Available online at www.sciencedirect.com
ScienceDirect
Journal of Physiology - Paris 100 (2006) 171–181
Journal of Physiology Paris
www.elsevier.com/locate/jphysparis

Exploring large-scale brain networks in functional MRI

Guillaume Marrelec [a,b,c], Pierre Bellec [b,c], Habib Benali [a,b,c,*]

[a] Université de Montréal, MIC/UNF, Montreal, Canada H3W 1W5
[b] Inserm, U678, Paris F-75013, France
[c] Université Pierre et Marie Curie, Faculté de médecine Pitié-Salpêtrière, Paris F-75013, France

Abstract

Increasing emphasis has been recently put on large-scale network processing of brain functions. To explore these networks, many approaches have been proposed in functional magnetic resonance imaging (fMRI). Their objective is to answer the following two questions: (1) what brain regions are involved in the functional process under investigation? and (2) how do these regions interact? We review some of the key concepts and corresponding methods to cope with both issues.

Keywords: Functional brain interactions; Large-scale networks; Functional MRI; Functional connectivity; Effective connectivity; Conditional correlation

1. Introduction

During the past decade, investigation of brain activity has put increasing emphasis on the analysis of interactions within large-scale networks of brain areas (Horwitz et al., 1999; Varela et al., 2001). By defining neural assemblies as distributed local networks transiently linked by large-scale reciprocal dynamical connections, Varela et al. (2001) make a clear distinction between *local* and *large* scale. On the one hand, a local network is defined as a large patch ($\lesssim$1 cm) of neural tissue that synchronizes its activity through local cytoarchitecture. This definition strongly relates to that of Hebbian cell assemblies, i.e., groups of entities (neurons) acting together in a coherent fashion (Hebb, 1949). On the other hand, large-scale dynamic connections are defined as interactions based on large fiber pathways between regions that are located far apart from one another (>1 cm). The dichotomy between local and large-scale networks serves as a neural basis for the key assumption that brain functional architecture abides by two principles: functional segregation and functional integration. While the segregation principle states that some functional processes specifically engage well localized and specialized brain regions, it is now believed that brain functions are most likely to emerge through integration of information flows across distributed regions (Tononi et al., 1998a; Varela et al., 2001; Frackowiak et al., 2004; Sporns et al., 2004). In this approach, it is not only a collection of brain areas that is hypothesized to process functional tasks, but rather a large-scale network, i.e., a set of brain regions dynamically interacting with one another.

The concepts of segregation and integration quickly became central in neuroimaging, which was able to sample evoked responses over the entire brain at the same time (Tononi et al., 1992; Friston et al., 1993a). Neuroimaging encompasses various techniques that allow to dynamically and noninvasively follow various markers of brain activity. Functional magnetic resonance imaging (fMRI), on which this paper will focus, is one of the major methods currently used in research and clinical routine. Most fMRI acquisitions rely on the so-called blood oxygen level dependent (BOLD) contrast, which measures metabolic and hemodynamic consequences of brain activity (Chen and Ogawa, 1999; Huettel et al., 2004). Recently, compelling experimental evidence has been brought to support the fact that BOLD signal roughly reflects the slow fluctuations of local

* Corresponding author. Address: CRIUGM, 4545 Chemin Queen-Mary, Montréal, QC, Canada H3W 1W5.
E-mail address: habib.benali@imed.jussieu.fr (H. Benali).

0928-4257/$ - see front matter © 2007 Published by Elsevier Ltd.
doi:10.1016/j.jphysparis.2007.01.003

field potentials (Logothetis et al., 2001), i.e., chiefly the integrated synaptic activity over a few millimeters of neural tissue. While many approaches have been proposed so far to model neural activity (both at a cellular and a population level), induced metabolic and vascular responses, and the physical process leading to fMRI measurements (Li et al., 2000; Attwell and Iadecola, 2002; Aubert and Costalat, 2002; Shulman et al., 2004), there still is no unanimously accepted model connecting neuronal activation, EEG/MEG signals and BOLD fMRI measurements (Horwitz, 2005).

On the basis of the information conveyed by the BOLD signal regarding neuronal activity, it is increasingly accepted that brain interactions can be captured by fMRI, at least to a certain extent. This has resulted in a new investigation field that has drawn increasing interest during recent years (see, e.g., Stone and Kötter, 2002; Horwitz, 2003; Lee et al., 2003a,b; Horwitz et al., 2005). Its objective is to capture the dynamic, context-dependent processes leading to the preferential involvement of some regions over others (McIntosh and Gonzalez-Lima, 1994; Varela et al., 2001; Frackowiak et al., 2004; Ramnani et al., 2004; Sporns et al., 2004). Basically, investigation of functional networks requires answering of the following two questions: (1) what brain regions are involved in the functional process under investigation? and (2) how do these regions interact? The various methods that have been proposed to tackle brain interactions can be classified according to which of these two questions they address. In this paper, we review some of the key concepts, methods, and recent advances relative to both issues.

Answering Question 1 boils down to determining what regions form the spatial support of the network investigated. In the next section, we introduce the most common techniques, which are based on either anatomical information and/or activation maps. We also review other approaches that do not rely on a comparison to a control condition but attempt to identify salient large-scale networks in an exploratory way; such methods include connectivity maps (Biswal et al., 1995) or independent component analysis (McKeown et al., 1998). In this review, we pay special attention to the mathematical approaches utilized to characterize the existence of large-scale networks. We also focus on recent advances that aim to explicitly model and take into account the main acknowledged features of large-scale neural networks as observed in fMRI (Bellec et al., 2006).

Once the network regions have been determined, Question 2 deals with quantitative measures of inter-regional interactions. As emphasized by Horwitz (2003), brain interactions have many definitions in the existing literature, depending on the authors' perspective. Nonetheless, most approaches rely on either functional or effective connectivity for reviews and discussions, see, e.g., (Stone and Kötter, 2002; Horwitz, 2003; Lee et al., 2003b). The functional connectivity between two voxels or regions is defined as the temporal correlation between the voxel or region time courses (Friston et al., 1993b). As for effective connectivity, it rather considers the influence that regions exert on each other (Friston et al., 1993a; McIntosh and Gonzalez-Lima, 1994). Retracing the major differences between effective and functional connectivity, we review tentative approaches to bridge the gap between both frameworks.

2. Identification of large-scale networks

As emphasized in the introduction, examination of functional interactions within a distributed network can be decomposed into two steps. Step 1 consists of defining the underlying support of the network, i.e., the set of regions that are involved in the brain function under scrutiny. Depending on the amount of information available, various methods can be applied to this end.

2.1. From knowledge-based to exploratory methods

A first approach is to make use of relevant information that has been gathered by previous experience or experiments. Such information can be obtained either from anatomical or functional considerations. Anatomically, regions that are known to be part of a network can be segmented according to structural features, e.g., in terms of their cytoarchitecture, as is the case for Brodmann's areas, (e.g., McIntosh et al., 1996; Salvador et al., 2005). Functionally, regions detected by conventional activation analyses share a temporal pattern that is specifically modulated by a task as compared to a control condition (Friston et al., 1995). Such regions should be considered for inclusion into the network (e.g., Büchel and Friston, 1997). This approach can be restrictive because it requires explicit modeling of the temporal activity with regard to both the task and the control condition. Furthermore, it often only extracts a restricted number of highly specific regions, while neglecting regions that are part of the network but less strongly related to the experimental protocol (Horwitz et al., 1999; Gusnard and Raichle, 2001). There are also conditions for which we have only little spatial or functional prior information; such is the case for, e.g., data gathered at rest (Biswal et al., 1995), or data of patients in coma (Laureys et al., 2004).

Others methods exist to select the support of a functional network. Pioneered by early works in PET imaging (Clark et al., 1984; Metter et al., 1984; Horwitz et al., 1984), functional connectivity makes it possible to explore which regions strongly interact with a defined region during a given condition without reference to a control condition. Biswal et al. (1995) introduced functional connectivity maps to explore the network of regions that were functionally related to a seed region located in the primary motor cortex. A functional 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, determined either empirically or statistically, is applied to the map in order to identify the network

of brain regions functionally connected (in the sense of correlation) to the seed. It was suggested that such a network includes mostly regions possessing strong anatomical connections to the seed, either directly or indirectly (Xiong et al., 1999). This technique has been particularly influential for the analysis of resting-state datasets, where subjects were asked to refrain from overt activity. It has been applied for a variety of seed regions, located in motor (Biswal et al., 1995; Xiong et al., 1999), visual (Lowe et al., 1998), language (Cordes et al., 2000) and cingulate (Greicius et al., 2003) cortices, as well as sub-cortical regions (Stein et al., 2000). By contrast, only few studies have investigated correlation maps for subjects steadily performing a given task (e.g., Lowe et al., 2000; Greicius et al., 2003). Networks exhibited with this method tend to be much larger than the set of activated regions (Xiong et al., 1999).

Although connectivity maps have proved to be a powerful tool, their exact relationship with functionally relevant large-scale networks remains to be further investigated. Connectivity map exploration heavily relies on the choice of a seed region. Indeed, the use of regions of arbitrary shape, even restricted to small areas declared activated by fMRI activation analyses, can lead to spatially inhomogeneous regions (Baumgartner et al., 2000). It has furthermore been shown that two small and spatially close brain regions can lead to very dissimilar functional connectivity (e.g., Cordes et al., 2000; Gonçalves and Hall, 2003; Waites et al., 2005). In McIntosh et al. (1996) and Friston et al. (1997), psycho-physiological parameters are considered when using a general linear model or a partial least-squares model. Partial least-squares is a singular value decomposition of the cross-correlations between physiological responses and behavior. Psycho-physiological interaction analyses test for changes in the regression slope of activity, at every voxel on a seed voxel, that are induced by an experimental manipulation. Even if these techniques are more flexible than the method of correlation with a target area, they remain strongly dependent on the choice of the target region(s) and/or the psychophysical parameter(s).

Other approaches for identifying large-scale patterns of functional connectivity exist that do not rely on a seed region and, hence, reduce the dependence of the analysis on this parameter. Such methods include principal components analysis (PCA) (Friston et al., 1993b), independent components analysis (ICA) (McKeown et al., 1998) and fuzzy clustering, e.g., *k*-means (Baumgartner et al., 1998). They were initially developed in the general framework of multivariate statistics and assume various models of the fMRI data. While PCA and ICA suppose that fMRI data are a linear mixing of a given number of temporal factors with an associated factor-specific spatial distribution, *k*-means assumes that brain voxels can be grouped into clusters sharing similar activity patterns. Practical factor-image decomposition or clustering of the PCA, ICA, or *k*-means is achieved by optimizing the following mathematical criteria: maximal variance of the data after projection in an uncorrelated spatial basis, independence of the spatial distributions in a linear mixing, and intra-cluster homogeneity, respectively. These methods, and most notably independent component analysis, have led to promising results, e.g., performing blind identification of networks that had already been exhibited with functional connectivity maps and a well-defined seed region (Greicius et al., 2004). It is nevertheless necessary to specify the number of factors or clusters. Moreover, among all patterns identified using such approaches, not all of them have a clear relationship with brain neural activity (McKeown et al., 1998). Final identification of the relevant factors or clusters is mainly performed visually, even if some automatic procedures have recently been proposed that apply in some particular contexts (e.g., Greicius et al., 2004). The lack of a clear and systematic relationship between the mathematical criteria optimized by the aforementioned approaches and the neuroscientific concept of large-scale networks makes the corresponding results arduous to interpret from a general perspective. These techniques were developed in a more general context and for different purposes; the mathematical criteria optimized often cannot be expressed in terms of large-scale functional connectivity. Furthermore, such methods are completely independent of the relative localization of the selected regions with each other – a central element for large-scale networks.

2.2. Large-scale networks and statistical modeling

Recent works have proposed to unambiguously embed neuroscientific considerations into statistical models. For example, Tononi et al. (1998b) used a measure of functional integration derived from information theory to define the concept of functional clusters; Goutte et al. (2001) clustered brain voxels on the basis of the similarity of their hemodynamic response. Concerning such features of neural assemblies as local and large-scale, what has been mentioned in the previous section is compelling evidence in favor of the fact that correlation may partly reflect neuronal interactions. In this framework, synchrony of neuronal activity within local networks imposes that each region of the network should be homogeneous, i.e., composed of voxels whose time series are highly correlated, e.g., according to the Kendall coefficient of concordance (Zang et al., 2004).

The existence of large-scale interactions moreover implies that each region in the network exhibits strong correlation with other distant region(s) in the network. As stressed by Lund (2001) and Cordes et al. (2002), a major source of confound for large-scale functional connectivity is thus the spurious spatial correlation induced by fMRI noise (e.g., cardiac or respiratory fluctuations, local interactions), which needs to be properly corrected and/or modeled to achieve accurate identification of large-scale networks. With this aim, Bellec et al. (2006) used a Gaussian process with a stationary spatial correlation matrix function of the sole distance between regions, or lag (Cressie, 1993). Three parameters were involved in the modeling and estimated from the data (see Fig. 1): a local

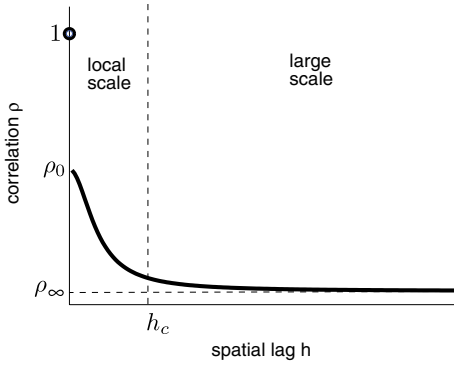

Fig. 1. Rational-quadratic correlogram: the spurious correlations between the time series of two regions are modeled as a parametric function of the spatial distance (or lag) between the regions. Graphical interpretation of the parameters ρ_0, ρ_∞, and h_c is shown in the figure (from Bellec et al., 2006).

correlation level ρ_0, an asymptotic level ρ_∞ at large-scale, and a critical distance h_c characteristic of the transition between local and large-scale. A pair of regions was then said to be on a large scale if their distance was larger than h_c. The set of regions exhibiting an outlier correlation with at least one distant region then composed the large-scale network. With fMRI datasets acquired during simple motor protocols, this approach performed blind identification of networks both related to the task and subserving spontaneous activity, such as the default-mode network described at rest (Greicius et al., 2003); the networks discovered shared many similarities with those extracted using other exploratory approaches, such as ICA (Greicius et al., 2004). Yet, unlike ICA, all regions involved in a salient large-scale network could be identified at once, which typically represent tens or hundreds of regions for a region size of approximately 1 cm^2 of cortex. This unique behavior allowed first time blind exploration of all functional brain networks engaged in functional processing during an fMRI acquisition. Unfortunately, apprehending very large networks is necessarily a complex and tedious task for which existing tools are not well adapted. Consequently, routine tasks, such as validating the results in the light of previous anatomical/functional information or comparing them between subjects, can quickly prove a challenge. A possible remedy is to use additional data analyses to summarize the rich interaction patterns, such as hierarchical clustering (Goutte et al., 1999; Cordes et al., 2002), or multi-dimensional scaling (Friston et al., 1996; Welchew et al., 2002; Salvador et al., 2005).

3. Measuring functional interactions

Once the network regions have been selected, functional connectivity appears as a simple way to measure interactions within the resulting large-scale network. Yet, a significant functional connectivity between two regions can have various exogenous origins: a common response to the same external stimulation, a common input, or an indirect interaction mediated by a third region (Xiong et al., 1999; Marrelec et al., 2005a). When simply assessing the presence or absence of functional interactions regardless of their nature (direct, indirect, or stimulus-locked), functional connectivity might prove sufficient; as evidenced in the previous section, many methods indeed rely on it to extract the large-scale network. However, the full understanding of the network interaction structure requires to further disambiguate the origin of the observed functional connectivity and to determine patterns of effective connectivity, i.e., the effect that regions exert on one another (Friston et al., 1993a). Unfortunately, while functional connectivity is data-driven, major methods implemented to investigate effective connectivity are model-based, hindering their use.

In this perspective, we first review the major approaches developed to examine effective connectivity. We then report some first steps toward measures of connectivity that, like functional connectivity, are data-driven and moreover share some key aspects of effective connectivity that functional connectivity fails to apprehend.

3.1. Effective connectivity, SEM, and DCM

Structural equation modeling (SEM) is the most widespread way to model effective connectivity (McIntosh and Gonzalez-Lima, 1994; Gonzalez-Lima and McIntosh, 1995; Bullmore et al., 2000). Starting from a set of D regions, a model is set a priori, expressing the time course $y_d(t)$ of each region as a linear function of the time courses of other regions:

$$y_d(t) = \sum_{e \neq d} \lambda_{d,e} y_e(t),$$

some coefficients $\lambda_{d,e}$ being constrained to 0, the others being free to vary. $\lambda_{d,e}$ quantifies the strength that region e exerts on region d, whence the name of effective connectivity. The free parameters $\lambda_{d,e}$ can be translated in a graphical model as arrows joining node d to node e (see Fig. 2). The graph links are usually thought of as anatomical connections that are functionally relevant for the experiment under consideration. The model of effective connectivity, fully determined by the $\lambda_{d,e}$, is associated with a unique

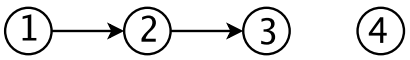

regions	interaction?	correlation	conditional correlation
1 and 4	no	$\mathrm{Corr}[y_1, y_4] = 0$	
1 and 2	yes, direct	$\mathrm{Corr}[y_1, y_2] \neq 0$	$\mathrm{Corr}[y_1, y_2 \vert y_3] \neq 0$
1 and 3	yes, indirect (mediated by 2)	$\mathrm{Corr}[y_1, y_3] \neq 0$	$\mathrm{Corr}[y_1, y_3 \vert y_2] = 0$

Fig. 2. Example: structural model, e.g., for movement-related information processing in the visual pathway, from (1) V1 to (2) V2 to (3) V5, and with (4) the primary auditory region on Heschl's gyrus. Measures of functional interaction corresponding to the various interaction patterns graphically represented (from Marrelec et al., 2005a).

matrix of inter-regional covariances. This parametric matrix can be compared to the one observed empirically in the data; by minimizing the discrepancy between both, one can then estimate the model coefficients (Cudeck et al., 1993; Bullmore et al., 2000). Performing this same step for two different experimental settings, it is possible to analyze the influence of the protocol modification on $\lambda_{d,e}$ and, consequently, on the actions exerted by region e on region d. To release the constraint of linear relationships between variables, the same framework has been expanded in order to investigate how a factor (regional activation or experimental parameter) modulates the influence of another factor on a regional response (Büchel and Friston, 1997; Friston et al., 1997).

The major drawback of SEM is that it is strongly model-based. Apart from the fact that SEM usually assumes linear interactions between regions and does not account for the dynamic aspect of hemodynamic responses, it moreover requires the setting of an a priori model, that will — and can — difficultly be challenged. SEM cannot be used in cases where no information of functional interactions is available. Neither can it deal with cases where the number of constraints imposed on the connections is low, since one needs a very sparse connectivity structure to make the effective connectivity parameters estimable. This places strong constraints on the structural equation models that can be used. An emerging literature trying to cope with these issues has recently emerged (Bullmore et al., 2000; Mechelli et al., 2002). These methods are useful when some information is available, such as a small set of potential structural models or partial information relative to the connectivity. They remain not well adapted for complex and/or exploratory analyses. Another technical point that, to our knowledge, has never been reported in the fMRI literature, is the commonplace use of the correlation matrix instead of the covariance matrix for SEM analysis. This change is largely justified by the very nature of BOLD signal, which can only be interpretated in terms of percentage of increase or decrease compared to a baseline. Still, from a theoretical perspective, using either the correlation or the covariance matrix defines two different inference processes that can potentially lead to different results (Cudeck, 1989). The bearing of this issue on fMRI data analysis remains to be further investigated.

Another framework for effective connectivity analysis, dynamical causal modeling (DCM) has recently been developed as a generalization of both convolution models and SEM (Friston, 2003; Penny et al., 2004b). DCM also relies on the definition of a structural model in the form of a directed graph prior to the analysis. The model then assumes a dynamic neuronal model of interacting brain regions, whereby neuronal activity in a given brain region causes changes in neuronal activity in other regions according to the graphical model. This neuronal model is then supplemented with a forward model of how neuronal activity generates a measured BOLD response through the balloon model (Friston, 2003). Last, a Bayesian inference scheme is devised to infer the model parameters from the data. It is suspected that DCM might be less sensitive than SEM to the number of degrees of freedom (i.e., the sparsity of the connectivity matrix), even though this remains to be formally proved. More generally, DCM is quite recent and has only led to few studies so far (Mechelli et al., 2003b; Ethofer et al., 2006; Lee et al., 2006). Further research is thus still required to better ascertain its main strengths and weaknesses.

By construction, the mathematical framework of DCM takes nonlinearities and temporal correlations into account. It also quantifies the interaction strength that one brain region exerts on another brain region at the neuronal level, whereas SEM remains at the level of the observed BOLD signal. Unlike SEM, DCM also models the effect of experimental, external, and/or modulatory inputs on network dynamics. A critical feature of the proposed forward balloon model is the relationship between blood flow changes and oxygen metabolism changes during activation. The given forward model might only be well adapted to the steady state condition (Aubert and Costalat, 2002). Further issues regarding identifiability and estimability of this complex model have yet to be coped with. Since DCM takes dynamics and modulations into account in the model and some part of the uncertainty in the inference, this framework is even more complex than SEM. As a consequence, DCM is computationally limited by the number of regions that can be included in the analysis (maximum of eight according to Penny et al. (2004b); three in Mechelli et al. (2003a), Penny et al. (2004b), and Ethofer et al. (2006); three and five in Lee et al. (2006)). To cope with the thorny issue of model selection, Penny et al. (2004a) proposed an extension of the DCM framework to perform model comparison within a set of graphs given a priori. How this approach can be generalized to allow for blind model selection from the whole set of structural models (i.e., with no structural model required a priori) remains a central, yet complex, issue.

A feature that is noteworthy about SEM and DCM is that they both try to simultaneously handle two different concepts in one step: direct interactions and causality. On the other hand, data-driven measures have so far concentrated on the exploration of either direct interactions or causal relationships.

3.2. Data-driven measures of direct interaction

As correlation was defined, it is possible to define conditional covariance and correlation (Anderson, 1958). The conditional correlation $\mathrm{corr}[y_i, y_j | y_{\mathscr{R}}]$ between regions i and j with respect to a set $\mathscr{R}$ of regions measures the residual correlation between the time courses of regions i and j once the (linear) effect of regions in $\mathscr{R}$ has been removed. It hence only considers the correlation between i and j that cannot be accounted for by the (linear) influence of any area in $\mathscr{R}$. Each structural model entails a unique pattern in terms of conditional correlation, which is closely related

to path coefficients (e.g., McIntosh and Gonzalez-Lima, 1994; Bullmore et al., 2000; Marrelec et al., 2005a,b). Marrelec et al. (2005a) demonstrated that, unlike (marginal) correlation, conditional correlation could successfully retrieve information of mediation from a structural model. Consequently, they hypothesized that, by conditioning the dependencies between two areas on other areas, the ensuing conditional correlation should be more closely related to direct interaction. Specifically, if functional interactions between regions i and j are mediated by a set $\mathscr{R}$ of regions, then $\mathrm{corr}[y_i, y_j | y_{\mathscr{R}}]$, the conditional correlation between regions i and j given $\mathscr{R}$, is zero. Conversely, if some interactions between the same two regions is not mediated by regions in $\mathscr{R}$, then a natural measure of how much interaction is not mediated by this set of regions is given by $\mathrm{corr}[y_i, y_j | y_{\mathscr{R}}]$ (see Fig. 2). This theoretical assumption remains to be confirmed on synthetic as well as real data. In order to do so, several issues need to be solved. Starting from N regions, $N(N-1)\cdot 2^{N-2}/2$ conditional correlation coefficients could potentially be calculated, many of which probably redundant with one another. Without a structural model to guide the investigation, calculation of this whole set remains lengthy, not to say untractable, and interpretation tedious. Marrelec et al. (2006) proposed to circumvent this issue using partial correlation (i.e., correlation between two regions conditioned on the set of remaining regions). Marrelec et al. (in press) provided a first study comparing this novel technique with SEM. However, more research is still required to clarify the link between both approaches. Furthermore, conditional correlation, like marginal correlation, also seems unable to provide relevant information regarding causality.

3.3. Data-driven measures of causality

In fact, exploratory investigation of causal relationships has proven at least as challenging as detection of direct interactions. Causality appears to be naturally embedded in SEM and DCM in the form of arrow directions. Inverting the direction of an arrow inverts the flow of information and, hence, causality. In other words, the search of causality in structural models is tantamount to setting arrow directions. In a more general setting, i.e., unrestricted to structural models, the question to address becomes far less obvious, though. It could nonetheless be argued that determining whether a given region i has a causal influence on another region j could be performed by assessing whether changes in region i affect region j. Despite some related research in electrophysiology, electroencephalography (EEG), and magnetoencephalography (MEG) (Baccala and Sameshima, 2001; Kamiński et al., 2001; Chávez et al., 2003), very few methods have been proposed in fMRI, with the notable exception of Goebel et al. (2003) and Roebroeck et al. (2005), using Granger causality. Granger's concept of causality (Granger, 1969) has been proposed in the context of linear regression models and received a great deal of attention. We say that a first time course causes another one if incorporating past values of the first time course improves the prediction of the current value of the second one (Geweke, 1982). Thus, temporal precedence is used to identify the direction of causality from the data. In Goebel et al. (2003), Granger causality was formalized and tested using vector autoregressive (VAR) models that capture the joint temporal dynamics of two time series. Recently, using VAR models, Roebroeck et al. (2005) have introduced Granger causality maps to explore the directed network of regions that are causally related to a seed region. If this exploratory approach is confirmed to have the capacity to infer changes of information flow between brain regions, it could prove an interesting addition to existing models of effective connectivity. To this aim, the causal nature of Granger causality (in the sense of SEM or DCM) must be confirmed. One should furthermore keep in mind that causation and direct interaction are two different concepts. Using direct transfer function (DTF) as a measure of Granger causality, Kamiński et al. (2001) indeed demonstrated that a nonzero DTF value did not necessarily imply direct causal influence between two regions: the effect could as well be mediated by another region or group of regions. Furthermore, the very issue of whether fMRI signals actually do convey causal information remains open. For instance, the intrinsic regional variability of the hemodynamic response and other acquisition artifacts imply that the hemodynamic response latency should be used with great care as a measure of temporal precedence (Friston et al., 1998; Miezin et al., 2000).

4. Discussion and perspectives

In this article, we reviewed some of the major concepts and tools used to explore functional networks according to the following two questions: (1) what brain regions are involved in the functional process under investigation? and (2) how do these regions interact? This artificial dichotomy emphasizes the practical division between procedures that are able to exhibit large-scale networks from measures quantifying functional interactions. However, both questions appear to be tightly linked and should, hence, not be regarded as independent. Many methods designed to identify large-scale networks in the brain indeed make use of a measure of interaction, e.g., functional connectivity (Biswal et al., 1995) or Granger causality (Roebroeck et al., 2005). On the other hand, to correctly measure direct interactions and causality within a network, it is mandatory to first include all regions potentially engaged in the network and able to mediate interactions. From a practical perspective, though, it is possible, and even common practice, to first use data-driven measures of connectivity before refining the analysis by application of model-based methods such as SEM (McIntosh et al., 1998; Huettel et al., 2004).

As many measures of connectivity involve at some point the correlation of fMRI time series, a major confounding

factor is spatially structured noise (Cordes et al., 2002). Major factors contributing to this correlation include partial volume effect, nonwhite measurement noise, pre-processing steps (such as slice-timing correction or spatial filtering), as well as physiology-induced fluctuations, such as respiratory and cardiac fluctuations (Dagli et al., 1999; Cordes et al., 2001; Windischberger et al., 2002) and motion-related artifacts (Cordes et al., 2002; Gavrilescu et al., 2004); see Woolrich et al. (2004) for fuller description of these factors. For a decade or so, the fact that measures such as functional connectivity could be mere reflections of structured noise could not be ruled out (Lund, 2001). Studies using dedicated fMRI acquisitions with short TR, allowing to correctly suppress physiological fluctuations, have showed that such factors do not suffice to explain the observed functional connectivity (Lowe et al., 1998; Cordes et al., 2001). Moreover, the spatial structure of functional connectivity maps are relevant in the face of current anatomical and functional knowledge of the brain architecture (De Luca et al., 2006). Together, these results are compelling evidence that functional connectivity is, at least partly, related to neuronal activity. Still, measures of connectivity can possibly be affected by fMRI noise. As such, it is important to develop and apply strategies to remove, or at least reduce, such spurious effects for routine fMRI acquisitions (e.g., Thomas et al., 2002; Perlbarg et al., 2007) in this context.

A key tool to assess the validity of large-scale network exploration in fMRI is knowledge of the underlying anatomical connections. Thanks to recent progress in MR imaging, anatomical connectivities can now be inferred from the anisotropic diffusion of water in white matter as measured by diffusion tensor imaging (DTI) (Basser et al., 2002). Using DTI, fiber pathways joining two regions can be tracked, producing results mostly in accordance with general anatomical knowledge (Wakana et al., 2004). Combination of DTI and fMRI will prove essential to discover to what extent the brain functional organization as investigated with fMRI reflects structural features of the brain and, hence, to better assess the relevance of fMRI to examine the relationship between functional and anatomical large-scale networks. Koch et al. (2002) compared anatomical and functional connectivity in healthy subject and found no simple relationship between these connectivities. Regions that were directly linked by fiber tracts exhibited higher functional connectivity; yet, the converse did not seem to hold. This result is of little surprise considered the nonspecific feature of correlation. As emphasized in Section 3.2, functional connectivity is not able to differentiate between direct interactions (that are supported by anatomical connections) and indirect interactions (that are not). As stressed by the authors (Koch et al., 2002), the correlations between pairs of regions could be mediated by indirect anatomical connections and in this case the BOLD signal may be correlated, although no connecting fiber is present in the anatomical network. In spite of this limitation, the prospect of using DTI to specify the underlying structural anatomical model to inform functional connectivity analyses is a promising direction of research. As a recent example (Lehéricy et al., 2005) demonstrated that motor representations shift from the associative to the sensorimotor territories of the striato-pallidal complex during the explicit learning of motor sequences, suggesting that motor skills are stored in the sensorimotor territory of the basal ganglia. Using DTI, the authors found (Lehéricy et al., 2004) that the sensorimotor territories of both the caudate nucleus and the putamen were connected to the sensorimotor cortex as well as premotor areas. Using DTI in a more extensive manner would definitely help to better understand the structural basis of large-scale networks (Ramnani et al., 2004). For instance, information originating from tractography could be used to constrain structural models. However, two issues must be tackled to this aim: the relevance of DTI and related tractography algorithms with regard to structural anatomy, and whether effective connectivity estimates indeed have anatomical correlates.

Another cogent source of validation for methods introduced in this review is the use of synthetic data. On the one hand, generative models rely on the theory of dynamical systems to simulate neuronal dynamics (Friston and Price, 2001a,b; David and Friston, 2003). They are strongly related to dynamical causal models and are hence already implicitely involved in effective connectivity investigation with DCM (Friston, 2003). Large-scale neural networks, on the other hand, have been introduced as a way to propose biologically plausible models of brain processing that can, to some extent, mimic its behavior at various temporal and spatial scales, depending on the model (Horwitz, 2004). Models for visual (Tagamets and Horwitz, 1998; Deco et al., 2004) and auditory (Husain et al., 2004) processing have been proposed so far. Such models were used to assess the relevance of both functional and effective connectivity (Horwitz et al., 1999, 2005; Lee et al., 2006). The great advantage of simulated data is that their complexity can be controlled. In synthetic neuronal models, and unlike what typically occurs in real brain acquisitions, various confounds (drift; habituation; movement; susceptibility, cardiac, and respiratory artifacts) can be either discarded as being irrelevant or specifically modeled to test the limits of the method under evaluation. Most importantly, the structure underlying data generation can be fully specified; what is expected can then be compared with what the method is able to retrieve from the simulated data. Crucially, the outcome of such evaluation will be all the more relevant for the investigation of brain interactions that the models used to produce the data are realistic.

A strongly connected issue is the search for a better interpretation of BOLD signal, i.e., the relationships between neuronal firing rates and neuronal activity with hemodynamic and metabolism. Despite extensive developments in brain functional imaging techniques, the physiological and

biochemical mechanisms involved in neural activity remain difficult to quantify. Several models have been developed to examine the relationships between synaptic activity (both excitatory and inhibitory), brain metabolic changes, vascular responses, and hemodynamic signals in cortical regions allowing one to relate results obtained at a systems-level with those obtained at the neural ensemble level (Aertsen et al., 1994; Arbib et al., 1995; Horwitz et al., 2000). Aertsen et al. (1994) used detailed simulations of interacting neuronal populations with Hodgkin Huxley-like dynamics to explore the relationship between synchronization and mean synaptic activity. Neural mass models afford a straightforward approach to model the activity of populations of neurons. These models are used to understand some macroscopic properties of MEG/EEG signals (Wendling et al., 2000; David and Friston, 2003). Several mathematical models link physiological processes and brain functional imaging data including oxygen exchanges between blood vessels and brain tissue (Buxton and Frank, 1997; Gjedde, 1997; Mintun et al., 2001), energy metabolism (Gruetter et al., 2001) and hemodynamic processes (Buxton and Frank, 1997; Friston et al., 2000; Aubert and Costalat, 2002). We anticipate that integration of the underlying processes will be a key step to better understand information flow.

Other methods than fMRI are used extensively as well to examine brain processes, including MEG and EEG. The signals measured by fMRI, MEG, or EEG stem from different features of brain activity, even though the exact underlying process is still under investigation (Li et al., 2000; Logothetis et al., 2001; Aubert and Costalat, 2002). Compared to BOLD fMRI, the signals obtained in MEG/EEG are more closely related to the neuronal currents (Hämäläinen et al., 1993). While fMRI signal would increase with neuronal firing rates, MEG/EEG is rather sensitive to post-synaptic activity (Nunez and Silberstein, 2000). Nonetheless, MEG/EEG have also proved useful to investigate brain networks. The (a)synchrony between two spatially remote neuronal spike trains can be interpreted with respect to the so-called ''binding problem'' (Treisman, 1996), still in debate, as the necessary brain mechanisms providing a coherent sensorial/emotional experience integration (Varela et al., 2001). Different tools have been proposed in MEG/EEG for connectivity analysis including temporal correlation, nonlinear correlation, mutual information, generalized synchronization, phase difference, coherence, and phase locking value. Despite very distinctive definitions of all these indices of functional interactions, they essentially measure the same quantity, namely correlation, under the premise that we are seeking linear relationships between i.i.d. multivariate Gaussian variables (Marrelec et al., 2005a). Such a finding provides a first step toward a common conceptual framework for joint fMRI and MEG/EEG functional brain connectivity exploration. This is all the more important that an increasing number of protocols include combined or even simultaneous fMRI and MEG/EEG acquisitions.

Acknowledgements

The authors wish to thank an anonymous referee for improving the quality of this manuscript and Barry Horwitz for providing us with an historical perspective on functional connectivity in PET. They are also thankful to Mélanie Pélégrini-Issac for proofreading the last version of this manuscript. G. Marrelec is supported partly by the Foundation Fyssen (Paris, France) and partly by the Bourse des Gouverneurs (Montreal, Canada).

References

Aertsen, A., Erb, M., Palm, G., 1994. Dynamics of functional coupling in the cerebral cortex: an attempt at a model-based interpretation. Physica D 75, 103–128.

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

Arbib, M.A., Bischoff, A., Fagg, A.H., Grafton, S.T., 1995. Synthetic PET: analyzing large-scale properties of neural networks. Human Brain Mapping 2, 225–233.

Attwell, D., Iadecola, C., 2002. The neural basis of functional brain imaging signals. Trends in Neurosciences 25, 621–625.

Aubert, A., Costalat, R., 2002. A model of the coupling between brain electrical activity, metabolism, and hemodynamics: application to the interpretation of functional neuroimaging. NeuroImage 17, 1162–1181.

Baccala, L.A., Sameshima, K., 2001. Partial directed coherence: a new concept in neural structure determination. Biological Cybernetics 84, 463–474.

Basser, P.J., Mattiello, J., Le Bihan, D., 2002. Estimation of the effective self-diffusion tensor from the NMR spin echo. Journal of Magnetic Resonance 103, 247–254.

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

Baumgartner, R., Somorjai, R., Summers, R., Richter, W., Ryner, L., 2000. Correlator beware: correlation has limited selectivity for fMRI data analysis. NeuroImage 12, 240–243.

Bellec, P., Perlbarg, V., Jbabdi, S., Pélégrini-Issac, M., Anton, J.-L., Doyon, J., Benali, H., 2006. Identification of large-scale networks in the brain using fMRI. NeuroImage 29, 1231–1243.

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. Magnetic Resonance in Medicine 34, 537–541.

Büchel, C., Friston, K.J., 1997. Modulation of connectivity in visual pathways by attention: cortical interactions evaluated with structural equation modelling and fMRI. Cerebral Cortex 7, 768–778.

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.

Buxton, R.B., Frank, L.R., 1997. A model for the coupling between cerebral blood flow and oxygen metabolism during neural stimulation. Journal of Cerebral Blood Flow and Metabolism 17, 64–72.

Chávez, M., Martinerie, J., Le Van Quyen, M., 2003. Statistical assessment of nonlinear causality: application to epileptic EEG signals. Journal of Neuroscience 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.

Clark, C.M., Kessler, R., Buchsbaum, M.S., Margolin, R.A., Holcomb, H.H., 1984. Correlational methods for determining regional coupling of cerebral glucose metabolism: a pilot study. Biological Psychiatry 19, 663–678.

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. American Journal of Neuroradiology 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. American Journal of Neuroradiology 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. Magnetic Resonance Imaging 20, 305–317.

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

Cudeck, R., 1989. Analysis of correlation matrices using covariance structure models. Psychological Bulletin 105, 317–327.

Cudeck, R., Klebe, K.J., Henly, S.J., 1993. A simple Gauss–Newton procedure for covariance structure analysis with high-level computer languages. Psychometrika 58, 211–232.

Dagli, M.S., Ingeholm, J.E., Haxby, J.V., 1999. Localization of cardiac-induced signal change in fMRI. Physics in Medicine and Biology 9, 407–415.

David, O., Friston, K.J., 2003. A neural mass model for MEG/EEG: coupling and neuronal dynamics. NeuroImage 20, 1743–1755.

Deco, G., Rolls, E.T., Horwitz, B., 2004. Integrating fMRI and single-cell data of visual working memory. Neurocomputing 58-60, 729–737.

De Luca, M., Beckmann, C.F., De Stefano, N., Matthews, P.M., Smith, S.M., 2006. fMRI resting state networks define distinct modes of long-distance interactions in the human brain. NeuroImage 29, 1359–1367.

Ethofer, T., Anders, S., Erb, M., Herbert, C., Wiethoff, S., Kissler, J., Grodd, W., Wildgruber, D., 2006. Cerebral pathways in processing of affective prosody: a dynamical causal modeling study. NeuroImage 30, 580–587.

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, second ed. Academic Press.

Friston, K.J., 2003. Learning and inference in the brain. Nature Neuroscience 16, 1325–1352.

Friston, K.J., Price, C.J., 2001a. Dynamic representations and generative models of brain function. Brain Research Bulletin 54, 275–285.

Friston, K.J., Price, C.J., 2001b. Generative models, brain function and neuroimaging. Scandinavian Journal of Psychology 42, 167–177.

Friston, K.J., Frith, C.D., Frackowiak, R.S.J., 1993a. Time-dependent changes in effective connectivity measured with PET. Human Brain Mapping 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. Journal of Cerebral Blood Flow and Metabolism 13, 5–14.

Friston, K.J., Holmes, A.P., Poline, J.-B., Grasby, P.J., Williams, S.C.R., Frackowiak, R.S.J., Turner, R., 1995. Analysis of fMRI time-series revisited. NeuroImage 2, 45–53.

Friston, K.J., Frith, C.D., Fletchet, P., Liddle, P.F., Frackowiak, R.S.J., 1996. Functional topography: multidimensional scaling and functional connectivity in the brain. Cerebral Cortex 6, 156–164.

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., Fletcher, P., Josephs, O., Holmes, A., Rugg, M.D., Turner, R., 1998. Event-related fMRI: characterizing differential responses. NeuroImage 7, 30–40.

Friston, K.J., Michelli, A., Turner, R., Price, C.J., 2000. Nonlinear responses in fMRI: the balloon model, Voltera kernels, and other hemodynamics. NeuroImage 12, 466–477.

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. Human Brain Mapping 21, 49–63.

Geweke, J.F., 1982. Measurement of linear dependence and feedback between multiple time series. Journal of American Statistical Association 77, 304–324.

Gjedde, A., 1997. The relation between brain function and cerebral blood flow and metabolism. In: Batjer, H.H. (Ed.), Cerebrovascular Disease. Lippincott-Raven, Philadelphia, pp. 23–40.

Goebel, R., Roebroeck, A., Kim, D.-S., Formisano, E., 2003. Investigating directed cortical interactions in time-resolved fMRI data using vector autoregressive modeling and Granger causality mapping. Magnetic Resonance Imaging 21, 1251–1261.

Gonçalves, M.S., Hall, D.A., 2003. Connectivity analysis with structural equation modelling: an example of the effects of voxel selection. NeuroImage 20, 1455–1467.

Gonzalez-Lima, F., McIntosh, A.R., 1995. Analysis of neural interactions related to associative learning using structural equation modeling. Mathematics and Computers in Simulation 40, 115–140.

Goutte, C., Toft, P., Rostrup, E., Nielsen, F.Å., Hansen, L.K., 1999. On clustering fMRI time series. NeuroImage 9, 298–310.

Goutte, C., Hansen, L.K., Liptrot, M.G., Rostrup, E., 2001. Feature-space clustering for fMRI meta-analysis. Human Brain Mapping 13, 165–183.

Granger, C.W.J., 1969. Investigating causal relations by econometric models and cross-spectral methods. Econometrica 37, 424–438.

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 USA 100, 253–258.

Greicius, M.D., Srivastava, G., Reiss, A.L., Menon, V., 2004. Default-mode network activity distinguishes Alzheimer's disease from healthy aging: evidence from functional MRI. Proceedings of the National Academy of Sciences of the USA 101, 4637–4642.

Gruetter, R., Seaquist, E.R., Ugurbil, K., 2001. A mathematical model of compartmentalized neurotransmitter metabolism in the human brain. American Journal of Physiology, Endocrinology and Metabolism 281, E100–E112.

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

Hämäläinen, M., Hari, R., Ilmoniemi, R., Knuutila, J., Lounasmaa, O.V., 1993. Magnetoencephalography – theory, instrumentation, and application to noninvasive studies of the working human brain. Reviews of Modern Physics 65, 413–497.

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.

Horwitz, B., 2004. Relating fMRI and PET signals to neural activity by means of large-scale neural models. Neuroinformatics 2, 251–266.

Horwitz, B., 2005. Integrating neuroscientific data across spatiotemporal scales. Comptes Rendus Biologies 328, 109–118.

Horwitz, B., Duara, R., Rapoport, S.I., 1984. Intercorrelations of glucose metabolic rates between brain regions: application to healthy males in a state of reduced sensory input. Journal of Cerebral Blood Flow and Metabolism 4, 484–499.

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

Horwitz, B., Friston, K.J., Taylor, J.G., 2000. Neural modeling and functional brain imaging: an overview. Neural Networks 13, 829–846.

Horwitz, B., Warner, B., Fitzer, J., Tagamets, M.-A., Husain, F.T., Long, T.W., 2005. Investigating the neural basis for functional and effective connectivity. Application to fMRI. Philosophical Transactions of the Royal Society of London 360, 1093–1108.

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

Husain, F.T., Tagamets, M.-A., Fromm, S.J., Braun, A.R., Horwitz, B., 2004. Relating neuronal dynamics for auditory object processing to

neuroimaging activity: a computational modeling and an fMRI study. NeuroImage 21, 1701–1720.

Kamiński, M., Ding, M., Truccolo, W.A., Bressler, S.L., 2001. Evaluating causal relations in neural systems: Granger causality, directed transfer function and statistical assessment of significance. Biological Cybernetics 85, 145–157.

Koch, M.A., Norris, D.G., Hund-Georgiadis, M., 2002. An investigation of functional and anatomical connectivity using magnetic resonance imaging. NeuroImage 16, 241–250.

Laureys, S., Owen, A.M., Schiff, N.D., 2004. Brain function in coma, vegetative state, and related disorders. The Lancet Neurology 3, 537–546.

Lee, L., Harrison, L.M., Mechelli, A., 2003a. The functional brain connectivity workshop: report and commentary. Network: Computation in Neural Systems 14, R1–R15.

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

Lee, L., Friston, K., Horwitz, B., 2006. Large-scale neural models and dynamic causal modelling. NeuroImage 30, 1243–1254.

Lehéricy, S., Ducros, M., Van de Moortele, P.-F., Francois, C., Thivard, L., Poupon, C., Swindale, N., Ugurbil, K., Kim, D.-S., 2004. Diffusion tensor fiber tracking shows distinct corticostriatal circuits in humans. Annals of Neurology 55, 522–529.

Lehéricy, S., Benali, H., Van de Moortele, P.-F., Pélégrini-Issac, M., Waechter, T., Ugurbil, K., Doyon, J., 2005. Distinct basal ganglia territories are engaged in early and advanced motor sequence learning. Proceedings of the National Academy of Sciences of the USA 102, 12566–12571.

Li, T.-Q., Haefelin, T.N., Chan, B., Kastrup, A., Jonsson, T., Glover, G.H., Moseley, M.E., 2000. Assessment of hemodynamic response during focal neural activity in human using bolus tracking, arterial spin labeling and BOLD techniques. NeuroImage 12, 442–451.

Logothetis, N.K., Pauls, J., Augath, M., Trinath, T., Oeltermann, A., 2001. Neurophysiological investigation of the basis of the fMRI signal. Nature 412, 150–157.

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.

Lund, T.E., 2001. fcMRI – mapping functional connectivity or correlating cardiac-induced noise? Magnetic Resonance in Medicine 46, 628.

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 Transactions on Signal Processing 53, 3503–3516.

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

Marrelec, G., Horwitz, B., Kim, J., Pélégrini-Issac, M., Benali, H., Doyon, J., in press. Using partial correlation to enhance structural equation modeling of functional MRI data. NeuroImage.

Marrelec, G., Krainik, A., Duffau, H., Pélégrini-Issac, M., Lehéricy, S., Doyon, J., Benali, H., 2006. Partial correlation for functional brain interactivity investigation in functional MRI. NeuroImage 32, 228–237.

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

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.

McIntosh, A.R., Cabeza, A.R., Lobaugh, N.J., 1998. Analysis of neural interactions explains the activation of the occipital cortex by an auditory stimulus. Journal of Neurophysiology 80, 2790–2796.

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. Human Brain Mapping 6, 160–188.

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.

Mechelli, A., Henson, R.N.A., Price, C.J., Friston, K.J., 2003a. Comparing event-related and epoch analysis in blocked design fMRI. NeuroImage 18, 806–810.

Mechelli, A., Price, C.J., Noppeney, U., Friston, K.J., 2003b. A dynamic causal modeling study on category effects: bottom-up or top-down mediation? Journal of Cognitive Neuroscience 15, 925–934.

Metter, E.J., Riege, W.H., Kuhl, D.E., Phelps, M.E., 1984. Cerebral metabolic relationships for selected brain regions in healthy adults. Journal of Cerebral Blood Flow and Metabolism 4, 1–7.

Miezin, F.M., Maccotta, L., Ollinger, J.M., Petersen, S.E., Buckner, R.L., 2000. Characterizing the hemodynamic response: effects of presentation rate, sampling procedure, and the possibility of ordering brain activity based on relative timing. NeuroImage 11, 735–759.

Mintun, M.A., Lundstrom, B.N., Snyder, A.Z., Vlassenko, A.G., Shulman, G.L., Raichle, M.E., 2001. Blood flow and oxygen delivery to human brain during functional activity: theoretical modeling and experimental data. Proceedings of the National Academy of Sciences of the USA 98, 6859–6864.

Nunez, P.L., Silberstein, R.B., 2000. On the relationship of synaptic activity to macroscopic measurements: does co-registration of EEG with fMRI make sense? Brain Topography 13, 79–96.

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.

Perlbarg, V., Bellec, P., Anton, J.-L., Pélégrini-Issac, M., Doyon, J., Benali, H., 2007. CORSICA: correction of structured noise in fMRI by automatic identification of ICA components. Magnetic Resonance Imaging 25, 35–46.

Ramnani, N., Behrens, T.E.J., Penny, W., Matthews, P.M., 2004. New approaches for exploring anatomical and functional connectivity in the human brain. Biological Psychiatry 56, 613–619.

Roebroeck, A., Formisano, E., Goebel, R., 2005. Mapping directed influence over the brain using Granger causality and fMRI. NeuroImage 25, 230–242.

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

Shulman, R.G., Rothman, D.L., Behar, K.L., Hyder, F., 2004. Energetic basis of brain activity: implication for neuroimaging. Trends in Neurosciences 27, 489–495.

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

Stein, T., Moritz, C., Quigley, M., Cordes, D., Haughton, V., Meyerand, E., 2000. Functional connectivity in the thalamus and hippocampus studied with functional MR imaging. American Journal of Neuroradiology 21, 1397–1401.

Stone, J.V., Kötter, R., 2002. Making connections about brain connectivity. Trends in Cognitive Sciences 6, 327–328.

Tagamets, M.-A., Horwitz, B., 1998. Integrating electrophysiological and anatomical experimental data to create a large-scale network that simulates a delayed match-to-sample human brain imaging study. Cerebral Cortex 8, 310–320.

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., Sporns, O., Edelman, G.M., 1992. Reentry and the problem of integrating multiple cortical areas: simulation of dynamics integration in the visual system. Cerebral Cortex 2, 310–335.

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

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

Treisman, A., 1996. The binding problem. Current Opinion in Neurobiology 6, 171–178.

Varela, F.J., Lachaux, J.-P., Rodriguez, E., Martinerie, J., 2001. The brainweb: phase synchronization and large-scale intergration. Nature Reviews Neuroscience 2, 229–239.

Waites, A.B., Stanislavsky, A., Abbott, D.F., Jackson, G.D., 2005. Effect of prior cognitive state on resting state networks measured with functional connectivity. Human Brain Mapping 24, 59–68.

Wakana, S., Jiang, H., Nagae-Poescher, L.M., van Zijl, P.C.M., Mori, S., 2004. Fiber tract-based atlas of human white matter anatomy. Radiology 230, 77–87.

Welchew, D.E., Honey, G.D., Sharma, T., Robbins, T.W., Bullmore, E.T., 2002. Multidimensional scaling of integrated neurocognitive function and schizophrenia as a disconnexion disorder. NeuroImage 17, 1227–1239.

Wendling, F., Bellanger, J.J., Bartolomei, F., Chauvel, P., 2000. Relevance of nonlinear lumped-parameter models in the analysis of depth-EEG epileptic signals. Biological Cybernetics 83, 367–378.

Windischberger, C., Langenberger, H., Sycha, T., Tschernko, E.M., Fuchsjäger-Mayerl, G., Schmetterer, L., Moser, E., 2002. On the origin of respiratory artifacts in BOLD-EPI of the human brain. Magnetic Resonance Imaging 20, 575–582.

Woolrich, M.W., Behrens, T.E.J., Beckmann, C.F., Jenkinson, M., Smith, S.M., 2004. Multilevel linear modelling for fMRI group analysis using Baysian inference. NeuroImage 21, 1732–1747.

Xiong, J., Parsons, L.M., Gao, J.-H., Fox, P.T., 1999. Interregional connectivity to primary motor cortex revealed using MRI resting state images. Human Brain Mapping 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.