macaca fascicularis*. Brain Behav. Evol. 15, 185–234.

Lauritzen, S.L., 1996. Graphical Models. Oxford Univ. Press, Oxford.

Lee, L., Harrison, L.M., Mechelli, A., 2003. A report of the functional connectivity workshop, Düsseldorf 2002. NeuroImage 19, 457–465.

Leichnetz, G.R., 1986. Afferent and efferent connections of the dorsolateral precentral gyrus (area 4, hand/arm region) in the macaque monkey, with comparisons to area 8. J. Comp. Neurol. 254, 460–492.

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.

Marrelec, G., Daunizeau, J., Pélégrini-Issac, M., Doyon, J., Benali, H., 2005a. Conditional correlation as a measure of mediated interactivity in fMRI and MEG/EEG. IEEE Trans. Signal Process. 53, 3503–3516.

Marrelec, G., Doyon, J., Pélégrini-Issac, M., Benali, H., 2005b. Heading for data-driven measures of effective connectivity in functional MRI. Proceedings of the International Joint Conference on Neural Networks, pp. 1528–1533.

McIntosh, A.R., Bookstein, F.L., Haxby, J.V., Grady, C.L., 1996. Spatial pattern analysis of functional brain images using partial least squares. NeuroImage 3, 143–157.

Mechelli, A., Penny, W.D., Price, C.J., Gitelman, D.R., Friston, K.J., 2002. Effective connectivity and intersubject variability: using a multisubject network to test differences and commonalities. NeuroImage 17, 1459–1469.

Naidich, T.P., Hof, P.R., I. Yousry, T.A.Y., 2001. The motor cortex: anatomic substrates of function. Neuroimaging Clin. N. Am. 11, 171–193.

Neal, R.M., 2000. On deducing conditional independence from d-separation in causal graphs with feedback. J. Artif. Intell. Res. 12, 87–91.

Pearl, J., 2001. Causality: Models, Reasoning, and Inference. Cambridge Univ. Press, Cambridge.

Pearl, J., Dechter, R., 1996. Identifying independencies in causal graphs with feedback. In: Horvitz, E., Jensen, E.F. (Eds.), Proceedings of the Twelfth Conference on Uncertainty in Artificial Intelligence. M. Kaufmann, San Francisco, pp. 240–246.

Penny, W.D., Stephan, K.E., Mechelli, A., Friston, K.J., 2004a. Comparing dynamic causal models. NeuroImage 22, 1157–1172.

Penny, W.D., Stephan, K.E., Mechelli, A., Friston, K.J., 2004b. Modelling functional integration: a comparison of structural equation and dynamic causal models. NeuroImage 23 (S1), S264–S274.

Rocca, M.A., Colombo, B., Falini, A., Ghezzi, A., Martinelli, V., Scotti, G., Comi, G., Filippi, M., 2005. Cortical adaptation in patients with MS: a cross-sectional functional MRI study of disease phenotypes. Lancet Neurol. 4, 618–626.

Rogers, B.P., Carew, J.D., Meyerand, M.E., 2004. Hemispheric asymmetry in supplementary motor area connectivity during unilateral finger movements. NeuroImage 22, 855–859.

Rossini, P.M., Calautti, C., Pauri, F., Baron, J.-C., 2003. Post-stroke plastic reorganisation in the adult brain. Lancet Neurol. 2, 493–502.

Ruanaidh, J.J.K.O., Fitzgerald, W.J., 1996. Numerical Bayesian Methods Applied to Signal Processing. Statistics and Computing. Springer, New York.

Spirtes, P., 1995. Directed cyclic graphical representations of feedback models. In: Besnard, P., Hanks, S. (Eds.), Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence. M. Kaufmann, San Mateo, pp. 491–498.

Sporns, O., Chialvo, D.R., Kaiser, M., Hilgetag, C.C., 2004. Organization, development and function of complex brain networks. Trends Cogn. Sci. 8, 418–425.

Stone, J.V., Kötter, R., 2002. Making connections about brain connectivity. Trends Cogn. Sci. 6, 327–328.

Sun, F.T., Miller, L.M., D'Esposito, M., 2004. Measuring interregional functional connectivity using coherence and partial coherence analyses of fMRI data. NeuroImage 21, 647–658.

Swinnen, S.P., 2002. Intermanual coordination: from behavioural principles to neural-network interactions. Nat. Rev., Neurosci. 3, 350–361.

Talairach, J., Tournoux, P., 1988. Co-Planar Stereotaxic Atlas of the Human Brain. Georg Thieme, Stuttgart.

Tononi, G., Edelman, G.M., Sporns, O., 1998. Complexity and coherence: integrating information in the brain. Trends Cogn. Sci. 2, 474–484.

Verstynen, T., Diedrichsen, J., Albert, N., Aparicio, P., Ivry, R.B., 2005. Ipsilateral motor cortex activity during unimanual hand movements relates to task complexity. J. Neurophysiol. 93, 1209–1222.

Ward, N.S., 2005. Plasticity and the functional reorganization of the human brain. Int. J. Psychophysiol. 58, 158–161.

Whittaker, J., 1990. Graphical Models in Applied Multivariate Statistics. J. Wiley and Sons, Chichester.

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.

Zigmond, M.J., Bloom, F.E., Landis, S.C., Roberts, J.L., Squire, L.R. (Eds.), 1999. Fundamental Neuroscience. Academic Press, London.