ACADEMIC PRESS

Available online at www.sciencedirect.com
SCIENCE DIRECT®

NeuroImage

NeuroImage 20 (2003) 962–974

www.elsevier.com/locate/ynimg

Functional connectivity: studying nonlinear, delayed interactions between BOLD signals

Pierre-Jean Lahaye,[a,b] Jean-Baptiste Poline,[a,b,*] Guillaume Flandin,[a,b,c] Silke Dodel,[a,b] and Line Garnero[a,d]

[a] *IFR 49, Imagerie NeuroFonctionnelle, France*
[b] *CEA/SHFJ, 4, place du genéral Leclerc, 91401 Orsay, France*
[c] *INRIA Sophia Antipolis, Epidaure Research Project, 2004, Route des lucioles, 06902 Sophia Antipolis, France*
[d] *LENA, Neurosciences cognitives et Imagerie cérébrale, Hopital de la Salpetriére, 47, Bd de l'hopital, 75661 Paris, France*

Received 10 January 2003; revised 16 May 2003; accepted 29 May 2003

Abstract

Correlation analysis has been widely used in the study of functional connectivity based on fMRI data. It assumes that the relevant information about the interactions of brain regions is reflected by a linear relationship between the values of two signals at the same time. However, this hypothesis has not been thoroughly investigated yet. In this work, we study in depth the information shared by BOLD signals of pairs of brain regions. In particular, we assess the amount of nonlinear and/or nonsynchronous interactions present in data. This is achieved by testing models reflecting linear, synchronous interactions against more general models, encompassing nonlinear, nonsynchronous interactions. Many factors influencing measured BOLD signals are critical for the study of connectivity, such as paradigm-induced BOLD responses, preprocessing, motion artifacts, and geometrical distortions. Interactions are also influenced by the proximity of brain regions. The influence of all these factors is taken into account and the nature of the interactions is studied using various experimental conditions such that the conclusions reached are robust with respect to variation of these factors. After defining nonlinear and/or nonsynchronous interaction models in the framework of general linear models, statistical tests are performed on different fMRI data sets to infer the nature of the interactions. Finally, a new connectivity metric is proposed which takes these inferences into account. We find that BOLD signal interactions are statistically more significant when taking into account the history of the distant signal, i.e., the signal from the interacting region, than when using a model of linear instantaneous interaction. Moreover, about 75% of the interactions are symmetric, as assessed with the proposed connectivity metric. The history-dependent part of the coupling between brain regions can explain a high percentage of the variance in the data sets studied. As these results are robust with respect to various confounding factors, this work suggests that models used to study the functional connectivity between brain areas should in general take the BOLD signal history into account.

Keywords: fMRI; Functional connectivity; Connectivity measures

1. Introduction

In the past decade, fMRI has provided a powerful approach to study in vivo the structure–function relationship in the human brain. While most of the work concentrated on detecting or estimating brain regions involved in specific cognitive or sensorimotor tasks, there is an increasing interest in understanding the relationships between brain regions depending, for example, on the state induced by an experimental paradigm.

To summarize, two basic kinds of relationships between brain regions have been investigated, the so-called functional and effective connectivities. In the former, no underlying anatomical model is assumed for the connections. In the latter, an a priori oriented graph is defined using prior knowledge from neuroscience, thus focusing on links among predefined anatomical regions.

Most techniques used to assess functional connectivity [zero-delay correlation coefficients (Biswal et al., 1995)

* Corresponding author. Fax: +33-1-69-86-78-68.
E-mail address: poline@shfj.cea.fr (J.-B. Poline).

1053-8119/03/$ – see front matter © 2003 Elsevier Inc. All rights reserved.
doi:10.1016/S1053-8119(03)00340-9

multivariate approaches (Della-Maggiore et al., 2000; McIntosh et al., 1996)] and effective connectivity [structural equation modeling (McIntosh and Gonzalez-Lima, 1994; Büchel and Friston, 1997)] are based on the correlation or the covariance matrix of the data.

Applications of these techniques in neuroscience are based on the observation of changes in connectivity when varying conditions or subject populations. They have proven useful to detect modulation of connectivity by attention (Friston and Beuchel, 2000) brain plasticity while learning (Poldrack, 2000) or alteration of connectivity between different populations (e.g., healthy subjects vs patients). The study of connectivity during the so-called resting state is an active field of research as well (see for instance (Lowe et al., 1998, 2000; Goldman et al., 2002) among others).

In functional or effective connectivity, the measure used as an indicator of the shared information among two regions is most often simply related to the zero-delay correlation of their functional signals.

Few studies using more complex indicators of connectivity have been proposed so far (see (Tononi et al., 1998) for an example). The first technique introducing a temporal memory in neuroimaging for the analysis of functional connectivity uses Volterra autoregressive models (see Appendix B). In this respect one can cite Harrison et al. (Harrison et al., 2002) and others.

To our knowledge, there is still a lack of understanding about the crucial information shared by fMRI time series originating from interacting brain regions—that is, the amount of information that one signal contains about another—and about the nature of their relation. The use of correlation limits a priori the shared information to linear and instantaneous interactions.

In this work, we address the question of whether the assumption of linear instantaneous interactions is sufficient to estimate functional connectivity or if nonlinear terms or lagged terms should be incorporated when evaluating interactions. More specifically, we study the amount of nonlinear and delayed interactions. The information explained by a type of interaction is evaluated as the variance explained by corresponding interaction terms. We also study whether these additional components vary with the strength of the connectivity as measured with correlation.

If the linear, instantaneous interactions explains most of signal interaction, this work would validate methods already proposed so far. If not, it may be that a fraction of the information shared by functionally or effectively connected regions is mediated by nonlinear and delayed interactions.

To ensure the robustness of our findings, we use several data sets (based on different paradigms, subjects, and scanners) and various methods of data preprocessing. Numerous factors may indeed influence the evaluation of connectivity measures, such as movements of the subject (if not appropriately corrected, movement can induce purely artifactual correlation) and the choice of low- and high-pass temporal filters (the lower frequency band can be contaminated by aliased respiratory and cardiac effects). Other factors such as the choice of the model used to detect the paradigm-induced signal could also affect the measured connectivity. There seems to be little known about the influence of those factors.

The results are given in terms of statistical significance and percentage of variance explained by models including nonlinear or noninstantaneous components. Special attention is granted to the robustness of the results depending on the various preprocessing steps.

The reasons for why interactions could be nonlinear and influenced by signal histories may be complex. Here we do not have the means to tease apart these two components, but as a first step, we quantify the amount of noninstantaneous interactions among brain regions. In the last section of the paper, possible causes of dependencies of signal history are discussed in light of our results.

2. Methods

2.1. Summary

The data used in this work consist of small brain regions covering the gray matter in each subject. The representation of the data is based on the parcellation algorithm described by Flandin et al. (Flandin et al., 2000). In our approach, interactions between all these regions are studied.

Models of different kinds of interactions are defined. To conclude whether a specific interaction between two given brain regions should take into account more information than simple correlation, complex models of interaction is tested against simpler models. The results are then extended to the whole set of interactions using a second level of testing.

2.2. Parcellation

Motivation

The high spatial resolution of fMRI images makes it difficult to study the connectivity among all pairs of voxels. Moreover, spatial correlation between neighboring voxels results in temporal correlation of the corresponding signals (Kiebel et al., 1999) which affects the resulting connectivity.

To solve these two issues, the selection of seed voxels or the definition of regions of interest is frequently used in connectivity studies. While the BOLD signal at seed voxels may be heavily affected by noise, the manual extraction of regions of interest can introduce bias.

In this paper, an alternative representation of fMRI data is used instead, as proposed by Flandin et al. (Flandin et al., 2000). It consists of an automatic parcellation of gray matter into a predefined number of regions. Fig. 1 shows an example of such a parcellation.

Fig. 1. Axial slice of parcelled grey matter, for one hemisphere.

Parcellation method

The parcellation scheme operates on a segmentation of gray matter detected on the basis of a histogram analysis (after image bias correction), as described by Mangin et al. (Mangin et al., 1998) An arbitrary number of seed voxels is dispatched so that they tend to be regularly spaced. Finally, gray matter is divided into parcels according to a voronoi diagram based on the seeds. Seeds and parcels are shown in Fig. 1. No functional constraint has been used to define parcellations.

Attributing functional signal to parcels

Functional images are coregistered with the anatomical image using SPM99 (mutual information) and then interpolated to its higher resolution with a sinc function. The strategy used to impute each parcel an equivalent signal uses two settings.

The first one uses gray matter parcel averaging, i.e., the equivalent signal is defined as the average of BOLD signals at each anatomical voxel belonging to the parcel. The second simply uses the BOLD signal at the seed voxel of each parcel. The latter setting is used to control so that there is no important loss of information in the averaging process.[1]

Properties of the parcellation

Parcellation can be seen as an informed spatial filter avoiding averaging between voxels belonging to different tissues. As discussed in the original work, the parcellation provides a good summary of functional data. The parcellation as a whole enables the observation of all putative gray matter–gray matter interactions at an arbitrary spatial resolution. By construction, the parcellation converges toward equal volume parcels. In our study, 100 parcels have been defined for each subject considered. The parcels have a mean volume of 4700 anatomical voxels, which is equivalent to 6.3 cm^3.

Choosing this sufficiently large mean volume should make the following study more robust to noise and effects of motion. Spatial correlation between parcels should also be reduced, independent of its origins (such as acquisition technique, vascular coupling, or similarity of neural response due to local homogeneity in function).

[1] These two settings later correspond to Settings 1 and 4 in Table 2.

2.3. Models

In order to investigate the nature of the parcel-to-parcel BOLD signal interactions, models of possible forms of interaction are designed. Considering the interaction $b \rightarrow a$, the BOLD signal of a is seen as a sum of terms, in particular the multiple connectivity terms due to the activity of b.

Some notations

We denote by a and b, respectively, two parcels as defined above, and by $x_a(t)$, $x_b(t)$, the corresponding BOLD signals at time t. In the following, an oriented relation $b \rightarrow a$ is called an interaction. $\mathcal{H}^i_{b\rightarrow a}$, $\mathcal{M}^i_{b\rightarrow a}$, and $\mathcal{G}^i_{b\rightarrow a}$ are, respectively, the hypothesis, the model, and the regressors of the model for the interaction $b \rightarrow a$.

Hypotheses and models for interactions

We now consider different hypotheses about the interaction $b \rightarrow a$, which will be tested one against the other:

$\mathcal{H}^0_{b\rightarrow a}$	Linear, instantaneous influence of x_b on x_a
$\mathcal{H}^{\mathrm{NL}}_{b\rightarrow a}$	Nonlinear, instantaneous influence of x_b on x_a
$\mathcal{H}^{\mathrm{HD}}_{b\rightarrow a}$	Linear influence of the past and present values of $x_b, x_b(s \leq t)$ on $x_a(t)$
$\mathcal{H}^{\mathrm{NL}\times\mathrm{HD}}_{b\rightarrow a}$	Nonlinear influence of the past and present values of $x_b, x_b(s \leq t)$ on $x_a(t)$

In the following paragraphs, we show how these terms can be assessed in the framework of general linear models (GLM) through the specification of some sets of regressors. These regressors are sets of vectors depending on time and can be seen as matrices.

Linear instantaneous interaction term. This term corresponds to $\{x_b(t)\}$. It accounts for a linear instantaneous part of the parcel-to-parcel relationship between a and b. In our models, it is considered a confound, since our goal is to test for more complex parcel-to-parcel relationships not explained by zero-delay correlation. All specified models therefore include this term.

Recent history term. As shown by Woolrich et al. (Woolrich et al., 2001) autocorrelation of BOLD signals is high, which implies that x_a is partly explained by its own history. The recent history term of a, approximated with $\{x_a(t - l\tau)\}_{1\leq l\leq n_{\mathrm{HD}}}$, where τ is the interscan time, allows one to model this effect.

If there is some instantaneous correlation between a and b, the history of a may then explain part of b. To prevent the tests from detecting such spurious relations, the history term is considered a confound when modeling $x_a(t)$ and is therefore included in all models.

Table 1
Summary of the specification of the various models used to test the hypothesis described in Section 2.3: $x_a = \mathcal{G}^i_{b\to a} \cdot \beta^i_{b\to a}$

		Instantaneous interaction	x_a recent history	Nonlinear and lagged interaction terms
$\mathcal{M}^0_{b\to a}$:	$\mathcal{G}^0_{b\to a}$	$= \{x_b(t),$	$x_a(t-k\tau)_{1\le k\le n_{HD}}\}$	
$\mathcal{M}^{NL}_{b\to a}$:	$\mathcal{G}^{NL}_{b\to a}$	$= \{x_b(t),$	$x_a(t-k\tau)_{1\le k\le n_{HD}},$	$x_b(t)^l_{1\le l\le n_{NL}}\}$
$\mathcal{M}^{HD}_{b\to a}$:	$\mathcal{G}^{HD}_{b\to a}$	$= \{x_b(t),$	$x_a(t-k\tau)_{1\le k\le n_{HD}},$	$x_b(t-k\tau)_{1\le k\le n_{HD}}\}$
$\mathcal{M}^{NL\times HD}_{b\to a}$:	$\mathcal{G}^{NL\times HD}_{b\to a}$	$= \{x_b(t),$	$x_a(t-k\tau)_{1\le k\le n_{HD}},$	$x_b(t-k\tau)^l_{1\le k\le n_{HD}}$ $_{0\le l\le n_{NL}}\}$
$\mathcal{M}^{NL+HD}_{b\to a}$:	$\mathcal{G}^{NL+HD}_{b\to a}$	$= \{x_b(t),$	$x_a(t-k\tau)_{1\le k\le n_{HD}},$	$x_b(t-k\tau)_{1\le k\le n_{HD}},$ $x_b(t)^l_{1\le l\le n_{NL}}\}$

Note. See Fig. 2 for a graphical representation of these models.

Linear lagged interaction term. Similarly, the "linear" influence of the history of signals x_b onto x_a is modeled using the past n_{HD} values. This corresponds to the terms $\{x_b(t-l\tau)\}_{1\le l\le n_{HD}}$, where τ is the interscan time.[2]

Nonlinear instantaneous interaction term. The nonlinear influence of $x_b(t)$ can be modeled using an expansion at order n_{NL}, with polynomial regressors $\{x_b(t)^k\}_{1\le k\le n_{NL}}$.

Nonlinear, lagged interaction term. Finally, the combination of the two possible effects, namely nonlinearity and dependence on signal history, can be modeled using the family of regressors $\{x_b(t-l\tau)^k\}^{1\le k\le n_{NL}}_{0\le l\le n_{HD}}$.

Using these approximations, the different models corresponding to the hypotheses above are defined, according to the terms they should include. The sets of regressors for each model are summarized in Table 1 and Fig. 2. The model $\mathcal{M}^{NL+HD}$ is designed to test the specific coupling between the historical and nonlinear components of a given interaction.

This set of models has a partly nested design, in the sense that $\mathcal{G}^0 \subset \mathcal{G}^{HD}$, $\mathcal{G}^{NL} \subset \mathcal{G}^{NL+HD} \subset \mathcal{G}^{NL\times HD}$.

The reader can refer to Appendix B in which we describe the difference between the proposed approach and the use of Volterra kernels.

Model fitting

Using the regressors defined in the preeceding section enables one to enter the GLM framework (Friston et al., 1995; Worsley and Friston, 1995). The regressors are specifically built for each specific *a* and *b* parcel, and the corresponding regression is performed. An approximated noise correlation matrix, $V = K'K$, is used, corresponding to

[2] This interaction term for interaction $b \to a$ (based on x_b) differs from the recent history term (based on x_a).

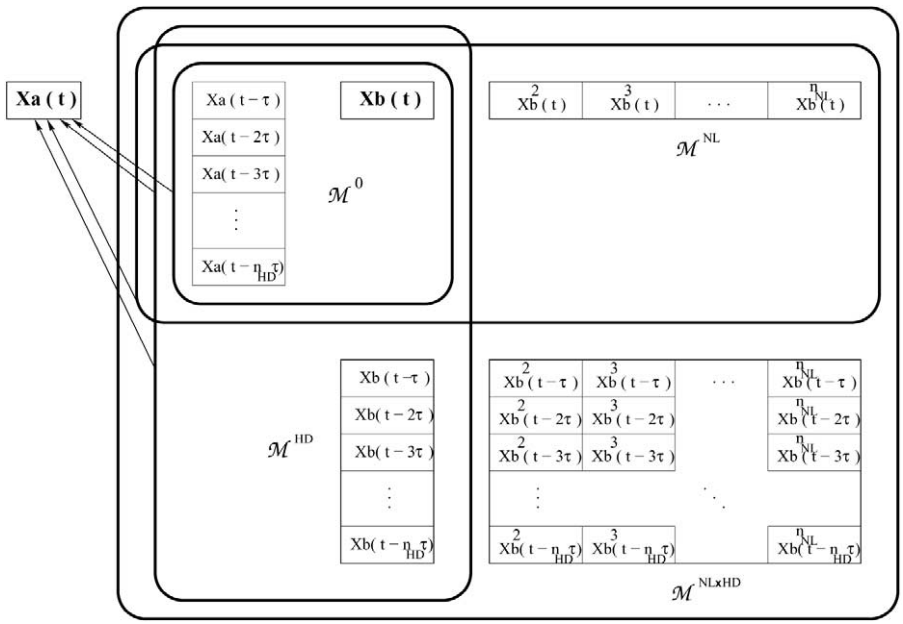

Fig. 2. The different models for interaction $x_a \to x_b$. Each model includes a linear interaction term and an auto-regressive term. Supplementary terms optionally account for nonlinearity and signal histories.

the correlation induced by convolution with a gaussian kernel K of FWHM = 4 s.

Fitting model $\mathcal{M}^i_{b\to a}$ reads as a matrix equation,

$$x_a = \mathcal{G}^i_{b\to a} \cdot \beta^i_{b\to a},$$

where $\mathcal{G}^i_{b\to a} \in \{\mathcal{G}^0_{b\to a}, \mathcal{G}^{\mathrm{HD}}_{b\to a}, \mathcal{G}^{\mathrm{NL}}_{b\to a}, \mathcal{G}^{\mathrm{NL+HD}}_{b\to a}, \mathcal{G}^{\mathrm{NL\times HD}}_{b\to a}\}$.

Model orders

Model orders n_{NL} and n_{HD} are chosen as follows.

Maximum lag of interactions. Given the typical length of the hemodynamic response, the span of history influencing BOLD signals should be within 10–20 s. The first minimum of signals' autocorrelation is classically used as time embedding order (see Appendix B). It was calculated for all signals from the rest data set and then averaged over signals. The mean value occurs for $\tau = 10$ s (that is $n_{\mathrm{HD}} = 5$ in the first data set described under Experiments), which is consistent with the typical orders of magnitude of the HRF time constants.

Maximum nonlinearity order. No clear biological prior on the order of nonlinearities is available. The model order $n_{\mathrm{NL}} = 5$ is chosen as a default setting for modeling nonlinearity, although orders 2 or 3 may be sufficient (see Results).

Validation. Different orders are used as control settings for validation (see Table 2). Their role is twofold. First, they enable one to check that hypothesis testing is robust to their change. Second, choosing n_{HD} corresponding to the maximum length of the hemodynamic delays (e.g., 20 s) enables one to check results while making sure that no local autoregressive effect remains, which could be artifactually attributed to other terms in the models.

2.4. Tests

First, the significance of each effect is assessed using the models specified previously for each pair of parcel using F tests (this amounts to $n(n-1)/2$ tests, with $n = 100$ in our setting). Second, the set of interactions is tested as a whole by comparing the distribution of the F test under the null hypothesis H^0 to the one obtained on the set of interactions using Kolmogorov–Smirnov tests. In this second step, the number of interactions ensures a good statistical power.

Hypothesis testing for a single interaction

In the following, we denote $F^{\mathcal{M}_i/\mathcal{M}_j}_{b\to a}$ the value of test F when testing model $\mathcal{M}^i_{b\to a}$ against model $\mathcal{M}^j_{b\to a}$ (that both model interaction $x_b \to x_a$).

The different hypotheses defined are tested by comparison of the corresponding model against the corresponding reduced model (M^0 or $M^{\mathrm{NL+HD}}$, depending on cases). The F values and the effective degrees of freedom are calculated in the GLM framework described by Friston and colleagues (Friston et al., 1995; Worseley and Friston, 1995) for each pair of parcel signals (as evaluated via the different settings, see Section 3.3).

The significance of nonlinear, historical, and nonlinear historical parts of interactions correspond to the following F tests:

$F^{\mathrm{NL\ DEF}} F^{\mathcal{M}^{\mathrm{NL}}/\mathcal{M}^0}$	Test for nonlinearity
$F^{\mathrm{HD\ DEF}} F^{\mathcal{M}^{\mathrm{HD}}/\mathcal{M}^0}$	Test for linearized historical dependence
$F^{\mathcal{M}^{\mathrm{NL\times HD}}/\mathcal{M}^0}$	Test for combined effects
$F^{\mathcal{M}^{\mathrm{NL\times HD}}/\mathcal{M}^{\mathrm{NL+HD}}}$	Test for the nonlinear part of the historical dependence

The F values and the effective degrees of freedom are computed in the GLM framework. Although informative, significance of F does not imply that the term tested explain a high percentage of the interaction. Under Results, we therefore also report the calculated amount of supplementary signal corresponding to each term of the interaction (e.g., nonlinear term, history term).

Hypothesis testing for whole data sets

As described above, the significance of each effect is estimated with F values for each particular interaction. Let $F^i_{b\to a}$ be the test associated with an effect of interest i. The cumulative distribution of F^i, denoted $p(F^i)$, is estimated from the set of interactions. This distribution is then compared to p^0, the distribution for an F test with same degrees of freedom under the null hypothesis, i.e., assuming that the effect tested (i.e., nonlinear interaction, or delayed interaction, or the coupling of both) is absent in the data.

A globally significant effect will result in a significant difference between the two cumulative distributions $p(F^i)$ and p^0. The comparison between these distributions is done using a Kolmogorov–Smirnov test.

All the interactions are included in the population considered for a given data set. Many of these could be neurobiologically related to one another via the neural networks mediating the activity we are examining. The number of independent interactions may therefore be smaller than the whole set. The Kolmogorov test used requires one to know the number N of independent interactions in the population, which cannot be evaluated with precision with the data used. However, we can set a lower bound to N using some hypotheses.

Weak hypothesis. Any interaction $b \to a$ is independent from interactions that do not involve a or b. A simple calculation shows that the number of independent interactions is at least greater to 50 in this case.

Strong hypothesis. Any two interactions except $b \rightarrow a$ and $a \rightarrow b$ are independent, which results in $N = 100 \times 99/2$ independent interactions.

Statistical power increases with the number of interactions considered. Therefore, results are reported in the results section for $N = 50$ and $N = 100 \times 99/2$. The D statistics are calculated and the corresponding types of interaction are considered significant if they survive 5% risk of error for the most stringent test ($N = 50$) for all subjects. The α value provided by this test represents the risk of wrongly assuming a significant effect on a given set of interactions for a given model.

3. Experiments

3.1. Data sets

Since one of our concerns is to obtain reproducible results about the nature of the coupling, we used data based on multiple subjects, sessions, and scanners.

Data set 1

This data set has been previously used and described by Simon et al. (Simon et al., 2002) This study was originally designed to identify subdivisions in the human parietal cortex and includes six tasks, performed by several subjects: grasping, pointing, visual saccades, attention, calculation, and phoneme detection. Each task was compared to a control condition designed as task-specific in block paradigm design. For each task, two sessions lasting 96×2 s are available. Each run was repeated twice yielding 12 runs per subject. Data were acquired on a 1.5-T Signa GE system.

Although computations have been made for a number of conditions, numerical results are only given in this paper for the grasping task for the six subjects exhibiting low motion parameters. Anatomical and functional voxel sizes are, respectively, (0.94 mm, 0.94 mm, 1.5 mm) and (3.75 mm, 3.75 mm, 3.8 mm).

Data set 2

The second data set has been acquired as a control on a 3-T Bruker scanner. It consists of a single run of so-called "resting state," i.e., the subject remained in a "rest" condition, eyes closed. This run lasts 348 s, TR = 3 s, yielding 116 scans. Anatomical and functional voxel sizes are, respectively, (1 mm, 1 mm, 1.2 mm) and (3.75 mm, 3.75 mm, 4.5 mm).

For all data used, subject motion parameters are inferior to 1 mm translation in all directions and 0.5° for all rotations.

While the first data set provides a statistical basis to draw task and subject-independent conclusions, the second is used as a control condition for the paradigm removal preprocessing. The following section describes those preprocessing.

3.2. Preprocessing and artifact correction

The analysis of fMRI signals generally requires a number of processing to correct for various artifacts or confounding factors.

T_1 effects

The first four scans were removed to avoid T_1 effects.

Distortion correction

Distortion correction was performed on data from the 3-T scanner. T_2^* image distortions were reduced by computing a phase map used to unwarp the images, using an algorithm similar to the one described by Jezzard and Balaban (Jezzard and Balaban, 1995). Distortion-corrected EPI images were then coregistered to a T_1 image. This step is important for the precision of the parcelling algorithm described above.

Slice timing correction

Slice timing was corrected with SPM99 using Fourier interpolation. However, since the TR is long, slice timing correction might not be possible, which can introduce artificial lags between different slices. We checked the absence of such a systematic effect, i.e., whether the amount of variance explained by lag terms increases when going to slices acquired later or earlier. Results are not shown here for the sake of concision.

Motion artifact correction

An incomplete correction of this effect is likely to induce positive or negative correlation between regions located close to the interfaces (gray/white matter, gray matter/CSF), as shown by Freire and Mangin (Freire and Mangin, 2001). We therefore checked that motion artifact removal, by regressing out subject motion parameters, as estimated with SPM99, did not alter our conclusions on the nature of connectivity links (i.e., similar α risks are obtained). In the following, only motion correction will be performed for the the removal of motion artifact, using fourier interpolation (i.e., no regression is made on motion parameters). Motion parameters are available under Experiments.

Cardiac and respiratory artifacts correction

Cardiac and respiratory artifacts were reduced using high-pass filtering. Indeed, fMRI data are most often acquired with relatively long repetition time (TR $>$ 1.5 s) in order to cover the whole brain, so that cardiac and respiratory effects are aliased and their influence partly falls in the low-frequency range. Nevertheless, studies have shown that most of the observed correlation is due to signal in the low-frequency range and have considered this band of frequency to study connectivity (Cordes et al., 2000; Lowe et al., 2000). Other studies such as that by Lund, (Lund, 2001) show the risk of correlating cardiac induced noise instead of pertinent information. Since this issue is still under study, a

Table 2
Definition of the different control settings used; each setting corresponds to a set of processing applied to fMRI data

	Setting 1 (reference)	Setting 2	Setting 3	Setting 4	Setting 5	Setting 6	Setting 7
Paradigm response removal	√	—	√	√	√	√	√
Motion artifacts removal	√	√	—	√	√	√	√
Spatial parcel averaging	√	√	√	—	√	√	√
Temporal high-pass filtering	√	√	√	√	—	√	√
Model orders ($n_{HD} = n_{NL}$)	5	5	5	5	5	2	10

strategy is used in this work in order to avoid misinterpretation of the results. In a first setting, we chose to restrict ourselves to frequencies above very low frequencies, i.e., $1/120$ s $= 3.4\ 10^{-2}$ Hz, and to use the corresponding high-pass filter. In a second setting, no temporal filter is applied.[3]

Intensity normalization

BOLD signal time courses are normalized to zero mean and unit variance.

Coactivation and paradigm removal

Connectivity has been studied both during rest and during cognitive tasks. The found connectivities may or may not be the same, as shown in (Lowe et al., 2000) and (Hampson et al., 2002). However, coactivation induced by protocol response raises a methodological question. The protocol response, as modeled using the GLM, can have a major contribution to BOLD signal at a given voxel. However, the GLM assumes no connectivity when estimating this contribution. The coactivation the GLM accounts for is a valuable information but is *not* considered an indicator of connectivity in this work. In the interaction $b \rightarrow a$, the BOLD signal of a is seen as a sum of terms, including the GLM paradigm response term, and multiple connectivity terms due to the activity of b. The activity induced by the experimental paradigm can be reduced by removing the expected signal. The simplest correction is therefore to specify a model for the studied paradigm and, after the effects estimation, to subtract the expected signal to the observed one. In the following, a very flexible model, using 11 pairs of Fourier basis set regressors, is used in order to remove as many task-related components as possible. However, if this model is not well specified and does not account for all effects, one may attribute the signal due to coactivation to connectivity. Consequently, the resting state paradigm is used, for which no response can be removed.

3.3. Control settings and validation

All settings, each corresponding to a set of preprocessing steps and model orders, are summarized in Table 2.

Our basic investigation—i.e., what is the nature of connectivity links?—is performed using a reference setting, on all subjects from the two data sets. Then, validation tests are done, using the multiple control settings on all subjects. The results obtained with the different settings are then compared. Results are considered stable if the same conclusions are drawn from the reference and the control settings.

Reference setting

The reference setting is obtained by applying the most comprehensive set of potentially useful corrections, while knowing that they might cause artifact. This justifies choosing a number of control settings to ensure robustness of our findings. The reference setting includes paradigm response removal, motion correction, gray matter averaging, and high-pass filtering.

Control setting

The multiplicity of preprocessing factors—listed above—precludes the study of all possible conditions. Instead, each potential artifact is controlled separately, yielding seven control settings.[4] These settings are designed to check that each single parameter change does not alter the results of the statistical study relative to the reference.

4. Results

In this section, we first describe the significance of the interaction terms presented in Section 2.3 and then their respective magnitude in terms of percentage of signal they are accounting for. Further, we present the robustness of those results with respect to different settings. Based on these results, a new measure of connectivity is proposed and its sensitivity is discussed.

4.1. Significance of the terms composing the interaction

The reference setting, which corresponds to the most comprehensive set of preprocessing (see Table 2), is used to compute the α risk values of observing the computed F values under the hypothesis of no effect. $N = 50$ corresponds to the weak hypothesis, and $N = 4950$ corresponds

[3] These two settings later correspond to Settings 1 and 5 in Table 2.

[4] Note that setting 2 is equivalent to setting 1 for the 'rest' data set for which no paradigm response can be removed.

Table 3
Mean percentage of data variance explained by the different terms of $b \rightarrow a$ interactions

	N_r	Average ($n = 6$)	Subject 1	Subject 2	Subject 3	Subject 4	Subject 5	Subject 6
Paradigm response	23	46.8%	45.8%	51.1%	42.4%	48.3%	52.5%	40.9%
Local signal history	4	10.2%	9.4%	11.6%	8.2%	9.3%	12.1%	11.0%
$b \rightarrow a$ interaction (all components)	25	16.4%	17.7%	13.9%	18.1%	16.5%	14.1%	18.0%
Residual "noise"	–	26.5%	27.2%	23.4%	31.3%	25.8%	21.3%	30.1%

Note. N_r is the number of regressors used to model each effect. Columns sum to 100% for each subject. Results are obtained with the reference setting (setting 1).

to the strong hypothesis (see Methods). For all subjects (from the two data sets) we observed the following:

$$N = 4950\text{: } \alpha_{NL} < 10^{-3},\ \alpha_{HD} < 10^{-10},\ \alpha_{NL\times HD} < 10^{-10}$$

$$N = 50\text{: } \alpha_{NL} > 5\%,\ \alpha_{HD} < 6\ 10^{-4},\ \alpha_{NL\times HD} < 5\ 10^{-3}.$$

The presence of delayed interactions, and the combination of these two effects, can be assumed with a very low risk of error in the data sets studied. The fact that very low α values are reached is due to the high number of interactions considered.

4.2. Magnitude of the interaction terms

The magnitudes of the different interaction terms, evaluated using setting 1, are presented in Tables 3 and 4.

Table 3 summarizes the magnitude of the different components that constitute a given parcel signal. For instance, let's consider the interaction $b \rightarrow a$. The signal of a is decomposed in the different terms presented in the first column. Nr is the number of regressors used to model the corresponding effects. Results are averaged over interactions.[5] The columns sum to 100% since the noise term is presented as well. The large part of the variance captured by the paradigm response is due to two factors: (a) the model chosen for this effect contains a large number of regressors (see Methods) and (b) very low frequencies have been removed. A large number of regressors was chosen for this term to prevent the interaction term from containing some information attributable to the experimental paradigm. This is at the expense of some overfitting of the activation signal. We note also that the local signal history is not negligible, about 75% of the interaction term.

Interestingly, the variance explained by the different terms seems to be reproducible across subjects.

Table 4 shows how the interaction term (third row in Table 3) is decomposed. We note that nonlinear and lagged terms explain a nonnegligible part of the data compared to what the instantaneous and linear term (analogous to correlation) does. The lagged interaction term is more important than the nonlinear interaction term. The nonlinear lagged interaction term explains the largest part of the variance, but this has to be seen given the high number of regressors used to model this term.

To summarize, functional connectivity is likely to contain more than simply instantaneous and linear interactions, and effects such as local recent history of the signal studied, nonlinear couplings, and history-dependent couplings may have to be taken into account (see below, Validation).

We then studied the relation between the amount of nonlinear and history-dependent terms with the amount of connectivity as measured by the proportion of variance explained (corresponding to line 3 in Table 3). We found that the proportion of the lagged part of the interaction increases with the amount of interaction (mean correlation $C = +0.15$, $n_{\text{interactions}} \approx 10^4$, $\alpha < 10^{-10}$), while the nonlinear instantaneous or nonlinear lagged interaction terms do not ($\alpha < 1\%$). This result tends to demonstrate that the linear lagged term has an important role when modeling functional connectivity, while the role of nonlinear terms is less clear.

As a consequence, one can reasonably choose to take the

[5] Note that the terms paradigm response, local history, and residual variance are specific to a given parcel and not to an interaction, such that those can be computed across parcels rather than across interactions.

Table 4
Percentage of data variance explained by specific types of interactions, relative to the total variance explained by all interactions

	N_r	Average ($n = 6$)	Subject 1	Subject 2	Subject 3	Subject 4	Subject 5	Subject 6
Linear instantaneous	1	20.8%	20.3%	18.4%	15.5%	23.0%	28.1%	19.7%
Nonlinear, instantaneous	4	15.7%	15.6%	16.1%	15.3%	15.9%	15.3%	16.0%
Linear, lagged	4	25.8%	28.2%	24.2%	25.5%	26.1%	24.5%	26.4%
Other nonlinear, lagged	16	37.6%	36.0%	41.3%	43.6%	34.9%	32.0%	37.8%

Note. The sum of the terms equals 100%. Setting 1 is used.

Table 5
The two models used for the new connectivity measure $\mathcal{F}$

$\mathcal{M}^{AR}_{b\to a}$:	$\mathcal{G}^{AR}_{b\to a} = \{x_a(t-k\tau)_{1<k<M}\}$
$\mathcal{M}^{HD}_{b\to a}$:	$\mathcal{G}^{HD}_{b\to a} = \{x_a(t-k\tau)_{1<k<M}, x_b(t), x_b(t-k\tau)_{1\le k\le M}\}$

Note that $\mathcal{M}_{HD}$ was defined in Table 1.

supplementary information provided by history in account when assessing connectivity.

4.3. Validation: robustness of results to methodological settings

In order to check that the results of statistical tests are not altered when choosing different methodological settings, the same study is now performed for all settings defined in Table 2.

Lagged interactions

Independent of the control setting setting used, we get across subjects

$$N = 4950: \qquad \alpha_{HD} < 10^{-10}, N = 50: \alpha_{HD} < 3\,10^{-3}.$$

Therefore, the results on the presence of lagged interactions are robust with the effects of paradigm removal, motion, spatial, and temporal filtering and model orders considered.

Nonlinear interactions

Independent of the number of $b \to a$ interactions considered, the nonlinear interaction term is not significant for at least one subject in settings 2, 3, 4, 5, and 7 ($\alpha > 5 \cdot 10^{-2}$ for at least one subject) and significant at $\alpha = 10^{-3}$ for setting 6 (all subjects).

Nonlinear lagged interaction is significant at $\alpha = 10^{-2}$ across subjects, independent of the setting used.

4.4. An updated connectivity measure

What information should be used?

The results described above suggest the definition of an updated connectivity measure. Since the linear lagged term is important in magnitude, related to the instantaneous term (see Section 4.2), and consistently significant across control settings, this term is included in the definition of this measure. In the following, we define $\mathcal{F}$ as a new connectivity measure.

Definition of a new connectivity measure: $\mathcal{F}$

The results above suggest taking history into account when assessing functional connectivity. To do so, we use two models, the autoregressive model $\mathcal{M}^{AR}$ and $\mathcal{M}^{HD}$, which was designed above (see Table 1) under Methods.

The autoregressive model $\mathcal{M}^{AR}$ serves as a null model for comparison with $\mathcal{M}^{HD}$.

The additional variance explained by model $\mathcal{M}^{HD}$ relative to the autoregressive model $\mathcal{M}^{AR}$ can be seen as the global influence of the distant region b on region a. This variance can naturally be interpreted as a measure of connectivity taking into account the signals' history. The corresponding F test is a measure of this influence. This measure is denoted

$$\mathcal{F}_{b\to a} \overset{def}{=} \mathcal{F}^{\mathcal{M}^{HD}/\mathcal{M}^{AR}}_{b\to a}.$$

This measure is oriented and causal in the sense of Granger (see note in Appendix B for a definition). The $\alpha_{\mathcal{F}}$ and α_C risks for an interaction can be determined (see Appendix A and (Cao and Worsley, 1999)). These α risks are of the same nature and can be used as a measure of differential sensitivity between correlation C and $\mathcal{F}$.

Properties of $\mathcal{F}$

Sensitivity. The sensitivities of C and $\mathcal{F}$ are compared through the number of significant interactions at a given α level. The proposed measure $\mathcal{F}$ shows a higher sensitivity relative to correlation for any α below 5% for all subjects (reference setting).

Comparison of the connectivity maps. The increased sensitivity is illustrated in Fig. 3. To construct this figure, a reference parcel is chosen at the maximum of activity (sub-

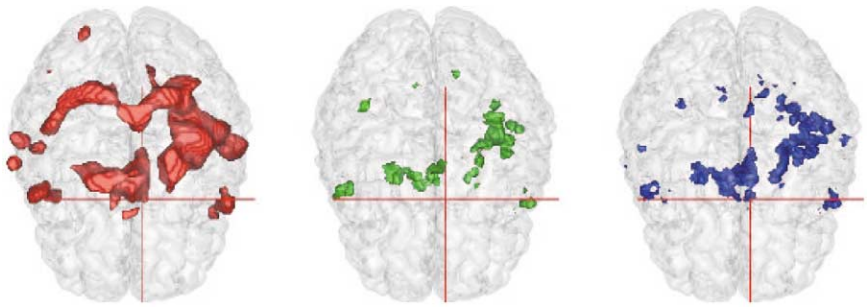

Fig. 3. Activation map and connectivity maps assessed with C and $\mathcal{F}$. (Grasping paradigm, subject 1, setting 1). The maximum activation parcel is the connectivity reference. (a) : Activation at $\alpha = 10^{-10}$ (uncorrected p value) (b) : Significant correlates at $\alpha = 10^{-3}$ (c) : Significant $\mathcal{F}$ values at $\alpha = 10^{-3}$

ject 1) found in the motor cortex. The SPM activation map is shown on the left, the connected parcels as measured with C in the middle, and those with with $\mathcal{F}$ on the right. Connectivity results are displayed for $\alpha \leq 10^{-3}$. Recall that connectivity maps are computed from the part of signals not explained by the very flexible activation model (Setting 1). The connectivity maps are restricted to the activation location although other regions are involved as well. There are very few areas found to be connected with C and not with $\mathcal{F}$. On the contrary, some regions such as the contralateral supplementary motor area are found connected to the seed region with $\mathcal{F}$ and not with C (see crosshair in Fig. 3).

Orientation. The percentage of symmetrical relations, i.e., for which directional interactions are both significant, is evaluated, depending on the threshold value of α. This may give elements to determine whether causality could be retrieved from fMRI data. We evaluated the percentage of symmetrical and asymmetrical relations by counting the part of α-significant $\mathcal{F}_{b\to a}$ for which $\mathcal{F}_{a\to b}$ is α-significant too.

For $\alpha \leq 10^{-3}$ and across subjects, more than 79% of interactions are found to be symmetrical. This indicates that causality might not be easily retrieved from data.

5. Discussion

We found that connectivity measured with BOLD fMRI signals is unlikely to contain only linear instantaneous information. This finding appeared robust to an important number of preprocessing steps applied on the data. A new connectivity measure has been designed that includes lagged signals of both the local and the distant regions. The past signal (time strictly less than t) of a region can therefore partly predict the response of an other area at time t, independent of the past of that area. The biological origins of this effect are not straightforward. In the following, we examine some possible causes and discuss related works and questions.

5.1. Possible origins of delayed interaction

Models of neural information from BOLD signals

While it is not clear whether the origin of noninstantaneous coupling is neuronal or vascular, there are recent attempts to consider neural response and the vascular responses in a common model in which parameters have to be identified from the BOLD responses (Friston et al., 2002). The observed BOLD signal is proposed to be the output of a vascular model of which the input are neural responses modulated by the experimental paradigm. This has the important advantage that the coupling can then be explicitly formulated at the neural level. However, it relies on the form of the models and on an appropriate parameter identification. Moreover, a growing literature now explicitly employs neural network models to simulate both neuronal activities and fMRI data, e.g. (Tagamets and Horwitz, 2001; Arbib et al., 2000). These approaches could be used to help quantify the links between the phenomena at the neural level and their consequence on the measured hemodynamic response (HRF). In order to investigate some possible causes of the effects we observed, let us consider the signals both at the neural and at the hemodynamic levels. Let us suppose—in a first approximation—that each neural signal, e_i, and its corresponding BOLD response, x_i, are linearly coupled by convolution with an unknown hemodynamic response function. The following diagram summarizes the relations between the neural and BOLD signals at parcels a and b.

$$\begin{array}{ccc} e_b & \xrightarrow[r_{b\to a}]{} & e_a \\ \downarrow *H_b & & \downarrow *H_a \\ x_b & \xrightarrow[R_{b\to a}]{} & x_a \end{array}$$

Three possible reasons of the role of history in BOLD signals interactions are now reviewed.

Neural signals coupling

First, results may be explained by coupling at the neural level, such that past signal from region b partly predicts signal of region a. Let us suppose that the coupling between neural signals e_a and e_b is of the form $e_a(t) = A * e_b$. Let us also suppose the linearity of transduction from the neural to the BOLD level, modeled by convolution of a constant HRF. This leads to a relation between BOLD signals which explains influence of history in interactions: $x_a(t) = A * x_b$. However, the time span influencing interaction at the neural level would have be of the order of the hemodynamic (second). This may involve complex feedback loops between several regions.

Better estimation of noisy BOLD signal

Second, noise temporal autocorrelation of BOLD signal may partly explain this effect. Indeed, the estimation of x_b at time t can be improved relative to its measure $x_b(t)$ using the information contained in its previous (and next) values. Consequently, the influence of $x_b(t)$ on $x_a(t)$ may be fitted more precisely, only because of the effect of noise averaging through time and in the absence of time component in the interactions. However, the reported autocorrelation of noise in BOLD signals is much shorter than the autocorrelation of BOLD signals, which corresponds to the time spans we considered. This possible origin does not seem to be sufficient to explain the observed dependence of interactions to history.

Variability of local HRFs

Third, the difference between local vascular responses may cause this effect, especially if they present different latencies. Let us suppose that regions a and b share the same neural activity $e(t)$. Supposing that $H_a \neq H_b$, it comes that correlation $\mathrm{corr}(e_a, e_b) < 1$, although neural activities are equal. This can explain why information can be lost when evaluating connectivity using correlation. Using Fourier transforms it can be shown that there exists $H_{b\to a}$ such as $x_a = H_{b\to a} * x_b$.[6]

In the hypotheses of neural signals coupling and variability of local HRFs, BOLD signals would be linked through a convolution operation, which the proposed measure of connectivity can account for. A large part of the lag may also stem from the fact that later signals come from small draining veins. At last, it is possible that the effect observed is due to a combination of these causes.

5.2. n-to-n connectivity

This study only takes into account pairwise connectivity. Systematical research of n-to-1 connectivity would be difficult to handle, and it is likely that not enough data would be available to detect these relations with sufficient significance (as evaluated by the value of F tests). Moreover, n-to-n connectivity can be studied using 1-to-1 connectivity measures using certain hypotheses (see (Lohmann and von Cramon, 2001)).

5.3. Improving the estimation of functional connectivity

Two main strategies to determine functional connectivity seem to emerge. The first one, presented in this paper, builds and validates a measure of connectivity acknowledging the effects of history. The second would use "deconvolution" of the BOLD signal as a prior step, i.e., simultaneous estimation of the local HRF and the corresponding "neural" signal. The study of connectivity could then be done on the approximated neural signal. The issue is to determine whether deconvolution introduces sufficiently low incertitude on the resulting signal. As far as we know, published studies about deconvolution estimate the local HRF without evaluating the neural component. Such deconvolution is addressed by Glover (Glover, 2001).

6. Conclusion

In the light of hypothesis testing, the recent history of BOLD signals seems to contain information that is useful to evaluate functional connectivity between brain regions. These results have been found robust with a number of preprocessing steps applied to the data. From these findings, a measure of connectivity that takes signals history into account has been proposed. This measure is shown to have a better sensitivity than correlation. The relations found using this measure are mostly reciprocal.

To further investigate the influence of the neural or vascular components in functional connectivity, it is hoped that conjoint fMRI/EEG may in the near future provides good priors for the neural signal and its localization and yield functional connectivity measures close to neural activity with good temporal and spatial characteristics.

7. Appendix A

The α risk for an interaction, whose F value is f, is given by

$$\alpha_f = P(F \geq f) = 1 - F_{\mathrm{cdf}}(f, \nu_1, \nu_2),$$

where F_{cdf} denotes the F distribution cumulative density function, (ν_1, ν_2) are the degrees of freedom of the F test, and f is a given threshold.

In the same way, the α risk of assessing a functional link given an observed correlation ρ can be calculated (see Cao and Worsley, 1999).

$$\begin{aligned}\alpha_\rho &= P(C \geq \rho)\\ &= T_{\mathrm{cdf}}(-\sqrt{\nu - 1}\rho(1-\rho^2)^{-1/2}, (\nu - 1)),\end{aligned}$$

where T_{cdf} denotes the T distribution cumulative density function, ν is the degrees of freedom of the considered time courses, and ρ is a given threshold.

8. Appendix B

8.1. Time embedding

Classically, nonlinear relations between systems can be studied in an embedding space (see (Broomhead and King, 1987)). Time embedding focuses on the recent history and/or future of the variable of interest at a given time, e.g.,

$$\begin{aligned}X_a(t) = [&x_a(t), x_a(t - k\tau),\\ &x_a(t - 2k\tau), \ldots, x_a(t - (n-1)\tau)],\end{aligned}$$

rather than its present value alone $x_a(t)$. X_a is called a state, and n is the embedding dimension.

Modeling signal x_a is done through its values in the state space. The relation sought in the most general case is of the form $\forall t, [x_a(t), x_b(t)] = F([X_a(t), X_b(t)])$.

Both Volterra kernels and our approach implement time embedding. Note that nonlinear interactions in the original space of signals values may or may not be linear in an appropriate embedding space.

8.2 Volterra autoregressive models

Volterra autoregressive models are based on time embedding. They have been used to evidence nonlinearities between

[6] $H_{b\to a} = R^{-1}\{[R(H_a)]/[R(H_b)]\}$, where R is Fourier transform.

EEG channels (Schiff et al., 1995) or BOLD signals (Friston et al., 1998; Harrison et al., 2002). Using the second order Volterra kernels to form a nonlinear autoregressive model (NLAR), where $x = (x_a, x_b)$ is regressed, and the q and k are the interaction parameters to be determined, we can model the $x_a \leftrightarrow x_b$ coupling as follows for x_a (equations for x_b are symmetrical):

$$
\begin{aligned}
x_a(t) = & \int_{-\infty}^{\infty} g_a(\tau)\, x_a(t-\tau)\, d\tau && \text{Recent history term (a)}\\
& + \int_{-\infty}^{\infty} g_b(\tau)\, x_b(t-\tau)\, d\tau && b \rightarrow a \text{ interaction term (b)}\\
& + \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} k_{a,b}(\tau_1, \tau_2)\, x_a(t-\tau_1)\, x_b(t-\tau_2)\, d\tau_1 d\tau_2 && \text{Complex interaction term (c)}\\
& + \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} k_{a,a}(\tau_1, \tau_2)\, x_a(t-\tau_1)\, x_a(t-\tau_2)\, d\tau_1 d\tau_2 && \text{Other nonlinear terms (d)}\\
& + \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} k_{b,b}(\tau_1, \tau_2)\, x_b(t-\tau_1)\, x_b(t-\tau_2)\, d\tau_1 d\tau_2 && \text{Other nonlinear terms (e)}
\end{aligned}
$$

Volterra kernels provide a flexible model, but at the cost of a potentially high number of regressors, and therefore should be used conjointly with model selection criteria to avoid data overfitting.

Moreover, this model has two drawbacks that make it seemingly unappropriate for our study. First, it does not enable us to disentangle orientation of relation $a \leftrightarrow b$ from the retrieved parameters. Second, it does not respect the causality in the sense of Granger,[7] which seems a reasonable prior for an interaction. We therefore chose to focus on orientated and Granger causal interactions, and to build corresponding models.

8.3. Remark on model design

The models defined under Methods can be seen as particular cases of Volterra autoregressive models, where certain g and k parameters are fixed to 0 by construction. Our aim here is to determine what data could be used with practical interest to provide a better measure of connectivity. Our definition of models is therefore guided by the hypotheses to be tested, although they may be restrictive hypotheses.

The different models should be able to represent, in the most general manner possible, the hypotheses made on the nature connectivity links. However, a too-flexible model would reduce the statistical power of the study.

Our first interest is to determine whether lagged interaction and nonlinear instantaneous interaction are significant (independent of each other). Besides lagged interaction and nonlinear interaction, only the possible coupling between history and nonlinearity seemed to us to be of interest. Although parts of terms (a), (b), and (e) in the equation above are included in our models since they reflect effects of interest, terms (c) and (d) have a difficult interpretation and have therefore not been taken into account.

Orientation

Considering all oriented interactions (i.e., both $b \rightarrow a$ and $a \rightarrow b$) allows us to retrieve the putative orientation of a given connectivity link. Whether this information is retrievable may be of interest. Moreover, this choice makes it possible to check a posteriori whether found relations are reciprocal.

Granger causality

Assuming precedence of signal x_b relative to signal x_a in a putative $b \rightarrow a$ relation implies taking into account lag terms but no prospective term. As a consequence, we use two models whose regressors are subsets of Volterra kernels including past terms only.

[7] Granger causality states that in the $b \rightarrow a$ causality, $a(t)$ depends on the past values of b, $b(s \leq t)$, and is independent of its future values $b(s > t)$.

Acknowledgments

The authors are deeply in debt to Alexis Roche for invaluable comments on the manuscript. This work was partly funded by "Action Concertée Incitative Plasticité" (Ministère de la Recherche) and Programme cognition et traitements de l'information (CNRS).

References

Arbib, M.A., Billard, A., Iacoboni, M., Oztop, E., 2000. Synthetic brain imaging: grasping, mirror neurons and imitation. Neural Netw. 13 (8–9), 975–997.

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

Broomhead, D.S., King, G.P., 1987. Extracting qualitative dynamics from experimental data. Phys. D. 20, 217–236.

Büchel, C., Friston, K.J., 1997. Modulation of connectivity in visual pathways by attention: cortical interactions evaluated with structural equations modelling and fMRI. Cereb. Cortex 7 (8), 768–778.

Cao, J., Worsley, K., 1999. The geometry of correlation fields with an application to functional connectivity of the brain. Ann. Appl. Probabil. 9 (4), 1021–1057.

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.

Della-Maggiore, V., Sekuler, A.B., Grady, C.L., Bennett, P.J., Sekuler, R., McIntosh, A.R., 2000. Corticolimbic interactions associated with performance on a short-term memory task are modified by age. J. Neurosci. 20 (22), 8410–8416.

Flandin, G., Kherif, F., Pennec, X., Malandain, G., Ayache, N., Poline, J.-B. Improved detection sensitivity of functional MRI data using a brain parcellation technique, in: Proceedings of the 5th International Conference on Medical Image Computing and Computer Assisted Intervention, LNCS 2488 (Part I), Tokyo, Japan, September 2002. Springer-Verlag, Berlin, pp. 467–474.

Freire, L., Mangin, J.-F., 2001. Motion correction algorithms may create spurious brain activations in the absence of subject motion. NeuroImage 14 (3), 709–722.

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., Joseph, O., Rees, G., Turner, R., 1998. Non-linear event-related response in fMRI. Magn. Reson. Med. 39 (1), 41–52.

Friston, K.J., Buechel, C., 2000. Attentional modulation of effective connectivity from V2 to V5/MT in humans. Proc. Natl. Acad. Sci. USA 97 (13), 7591–7596.

Friston, K.J., Harrison, L., Penny, W., 2002. Dynamic causal modelling. NeuroImage, submitted.

Glover, G.H., 2001. Physiological noise in oxygenation-sensitive magnetic resonance imaging. Magn. Reson. Med. 46 (4), 631–637.

Goldman, R.I., Stern, J.M., Engel, J., Cohen, M.S., MS. 2002. Simultaneous EEG and fMRI of the alpha rhythm. NeuroReport, 13(18).

Hampson, M., Peterson, B.S., Skudlarski, P., Gatenby, J.C., Gore, J.C., 2002. Detection of functional connectivity using temporal correlations in MR images. Hum. Brain Mapp. 15 (4), 247–262.

Harrison, L., Penny, W., Friston, K., 2002. Multivariate autoregressive modelling of fMRI time series, in: NeuroImage (HBM'02), Sendai, Japan.

Jezzard, P., Balaban, R.S., 1995. Correction for geometric distortion in echo planar images from b_0 field variations. Magn. Reson. Med. 34 (1), 65–73.

Kiebel, S., Poline, J.-B., Friston, K.J., Holmes, A.P., Worsley, K., 1999. Robust smoothness estimation in statistical parametric maps using standardized residuals from the general linear model. NeuroImage 10 (6), 756–766.

Lohmann, G., von Cramon, D., 2001. Detecting functionally coherent networks in fMRI data of the human brain using replicator dynamics, in: International Conference on Information Processing in Medical Imaging, pp. 218–224.

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.

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.

Lund, T.E., 2001. fcMRI—mapping functional connectivity or correlating cardiac-induced noise? Magn. Reson. Med. 46 (3), 628–629.

Mangin, J.-F., Coulon, O., Frouin, V., Robust brain segmentation using histogram scale-space analysis and mathematical morphology, in: Wells W.M., Colchester, A., Delp, S. (Eds.), Proceedings of the 1st International Conference on Medical Image Computing and Computer Assised Intervention, LNCS-1496, MIT, Boston, October 1998. Springer-Verlag, Berlin, pp. 1230–1241.

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

Poldrack, R., 2000. Imaging brain plasticity: a theoretical review. NeuroImage 12 (1), 1–13.

Schiff, N.D., Victor, J.D., Canel, A., Labar, D.R., 1995. Characteristic nonlinearities of the 3/second ictal EEG identified by nonlinear autoregressive analysis. Biol. Cybernet. 72, 519–526.

Simon, O., Mangin, J.-F., Cohen, L., Le Bihan, D., Dehaene, S., 2002. Topographical layout of hand, eye, calculation, and language-related areas in human parietal lobe. Neuron 33, 475–487.

Tagamets, M.A., Horwitz, B., 2001. Interpreting pet and fMRI measures of functional neural activity: the effects of synaptic inhibition on cortical activation in human imaging studies. Brain Res. Bull. 54 (3), 267–273.

Tononi, G., McIntosh, A., Russell, D., Edelman, G., 1998. Functional clustering: identifying strongly interactive brain regions in neuroimaging data. NeuroImage 7 (2), 133–149.

Woolrich, M., Ripley, B., Brady, M., Smith, S., 2001. Temporal autocorrelation in univariate linear modelling of fMRI data. NeuroImage 14 (6), 1370–1386.

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