A functional connectivity primer

Daniel Y. Kimberg, PhD
Center for Functional Neuroimaging
University of Pennsylvania

Revision 1, July 2008: initial revision

Introduction

The use of in-vivo imaging methods to study connectivity in the brain has grown dramatically over the past several years. While a large part of this growth is due to the availability of diffusion tensor imaging (DTI), methods for examining connectivity using widely available BOLD contrast data have also been growing. The umbrella term for the latter, "functional connectivity," was defined by Friston et al. (1993) as "temporal correlation between spatially remote neurophysiological events." This definition captures the essentially correlative nature of these methods – regions are considered *functionally connected* if their activity is in some way correlated, regardless of the mechanism underlying the correlation. As contrasted with effective connectivity ("the influence that one neural system exerts over another either directly or indirectly"), functional connectivity doesn't necessarily imply a physical pathway, and potentially includes patterns of connectivity that are entirely mediated by the common influence of some external event on distant neural areas.

Much of the study of functional connectivity (FC) is carried out by examining inter-regional correlations in resting BOLD data, an approach that is traced to a study by Biswal et al. (1995; see Rogers et al., 2008 for a brief note on earlier related methods), who observed correlations between activity in left and right somatosensory cortex during resting BOLD. Temporal correlations in resting data are of special interest because they are not easily explained by externally imposed task demands (although they may be influenced by endogenously driven behavior at rest). At the same time, a growing literature on neural activity during the kinds of passive experimental conditions ordinarily used as control or rest conditions during fMRI suggests that resting BOLD is more than just the absence of cognitively evoked activity. By contrast, connectivity measures from data collected during task performance risk discovering trivial associations of little novel interest – there is an obvious reason why regions in the left and right motor cortex should be correlated during bimanual finger tapping, and a functional connectivity analysis sheds no new light on these processes.

A major practical advantage of connectivity studies carried out in resting data is that the same data may be used repeatedly. At this point, a great deal of resting BOLD data is available either publicly or through local repositories, and many questions can be addressed without the collection of new data. Further, Fair et al. (2006) have demonstrated that rest periods from blocked-design BOLD studies yield similar results to what can be obtained from simple resting BOLD data, which opens up the possibility of drawing on an even larger store of existing data.

Understanding spatial patterns of intercorrelation may be important for another reason. Fox et al. (2006) have shown that some of the coherent signal in resting BOLD contributes roughly linearly to task-evoked BOLD. In this way, better understanding of the task-independent component of the signal can lead to markedly better sensitivity to detect task-evoked activation.

The purpose of this article is to convey the essential details of the most common

approaches to studying functional connectivity in BOLD data, mostly resting BOLD. I first describe typical processing steps involved in these analyses, then describe a few variants of the basic method, and finally describe available software tools for carrying out these analyses.

Inter-region time series correlations

The most straightforward approach to measuring functional connectivity in resting BOLD data is to examine the correlations between time series data acquired from different voxels or regions. Typically, the signal is first extracted from some seed region or voxel, identified either functionally (it was activated in a separate dataset) or anatomically. A correlation map may then be constructed by calculating the correlation between that seed signal and the signal from each other voxel in the brain. Here I describe some of the design and analysis considerations specific to functional connectivity analysis.

Data collection. Data may be collected with typical BOLD acquisition parameters. The amount of data required depends on many factors, including the field strength of the scanner, the other details of the scanning protocol and hardware, the nature of the paradigm if any, the nature of the subject population, and the planned analyses. Previous studies provide some points of reference, with typical studies of resting BOLD at 1.5T including roughly 20-30 minutes of scanning. This appears to be somewhat arbitrary – i.e., adequate solely because it has turned up meaningful results in the past. The basis for a proper power analysis is unclear, but it would probably depend on some *a priori* reasoning to identify a minimum inter-regional correlation coefficient of interest.

Preprocessing. Preprocessing of BOLD data for resting connectivity studies is similar to that for functional studies, with some additional considerations. Here I describe the relevant steps.

Correction for slice timing. If a whole brain volume is acquired over 3 seconds, then nearly 3 seconds may separate the first and last slices acquired. Naïvely calculating correlations between voxels is therefore likely to suffer from a temporal offset between the voxels. Although restricting the analysis to low frequencies (see “frequency domain filtering” below) minimizes the effects of this error, FC studies (and most fMRI studies) typically interpolate the data to correct for this temporal offset.

Intensity correction. Due to the order of acquisition, intensity scaling may differ between slices. Because FC methods draw direct comparisons between slices, correction of inter-slice intensity differences may be more important than it would be for BOLD studies, where each voxel is treated independently. Although simple scaling and offset differences between voxels would not affect correlations, these intensity differences can have an undesired effect of weighting the contributions of different slices when smoothing or spatially transforming images.

There may also be global intensity differences between BOLD runs where more than one run per subject is used to calculate the correlation. Scaling (multiplicative) differences may be factored out by scaling each run to a global mean. Offset (additive) differences can be factored out in the model stage of activation studies, or by subtracting off mean differences if a regression isn't otherwise needed.

Rigid realignment. Rigid realignment of the brain is required to correct for head movement during the scan. Rigid registration allows a subset of possible affine transforms: translations in and rotations around all three axes. Head movement is a generally worrisome problem for all BOLD studies, and subjects with excessive head movement are typically excluded. What qualifies as “excessive” is an open question,

but it is presumed that the same heuristic guidelines used for fMRI may be used for FC studies.

Spatial normalization. For multi-subject studies, spatial normalization (aka "brain warping" or "nonlinear registration") to a common template is required to align corresponding structures (as well as possible). When that common template is in a standard space (typically MNI space, see Collins, 1994), labellings from atlases in that space may be used to query specific anatomical regions. Specifically, there are several atlases in MNI space that assign neuroanatomical names to each voxel, and at least one map of Brodmann's areas.

Resampling. Both rigid realignment and spatial normalization entail first calculating a spatial transformation of each volume, a coordinate transformation that specifies, for each voxel in the new volume, the location in the old volume from which to take the signal value. This location is generally between voxels, and calculating values for between-voxel locations entails an interpolation step that loses some of the information from the original image. To avoid compounding this effect by first realigning and then normalizing, the two transformations may be combined into a single compound transformation, and the image resampled once. Note that although it is common to create resampled data at the original resolution (e.g., 3x3x3 mm), any resolution might be used.

Frequency domain filtering. Removing frequency components from the time series signal is an effective way to remove noise. Since Biswal (1995), it has been common in FC studies to remove frequencies higher than 0.08Hz (period of 12.5s). Higher frequencies than this are unlikely to contain much if any information about neural activity, given the slowness of the hemodynamic response measured by BOLD (Aguirre et al., 1997). Fox et al. (2005; 2006) also remove frequencies lower than 0.009Hz (period of 111s). Cordes et al. (2001) have demonstrated explicitly that lower frequencies underlie functional connectivity. Note that frequency domain filtering reduces the effective degrees of freedom in the data.

Spatial smoothing. Spatial smoothing almost invariably means smoothing with a 3D Gaussian kernel. At a spatial scale determined by the size of the Gaussian kernel, it improves the signal-to-noise ratio by averaging across voxels that may have some independent noise, but similar signals of interest. The cost of smoothing is a loss of power to detect effects that depend on more spatially localized signals, since these will not benefit from spatial averaging. Typical smoothing kernels for FC studies are 6-9 mm FWHM.

Regressing out patterns of no interest. Various noise signals can be removed from each time series in the brain by regression. This is particularly important for functional connectivity studies in order to remove degenerate sources of inter-region correlations. Fox et al. (2005; 2006) include the six parameters estimating head movement, global signal from the entire brain, signal drawn from a region in the ventricles, and signal drawn from a region in white matter. Fox et al. (2005) have also suggested an advantage to removing the first derivative of the global, ventricle, and white matter signals. The ventricle and white matter regions are typically drawn by hand, although with high-quality normalization, the regions can be done just once for a study.

Selecting seed signal. For practical reasons, we usually can't examine the connectivity for every pairwise combination of voxels – visualizing the data would require a separate full brain statistical map for each voxel. So studies using these methods usually test hypotheses concerning the connectivity of specific regions of interest, determined either anatomically (i.e., by delineating a structure of interest) or functionally (i.e., by delineating a

region of activation from separate data). The seed signal is typically an average of all the voxels in the region of interest, or sometimes a single voxel.

Creating correlation map. Correlation maps are usually created in the obvious way, by calculating a product-moment correlation with the seed signal for each voxel in the brain.

Extensions to this method

Rogers et al. (2005) note that it's possible to examine these correlations iteratively -- for each of the connected regions identified, we may then examine that region's network of connected regions. Hampson et al. (2002) have demonstrated that, given a blocked task performance dataset and a resting dataset, it's possible to go back and forth, using each dataset to generate hypotheses that may be tested in the other (e.g., regions activated in the performance data are hypothesized to be functionally connected in the resting data, and regions functionally connected are hypothesized to show a task-modulated connection in the performing data). Many different such arrangements are possible, and it is obviously important to consider the analyses carefully before new data are collected.

While resting BOLD data are often available, data from activation studies are even more readily available (e.g., from the fMRIDC, fmridc.org). Fair et al. 2007 have demonstrated that extracting rest-period data from blocked BOLD studies can be a viable substitute for pure resting BOLD data. However, they also observe that data from event-related paradigms cannot generally be used in place of resting data, because the contaminating effects of task activity are difficult to filter out.

The signal in one region may be best predicted by the past or future signal in another region, not just the time-locked signal. Correlations at a temporal lag are potentially informative (if not decisively so) as to the causal direction of a functional connection. That is, if BOLD signal in region A is best predicted by the past signal in region B, we may infer that the causal connection between the two is more likely to flow from B to A than the other way around. Analyses of this type fall under the umbrella term “multivariate autoregressive models,” (Harrison et al., 2003) and include Granger causality mapping (Goebel et al., 2003; Roebroeck et al., 2005).

Sun et al. (2004) have proposed an alternative method in which instead of cross-correlation, a measure of coherence is used instead. Coherence, described as the spectral analog of cross-correlation, considers the similarity between the distribution of frequencies present in the two signals, irrespective of phase. Because phase is ignored, the functional linkage between two regions may have an arbitrary temporal offset without affecting the coherence of the signals. Coherence is generally a more inclusive measure than multivariate autoregressive models, and may encompass relationships it would be difficult to characterize intuitively.

Differences in functional connectivity between two psychological states can be identified by modeling each voxel's signal as a function of both the seed region signal and some other variable state, as well as the interaction between the two. For example, we may observe that Wernicke's area and Broca's area are more highly correleated during a language task than during rest. This is a form of psychophysical interaction (PPI) that can be informative (Rogers et al., 2007 give several examples), but must be interpreted with caution. A significant interaction may tell us that the influence of one voxel on the other is stronger during the task, or that the influence of the task on a voxel is influenced by activity in the other voxel.

Software, discussion groups, and data archives

The software components required to carry out functional connectivity analyses are widely available in fMRI packages. At a minimum, the analyses would only require the ability to do typical preprocessing, frequency filtering, time series extraction, and multiple regression. Most if not all fMRI analysis packages can do all of these easily, so researchers tend to use the software with which they are already familiar. More esoteric methods are less well supported, and tend to require some programming on the part of the investigator. There are some resources specific to functional connectivity, which I list here.

A group at Washington University in St. Louis has pioneered the development of functional connectivity (e.g., Raichle et al., 2001; Fox et al., 2005), and they have been developing a web-based analysis/sharing platform called brainscape, which implements the methods described in their articles. The web site promises not only to allow researchers to analyze their own data, but also provides a way for researchers to share resting BOLD data. It is currently in beta testing and is available at:

http://www.brainscape.org/

The SPM web site lists a large collection of extensions. As of this writing, only one such package is described as being devoted to functional connectivity. Developed by Xiao-Wei Song, the REST toolkit provides several methods for examining functional connnectivity, and includes a graphical interface. This package requires MATLAB, and may be downloaded from:

http://resting-fmri.sourceforge.net/

Note that many other methods that fall under the umbrella of functional connectivity are also available in various toolkits. For a list of those associated with SPM (not all are SPM-dependent or even MATLAB-dependent), see:

http://www.fil.ion.ucl.ac.uk/spm/ext/

The Neuroimaging Tools and Informatics Resources Clearinghouse (NITRC) is a relatively new NIH-sponsored entity that promises to unify the community of developers of imaging-related (and especially fMRI-related) tools. The site currently hosts 93 projects, and is growing rapidly. The front page is at:

http://www.nitrc.org/

NITRC also hosts a discussion group devoted to task-independent fluctuations:

http://www.nitrc.org/projects/fluctuations/

The fMRI Data Center (fMRIDC) is an ambitious site devoted to archiving data from published fMRI studies. Submission to the fMRIDC was briefly required for several journals, which bolstered the contents of the archives, and it currently boasts 122 datasets. As of this writing (July 2008), the archive is in transition to a new location and not accepting new submissions. However, older datasets may apparently still be requested:

http://www.fmridc.org/

References

Biswal, B., F. Z. Yetkin, V. M. Haughton and J. S. Hyde. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI.. Magn Reson Med, 34(4):537-41, 1995.

Collins, D. Louis. 3D Model-based segmentation of individual brain structures from magnetic resonance imaging data. McGill University, Department of Biomedical Imaging, 1994.

Cordes, D., V. M. Haughton, K. Arfanakis, G. J. Wendt, P. A. Turski, C. H. Moritz, M. A. Quigley and M. E. Meyerand. Mapping functionally related regions of brain with functional connectivity MR imaging.. AJNR Am J Neuroradiol, 21(9):1636-44, 2000.

Cordes, D., V. M. Haughton, K. Arfanakis, J. D. Carew, P. A. Turski, C. H. Moritz, M. A. Quigley and M. E. Meyerand. Frequencies contributing to functional connectivity in the cerebral cortex in "resting-state" data.. AJNR Am J Neuroradiol, 22(7):1326-33, 2001.

Cordes, D., V. Haughton, J. D. Carew, K. Arfanakis and K. Maravilla. Hierarchical clustering to measure connectivity in fMRI resting-state data.. Magn Reson Imaging, 20(4):305-17, 2002.

Fair, D. A., A. L. Cohen, N. U. Dosenbach, J. A. Church, F. M. Miezin, D. M. Barch, M. E. Raichle, S. E. Petersen and B. L. Schlaggar. The maturing architecture of the brain's default network.. Proc Natl Acad Sci U S A, 105(10):4028-32, 2008.

Fox, M. D., M. Corbetta, A. Z. Snyder, J. L. Vincent and M. E. Raichle. Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems.. Proc Natl Acad Sci U S A, 103(26):10046-51, 2006.

Friston, K. J., L. Harrison and W. Penny. Dynamic causal modelling.. Neuroimage, 19(4):1273-302, 2003.

Fox, M. D. and M. E. Raichle. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging.. Nat Rev Neurosci, 8(9):700-11, 2007.

Fair, D. A., B. L. Schlaggar, A. L. Cohen, F. M. Miezin, N. U. Dosenbach, K. K. Wenger, M. D. Fox, A. Z. Snyder, M. E. Raichle and S. E. Petersen. A method for using blocked and event-related fMRI data to study "resting state" functional connectivity.. Neuroimage, 35(1):396-405, 2007.

Fox, M. D., A. Z. Snyder, J. L. Vincent, M. Corbetta, D. C. Van Essen and M. E. Raichle. The human brain is intrinsically organized into dynamic, anticorrelated functional networks.. Proc Natl Acad Sci U S A, 102(27):9673-8, 2005.

Fox, M. D., A. Z. Snyder, J. L. Vincent and M. E. Raichle. Intrinsic fluctuations within cortical systems account for intertrial variability in human behavior.. Neuron, 56(1):171-84, 2007.

Fox, M. D., A. Z. Snyder, J. M. Zacks and M. E. Raichle. Coherent spontaneous activity accounts for trial-to-trial variability in human evoked brain responses.. Nat Neurosci, 9(1):23-5, 2006.

Goebel, R., A. Roebroeck, D. S. Kim and E. Formisano. Investigating directed cortical interactions in time-resolved fMRI data using vector autoregressive modeling and Granger causality mapping.. Magn Reson Imaging, 21(10):1251-61, 2003.

Hampson, M., N. R. Driesen, P. Skudlarski, J. C. Gore and R. T. Constable. Brain connectivity related to working memory performance.. J Neurosci, 26(51):13338-43, 2006.

Hampson, M., I. R. Olson, H. C. Leung, P. Skudlarski and J. C. Gore. Changes in functional connectivity of human MT/V5 with visual motion input.. Neuroreport, 15(8):1315-9, 2004.

Harrison, L., W. D. Penny and K. Friston. Multivariate autoregressive modeling of fMRI time series.. Neuroimage, 19(4):1477-91, 2003.

Hampson, M., B. S. Peterson, P. Skudlarski, J. C. Gatenby and J. C. Gore. Detection of functional connectivity using temporal correlations in MR images.. Hum Brain Mapp, 15(4):247-62, 2002.

Hampson, M., F. Tokoglu, Z. Sun, R. J. Schafer, P. Skudlarski, J. C. Gore and R. T. Constable. Connectivity-behavior analysis reveals that functional connectivity between left BA39 and Broca's area varies with reading ability.. Neuroimage, 31(2):513-9, 2006.

Roebroeck, A., E. Formisano and R. Goebel. Mapping directed influence over the brain using Granger causality and fMRI.. Neuroimage, 25(1):230-42, 2005.

Rogers, B. P., V. L. Morgan, A. T. Newton and J. C. Gore. Assessing functional connectivity

in the human brain by fMRI.. Magn Reson Imaging, 25(10):1347-57, 2007.
Rogers, B. P., V. L. Morgan, A. T. Newton and J. C. Gore. Comment on "Assessing functional connectivity in the human brain by fMRI".. Magn Reson Imaging, 26(1):146, 2008.
Sun, F. T., L. M. Miller and M. D'Esposito. Measuring interregional functional connectivity using coherence and partial coherence analyses of fMRI data.. Neuroimage, 21(2):647-58, 2004.