Available online at www.sciencedirect.com
ScienceDirect
Magnetic Resonance Imaging 25 (2007) 1347–1357
MAGNETIC RESONANCE IMAGING

Original contributions

Assessing functional connectivity in the human brain by fMRI☆

Baxter P. Rogers[a,c], Victoria L. Morgan[a,c], Allen T. Newton[b,c], John C. Gore[a,b,c,*]

[a]Department of Radiology and Radiological Sciences, Vanderbilt University, Nashville, TN 37232, USA
[b]Department of Biomedical Engineering, Vanderbilt University, Nashville, TN 37232, USA
[c]Vanderbilt University Institute of Imaging Science, Vanderbilt University, Nashville, TN 37232, USA

Received 6 June 2007; accepted 11 June 2007

Abstract

Functional magnetic resonance imaging (fMRI) is widely used to detect and delineate regions of the brain that change their level of activation in response to specific stimuli and tasks. Simple activation maps depict only the average level of engagement of different regions within distributed systems. FMRI potentially can reveal additional information about the degree to which components of large-scale neural systems are functionally coupled together to achieve specific tasks. In order to better understand how brain regions contribute to functionally connected circuits, it is necessary to record activation maps either as a function of different conditions, at different times or in different subjects. Data obtained under different conditions may then be analyzed by a variety of techniques to infer correlations and couplings between nodes in networks. Several multivariate statistical methods have been adapted and applied to analyze variations within such data. An approach of particular interest that is suited to studies of connectivity within single subjects makes use of acquisitions of runs of MRI images obtained while the brain is in a so-called steady state, either at rest (i.e., without any specific stimulus or task) or in a condition of continuous activation. Interregional correlations between fluctuations of MRI signal potentially reveal functional connectivity. Recent studies have established that interregional correlations between different components of circuits in each of the visual, language, motor and working memory systems can be detected in the resting state. Correlations at baseline are changed during the performance of a continuous task. In this review, various methods available for assessing connectivity are described and evaluated.

Keywords: Functional connectivity; Human brain; fMRI

1. Introduction

Functional magnetic resonance imaging (fMRI) is well established as a method for the detection and delineation of regions of the brain that change their level of activation in response to specific experimental conditions. fMRI studies typically use "snapshot" imaging methods such as echo-planar sequences that are sensitive to changes in blood-oxygenation-level-dependent (BOLD) signal, which reflects neuronal activation, albeit indirectly [1]. fMRI studies produce activation maps that typically depict the average level of engagement during a specific task or in response to a specific stimulus of different regions in the brain. These may be compared between conditions or between subjects to evaluate the relative magnitudes of different responses. This is the basis of using fMRI for brain mapping and for comparing the activation patterns produced by different stimuli or between groups. However, appropriate fMRI data may also be analyzed in greater depth to reveal how components of large-scale distributed neural systems are coupled together in performing specific functions. The organization, interrelationship and integrated performance of these different regions are generally described by the term "functional connectivity."

For this review, we will restrict our use of "functional connectivity" to mean the quantification of the operational interactions of multiple spatially distinct brain regions that are engaged simultaneously in a task. We will further restrict our discussion to connectivity measures derived from fMRI activation data alone. Currently, there is no consensus on the most accurate or efficient method of detecting or measuring functional connectivity using fMRI [2], although there are considerable interest and activity in this field. However, while specific analytic techniques vary, a common feature of

☆ Presented at the International School on Magnetic Resonance and Brain Function meeting "Magnetic Resonance Combined With Electrophysiology and With Molecular and Pharmaceutical Action for the Study of Brain Function" held in Erice, Italy, May 21–28, 2006.

* Corresponding author. Vanderbilt University Institute of Imaging Science, Nashville, TN 37232, USA. Tel.: +1 615 322 8359.
E-mail address: john.gore@vanderbilt.edu (J.C. Gore).

0730-725X/$ – see front matter © 2007 Elsevier Inc. All rights reserved.
doi:10.1016/j.mri.2007.03.007

multiple fMRI assessments of connectivity is the use of correlations or covariances of activities derived from BOLD data. The objectives of this article are to introduce and explain several types of proposed analyses that attempt to quantify connectivity using such statistical properties. In addition, we will discuss the origins and nature of three primary sources of variance in these data: intersubject variance, task-related variance and intrinsic or steady-state variance. Finally, we will identify some confounding factors in the measurement of functional connectivity that may obfuscate conclusions in practical applications. We postulate that the ultimate value of fMRI in studies of brain function will depend not only on our ability to map activity patterns to reveal neural functional architecture but also on our ability to understand how brain regions work together to accomplish specific tasks and behaviors. Methods for assessing functional connectivity are key for obtaining such insights.

2. Statistical methods for the analysis of connectivity

Methodological approaches to the study of connectivity using BOLD data may be broadly divided into those that are more data driven and attempt to map connectivity in the whole brain (Fig. 1) and those that use prior knowledge or hypotheses to limit analysis to a restricted set of regions (Fig. 2). The first category of methods includes seed–voxel correlation maps, Granger causality maps derived from bivariate autoregressive models [3] and psychophysiological interaction (PPI) maps [4]. Other mapping techniques such as principal components analysis (PCA) and some applications of partial least squares (PLS) analysis create whole-brain maps of functional networks in which regions share some features of interest. The alternative to these mapping techniques is to use a model that attempts to describe the relationships between a number of selected regions of interest, wherein region-specific measurements such as MRI time series are extracted from whole-brain data prior to the connectivity modeling stage. This category includes structural equation modeling (SEM) [5] as well as multivariate autoregressive (MAR) modeling [6]. Correlation, PCA and PLS methods may be applied in this way as well, but these are less common approaches.

A second distinction may be made between methods that consider only correlation and ignore issues of causation and influence and methods that attempt to describe or make inferences about the direction of influence between regions. These two categories of analysis are often referred to as functional connectivity and effective connectivity, respectively [7]. Techniques in the first group that consider only correlations between regions include mapping using seed–voxel correlations, PCA and PLS methods. Techniques in the second group use more elaborate models and additional assumptions applied to calculated correlations or covariances to address questions about directional influences and include mapping based on PPIs, SEM, Granger causal mapping and MAR modeling.

Seed–voxel correlation mapping is one of the simplest techniques for studying functional connectivity: the correlation coefficient between the fMRI signal at different times and measurements of activation in a seed region is calculated separately for each voxel in the brain and may be displayed as a parametric image. This approach has most commonly been applied to steady-state time-series data, where experimental conditions do not change during an imaging run. It is possible to perform seed–voxel analysis iteratively so that connectivity maps made using one seed region can be used to identify other regions to be used as subsequent seeds, after which the entire process can be started again [8]. This iterative approach can reveal networks of functionally connected regions that would otherwise remain unidentified.

The method of SEM, in essence, takes a step beyond simply calculating correlations and allows the development of estimates of the directions of influences between variables. These estimates are typically calculated via a maximum likelihood procedure, which iteratively adjusts the parameters of the assumed causal model until predicted correlations match correlations in measured data as closely as possible. In a typical fMRI application, interconnections are specified between regions of interest based on prior anatomical knowledge or hypotheses, and the strengths of those connections are then estimated from region time-series data. Connection strengths, or path coefficients, can then be compared across experimental groups or conditions using a variety of statistical approaches. It is important to note that SEM cannot generally answer questions about the direction of influences between regions, but rather gives strengths of influences assuming that the specified causal structure is correct. Multiple regression models involving several regions, but with only a single dependent variable [9], are actually a special case of SEM. SEM has been used with fMRI data from a wide variety of experiments, including studies of visual attention [10–12], visual learning [12,13], grammar learning [14], tone listening [15], semantic and episodic memory [16], working memory [17], reading [18] and finger movement [19,20]. A number of variations of the technique have appeared, with many differing in the precise method used for the calculation of a correlation matrix (see Horwitz [2] for a discussion). Some of these are intended to address specific weaknesses of the approach, such as mathematical constraints on the number of free parameters [18] or the failure of some models to account for the time dependence of fMRI data [21].

A PPI [4] is a stimulus-dependent or context-dependent change in the influence of one brain region on another. A PPI can be identified using a linear regression model wherein a voxel's data are predicted by an influencing region's data, a predictor of stimulus-related signal changes and the product of the two terms (the interaction). If the contribution of the interaction term to voxel measurements is significant, that implies that the contribution of the influencing region depends on experimental context or,