Published in final edited form as:

Neuroimaging Clin N Am. 2007 November ; 17(4): 441–viii.

Statistical Approaches to Functional Neuroimaging Data

F DuBois Bowman, Ph.D.¹, Ying Guo, Ph.D.², and Gordana Derado, M.S.³

¹Associate Professor of Biostatistics, The Rollins School of Public Health, Emory University, Atlanta, GA

²Assistant Professor of Biostatistics, The Rollins School of Public Health, Emory University, Atlanta, GA

³Ph. D. Student, Department of Biostatistics, The Rollins School of Public Health, Emory University, Atlanta, GA

Synopsis

The field of statistics makes valuable contributions to functional neuroimaging research by establishing procedures for the design and conduct of neuroimaging experiments and by providing tools for objectively quantifying and measuring the strength of scientific evidence provided by the data. Two common functional neuroimaging research objectives include detecting brain regions that reveal task-related alterations in measured brain activity (*activations*) and identifying highly correlated brain regions that exhibit similar patterns of activity over time (*functional connectivity*). In this article, we highlight various statistical procedures for analyzing data from activation studies and from functional connectivity studies, focusing on functional magnetic resonance imaging (fMRI) and positron emission tomography (PET) data. We also discuss emerging statistical methods for prediction using fMRI and PET data, which stand to increase the translational significance of functional neuroimaging data to clinical practice.

Keywords

Linear model; spatial model; functional connectivity; independent component analysis; clustering; prediction; Bayesian hierarchical model

INTRODUCTION

There are a range of substantive applications for functional neuroimaging methods, as illustrated by subsequent chapters in this volume. For instance, functional neuroimaging has provided us with a better understanding of the neural basis for major psychiatric disorders, neurological disorders, and substance abuse. *Structural* or *anatomical* neuroimaging such as magnetic resonance imaging (MRI) has had a tremendous impact on medical practice for both clinical diagnoses, e.g., for detecting brain lesions or tumors, and for informing associated treatment decisions. Numerous *functional* neuroimaging studies, using functional MRI (fMRI) and cerebral blood flow positron emission tomography (PET), have identified neural processing differences or alterations in functional connectivity between patient groups and

Address correspondence, including proofs and reprints, to: ¹Dr. F. DuBois Bowman, Department of Biostatistics, Emory University, 1518 Clifton Road, N.E. Atlanta, GA 30322, Phone: (404) 712-9643, Fax: (404) 727-1370, E-mail: dbowma3@sph.emory.edu. Address for Co-authors: ²Dr. Ying Guo, Department of Biostatistics, Emory University, 1518 Clifton Road, N.E. Atlanta, GA 30322, tel: (404) 712-8646, fax: (404) 727-1370, yguo2@sph.emory.edu. ³Mrs. Gordana Derado, Department of Biostatistics, Emory University, 1518 Clifton Road, N.E. Atlanta, GA 30322, tel: (404) 727-8210, fax: (404) 727-1370, gderado@sph.emory.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

healthy controls, for instance concerning Alzheimer's disease [1]. While such group differences in neural processing suggest a role for functional neuroimaging in clinical practice, there has been a paucity of research targeting the development of standardized functional neuroimaging tools that aid in clinical decision-making, especially in terms of treatment decisions. Nonetheless, we envisage important clinical contributions for functional neuroimaging and will discuss these possibilities later in the chapter.

The field of statistics has played an integral role in the advancement of functional neuroimaging research and has the potential to vastly increase the translational significance of neuroimaging techniques to clinical practice. Current functional neuroimaging research focuses largely on *activation studies* and on studies of *functional connectivity*. Activation studies generally seek to characterize the neural responses to experimental tasks, which we may visualize as maps of distributed patterns of brain activity. Other common goals in activation studies are to detect differences in patterns of brain activity among various experimental stimuli, among different subgroups of subjects, and between two or more sessions, e.g., before and after treatment. While activation studies focus on localized brain activation, functional connectivity studies aim to determine multiple brain areas that exhibit similar temporal task-related (or resting state) activity profiles. In this chapter, we present statistical methodology for analyzing functional neuroimaging data in activation studies and studies of functional connectivity. Furthermore, we describe statistical techniques that show promise for clinical applications of functional neuroimaging data, enabling the *prediction* of behavior-related neural processing, of the neural response to treatment, and of the clinical symptom response to treatment.

In what follows, we highlight statistical procedures for fMRI and PET data that target activation studies, functional connectivity studies, and predictive techniques because a comprehensive review of statistical methods for functional neuroimaging data is beyond the scope of this chapter. To facilitate our presentation, we discuss the statistical methods in the context of fMRI data from a study of inhibitory control among cocaine-dependent men and PET data from a study of working memory in schizophrenia patients. Also, we briefly summarize general characteristics of fMRI and PET data.

DESCRIPTION OF EXPERIMENTAL DATA

Inhibitory Control in Cocaine Dependent Men

Impairment in the ability to exert inhibitory control (e.g., of drug-seeking behaviors) in spite of adverse consequences, represents a common characteristic of drug addiction [2]. The deficits in response inhibition that cocaine addicts often exhibit may generalize outside of a drug-seeking context. We consider an fMRI study that evaluates the effects of cocaine addiction and treatment-related cocaine abstinence on the neural representation of inhibitory control experimental tasks. The cocaine-dependent subjects all enrolled in an addiction treatment program.

Baseline fMRI scans were collected from 41 men, 24 of whom were classified as cocaine-dependent and 17 as cocaine-free. Scans were obtained while subjects performed a variation of the classic Go/No-Go task, a stop task, and a Stroop task, all of which are designed to assess inhibitory control processes. The study reassessed neural processing related to the inhibitory control tasks four weeks following the initiation of treatment and tracked the time to relapse weekly for each subject for one year. We consider data from the stop task in this chapter.

General Characteristics of fMRI and PET data

Functional neuroimaging using fMRI and PET measures hemodynamic correlates of localized brain activity. Each scan contains brain activity measurements at various spatial locations

called *voxels* that are small cubic regions of roughly a few mm in size. The number of voxels in an image depends on the scan resolution, but as an example, a typical three-dimensional (3-D) scan may consist of 2mm^3 voxels configured in a $79 \times 95 \times 69$ array after initial image preprocessing. We usually collect a series of 3-D scans for each individual, obtained under different experimental conditions. The series of 3-D scans represents a time series of image intensities at each voxel location, as depicted in Figure 1.

PET and fMRI data present numerous challenges to many conventional statistical procedures. The intercommunicating networks of neurons throughout the brain are quite extensive and not completely understood. Complex patterns of spatial correlations arise between localized measures of brain activity and are task-dependent. A major modeling challenge stems from the fact that white matter axonal connections (bundles) gives rise to both short and long-range functional dependencies, rendering spatial modeling procedures that assume decaying correlations with increasing distance inappropriate. The white matter connections are quite extensive and may join cortical areas within or between cortical areas, as well as link cerebral cortex to various subcortical structures. Therefore, distant regions linked by axonal fibers may have the potential to exhibit higher correlations than voxels that lack such connections. Furthermore, the enormous number of voxels may pose computational difficulties for implementing statistical techniques. At a given spatial location, the multiple measurements of brain activity from repeated scans exhibit temporal correlations. These temporal correlations are often driven by functional similarities in measured brain activity at different frequency bands. Temporal correlations may also relate to non-neurophysiological sources such as cardiac and/or respiratory artifacts and scanner drift effects.

PREPROCESSING

Functional neuroimaging data often undergo preprocessing prior to statistical analyses. The initial preprocessing step involves *preparing the functional data for viewing* to check for anomalous images, for example, that are improperly reconstructed, extremely noisy, or not correctly oriented (i.e., flipped or rotated). Scanners sequentially acquire the slices of a single 3-D functional image, resulting in brain oxygenation-level dependent (BOLD) signal samples at different layers of the brain corresponding to slightly different time points (msec apart). It is common to perform *slice timing correction*, using a phase shift, to make all of the slices from a single 3-D image correspond to the same point in time.

Head motion invariably occurs while subjects are in the scanner, despite attempts to prevent it. Therefore, it is necessary to ensure proper spatial alignment of each subject's serial brain scans. Using the method of least squares and a six-parameter rigid-body spatial transformation, we perform *realignment* on each subject's time-series of images to ensure that all the scans correspond to a single image (arbitrarily selected from the series) [3]. The set of realignment parameters can also be incorporated within a general linear model (see below) to adjust for movement differences between subjects. For PET, we also perform *co-registration*, which aligns the functional scans with anatomical MRI scans, enabling visualizations that superimpose PET information on an MRI image. The rigid-body transformation (in 3-D) can be parameterized by three translations and three rotations about the different axes [4].

A challenge confronting multi-subject functional neuroimaging studies involves the issue of combining data from different subjects, given neuroanatomical variation. Following realignment, it is common to perform *spatial normalization*, which attempts to deform or warp each individual's series of scans to a standard template brain [usually the Montreal Neurological Institute (MNI) template]. Spatial normalization allows pooling of data from different subjects, comparisons of measured activity between subjects, and comparisons of functional neuroimaging results between studies.

Another common data processing step is *spatial smoothing* (low-pass filtering), which convolves image volumes, usually with a Gaussian kernel of a specified width. Spatial smoothing usually blurs sharp edges because it suppresses the high frequency signal and enhances the low frequency signal, i.e., increases the signal-to-noise ratio (SNR) of the data. Conceptually, spatial smoothing is appealing because it mitigates the effects of neuroanatomic variation that may persist after spatial normalization. Another advantage of spatial smoothing is that it provides a stronger basis for applying random field theory for thresholding of “activation maps” (see below) by making random field theory assumptions more tenable. Spatial smoothing, however, is not uniformly supported. It may attenuate highly-localized effects (activations) due to the integration of BOLD responses across neighboring regions [5]; inflate correlations between neighboring voxels; and, some argue, is a suboptimal way to increase the SNR [6].

Cerebral hemodynamic responses to brief periods of neural activity are not instantaneous. There is a blurring and a delay of the peak response by several seconds (5–8 sec) as well as a transient increase of blood flow within the first 1,000 msec of neuronal activity. A simple way of capturing this is to assume that the fMRI response depends on the external stimulus by convolution with a hemodynamic response function $h(t)$, as follows: $x(t) = \int_0^{\infty} h(u)s(t-u)du$, where $x(t)$ is the (noise-free) fMRI response at a particular voxel at time t , and the external stimulus is given by $s(t)$ [7]. Numerous alternatives exist for specifying the hemodynamic response function $h(t)$, but a simple and common choice is a gamma function [8]. Some voxels in fMRI time-series data show considerable drift over time caused by physiological noise or scanner artifacts. One can remove drift either by *high-pass filtering* or by introducing low frequency drift terms, such as cosines, polynomials, or splines, into the statistical model. A related artifact that sometimes arises for fMRI data is a global gain, e.g. between runs or sessions. We *scale* the data to eliminate this effect and to give the data an interpretable (relative) scale. In *proportional scaling*, the gain for each image is estimated and each intensity is multiplied by the gain estimate. More often we use the *grand mean scaling*, where each volume is divided by a session specific gain. Macey et.al. [9] propose a new technique based on a voxel-level linear model of the global signal.

STATISTICAL METHODS FOR ACTIVATION STUDIES

The general objective of activation studies is to identify brain locations that are involved in the neural processing associated with tasks that subjects perform while in the scanner and possibly to compare these neural processing traits between tasks or between subgroups of individuals. We present a commonly employed two-stage statistical modeling procedure that is applicable to activation studies. We also give some extensions to the basic framework to handle several of the complex characteristics of functional neuroimaging data such as spatial correlation patterns. We present the modeling procedure for fMRI data, and PET analyses proceed using a similar framework. Figure 2 graphically depicts the general procedures for statistical modeling and inference.

Stage 1 Model

The first stage of the model characterizes distributed brain activity associated with various experimental conditions, separately for each individual. To set notation, consider functional neuroimaging data consisting of N subjects (indexed by $i = 1, \dots, N$), T scans (indexed by $t = 1, \dots, T$), and V voxels (indexed by $v = 1, \dots, V$). We arrange the serial BOLD fMRI responses for subject i in the vector $\mathbf{Y}_i(v) = (Y_{i1}(v), \dots, Y_{iT}(v))'$ (at a given voxel location). This is illustrated by Figure 1, which shows a voxel time series $\mathbf{Y}_i(v)$ from a cocaine addict in the inhibitory control study. The figure depicts the varying brain activity (BOLD responses) across 170 scans

and indicates the onsets of the experimental stimulus, which was an auditory tone signaling to the subject to refrain from responding to the “go” cue.

The first-stage linear model for the vector of BOLD responses (for subject i and voxel v) is given by

$$\mathbf{Y}_i(v) = \mathbf{X}_{iv} \mathbf{B}_i(v) + \mathbf{Z}_{iv} \boldsymbol{\alpha}_i(v) + \mathbf{e}_i(v). \quad (1)$$

The design matrix $\mathbf{X}_{iv}$ contains columns pertaining to the experimental stimuli. In its raw form, the design matrix often contains ones and zeros, where a one indicates that the scan associated with the row in which the one appears was obtained under the experimental stimulus associated with the column in which the one appears. For fMRI data, we convolve the raw design matrix with a hemodynamic response function prior to statistical modeling. The matrix $\mathbf{Z}_{iv}$ contains additional covariates that we control for in our statistical analysis, but that are not of substantive interest. Examples of such variables include global blood flow and signal determined by factors other than stimulus-induced neural processing. The parameters $\mathbf{B}_i(v)$ and $\boldsymbol{\alpha}_i(v)$ represent subject-specific effects corresponding to $\mathbf{X}_{iv}$ and $\mathbf{Z}_{iv}$, respectively. The vector $\mathbf{e}_i(v)$ denotes random errors representing characteristics of the measurement process that are unaccounted for by our statistical model, and we assume that the errors follow a zero-mean multivariate normal distribution $N(0, \sigma_i^2(v) \mathbf{I})$ where $\sigma_i^2(v)$ is the error variance at voxel v .

One notable feature of each individual’s data at a given voxel location is that the serial measurements of brain activity exhibit temporal correlations. The stage 1 model above does not explicitly address these serial correlations, but a few commonly employed strategies exist. First, the use of the two-stage modeling approach indirectly addresses within-subject correlations [10]. Second, one may pre-whiten the fMRI serial responses to remove the temporal correlations among multiple scans on each individual. Pre-whitening usually involves estimating a global correlation parameter from a first-order auto-regressive model [Statistical Parametric Mapping (SPM) 99]; and then using the resulting covariance matrix (based on localized variance estimates) to transform the data, yielding independent (transformed) scans if the assumed covariance model is correct. Third, one may implement temporal smoothing to increase the SNR by convolving the fMRI signal with a reference function and to introduce known (at least partially so) correlations, enabling estimation using the generalized least-squares method. For example, this approach is used with a Gaussian smoothing kernel (low-pass filter) matched to the hemodynamic response function along with a linear high-pass filter to impose an approximately known correlation structure [11].

For PET data, the serial measurements are obtained across much longer intervals than for fMRI, often separated by minutes or hours rather than seconds. Nonetheless, the data may still exhibit correlations between the repeated scans within an individual. The reduced number of scans for most PET studies enables the use of a single level mixed effects model analysis that includes various potential covariance structures to capture the temporal correlations [12].

Stage 2 Model

General Linear Model (GLM)—The stage 2 model combines the subject-specific effect $\mathbf{B}_i(v)$ from (1) to estimate the associated group-level or population parameters. Specifically, we model

$$\mathbf{B}_i(v) = \mathbf{W}_{iv} \boldsymbol{\beta}(v) + \boldsymbol{\varepsilon}_i(v), \quad (2)$$

where $\mathbf{W}_{iv}$ is the second-stage design matrix that often consists of subject-specific characteristics and $\boldsymbol{\beta}(v)$ contains the group-level parameters that represent the effects related to different sessions or the effects associated with various subpopulations such as different

treatment groups. In practice, researchers are often interested in making inferences pertaining to the differences between sessions or groups. One common approach is to define a contrast of the elements in $\mathbf{B}_i(v)$ and then to model the contrast in terms of the corresponding group-level contrast parameter [13]. For example, defining $\theta_i(v) = \mathbf{C}\mathbf{B}_i(v)$ with $\mathbf{C} = [1 \ -1 \ 0 \ \dots \ 0]$, takes the difference between the first and second elements in $\mathbf{B}_i(v)$. Modeling a contrast at the second stage, rather than the entire vector of subject-specific effects from model (1), is a convenient approach because it facilitates computations, but it also has some limitations addressed in the next subsection.

General Linear Multivariate Model (GLMM)—A major drawback of modeling contrasts at stage 2 is that it fails to incorporate correlations that may exist between elements of each individual's vector of effects, e.g., representing multiple memory tasks. In the inhibitory control study, for instance, a subject's baseline and follow-up (post-treatment) effects from stage 1 may exhibit positive correlations, and it is important for the stage 2 model to estimate these correlations and to incorporate them into all subsequent statistical inferences. By correctly accounting for these correlations, we can improve the statistical power of the hypothesis testing. An attractive alternative that overcomes the limitations of the GLM-based contrast modeling approach is to fit a GLMM for the repeated-measures at the second stage. The second-stage GLMM is given by

$$\mathbf{B}(v) = \mathbf{W}\beta(v) + \varepsilon(v), \quad (3)$$

where $\mathbf{B}(v)$ is a matrix with the i th row containing the entire vector of effects for subject i , i.e., the elements of $\mathbf{B}_i(v)$ arranged in a row. $\mathbf{W}$ is the multivariate model design matrix that contains between-subject effects such as subgroup assignment indicators (e.g., patients or control subjects), and the group-level parameters of interest are arrayed in the $q \times p$ matrix $\beta(v)$, in which the rows pertain to between-subject effects and the columns correspond to the multiple effects from the stage 1 model. The GLMM assumes that each row of $\mathbf{B}(v)$, containing all of the effects for a particular subject, has a covariance matrix that characterizes the correlations between the multiple effects. Estimation of the GLMM is fast and computationally efficient, even for a large number of voxels.

Spatial Models—Complex neural networks act as pathways that enable interactive neural processing. The GLM and GLMM approaches perform estimation separately for each voxel, which does not account for the neurophysiological associations that may exist between different brain regions. The common approach of estimating parameters from separate statistical models applied to each voxel implicitly assumes independence between voxels, clearly departing from the underlying neurophysiology. One popular approach applies random field-based methods for thresholding to determine significant activations to help remedy the spatial independence assumption involved during estimation [14,15], but this approach requires restrictive assumptions, sometimes involves alterations to the observed data by smoothing during preprocessing, and relies on (random field) theory that is often not supported by the data [16,17].

There is emerging recognition of the importance of modeling spatial correlations between voxels for both estimation and statistical inferences to address the non-independence between observed brain activity measurements in different voxels. The brain's intricate neurocircuitry subserving interactive neural processing is too extensive for comprehensive modeling in a statistical framework and is not completely understood. However, we may capture prominent features of the spatial networks using statistical models. Some authors attempt to incorporate correlations between the brain activity in a given voxel with the activity from neighboring voxels [18,19,20]. Penny et al. [21] presented a Bayesian model that assumes spatially correlated prior distributions for the regression coefficients from first-order neighbors,

comprised of the physically contiguous voxels (possibly extendable). They also assumed that the covariance matrix is known up to a multiplicative constant.

We have taken an alternative approach to formulate a spatial model that extends beyond (first-order) neighboring voxels [22]. This approach partitions voxels into functionally related neural processing clusters (regions) and applies a spatial simultaneous autoregressive model to capture intra-regional correlations. The spatial model is given by

$$\mathbf{B}_{ig}(v) = \mathbf{W}_{igv} \beta_g(v) + \left(\frac{\rho_g}{V_g - 1} \right) \sum_{v^* \in N_v} (\mathbf{B}_{ig}(v^*) - \mathbf{W}_{igv^*} \beta_g(v^*)) + \varepsilon_{ig}(v), \quad (4)$$

where $\varepsilon_{ig}(v)$ represents a vector of mean-zero random errors for voxel v (which belongs to region g), ρ_g is the spatial dependence parameter for region g (with $-1 < \rho_g < 1$), and N_v is the set of voxels within the same neural processing cluster as voxel v . In this model, the subject-specific effect at voxel v not only depends on the population effect at voxel v but also is affected by the subject's effect at voxels within the same cluster. Key practical advantages of this approach are that it allows the possibility of long range correlations due for example to white-matter axonal connections, and it leads to fast estimation. More generally, other benefits that may result from spatial modeling include more appropriate modeling assumptions, increased precision of voxel-specific brain activity estimators, and increased statistical power. Consider a spatio-temporal model for analyses targeting a specific region of interest (ROI) [23]. This ROI spatial model specifies correlations based on a functionally-defined distance metric so that the magnitudes of correlations decay with decreasing similarity in neural processing, rather than as a function of increasing physical distance between brain locations. Estimation of this model is computationally intensive, limiting its applicability to analyses targeting anatomically-focused hypotheses.

Inference Procedures

Statistical inferences for activation studies typically target the voxel-level (e.g., SPM [24,25]) or a broader regional level. We may seek voxel-level inferences using GLM, GLMM, or spatial modeling methods. Voxel-based procedures allow a high degree of localization for identifying task-related alterations in neural processing. Region-based methods are spatially more coarse, e.g., targeting a defined neuroanatomic structure or region such as a selected Brodmann area [26], and help to correct for *partial volume effects* that are particularly salient when dealing with anatomically abnormal and atrophic cerebral structures.

Voxel-Specific Inferences—Either the method of *least-squares* or *maximum likelihood* is applicable to obtain an estimator of the group-level parameter $\beta(v)$, denoted $\hat{\beta}(v)$, from voxel-level modeling approaches. One can also obtain model-based estimates of the mean activity in each voxel and for each subject, corresponding to different experimental stimuli, using $\mathbf{B}^{\wedge}_i(v) = \mathbf{W}_{iv} \hat{\beta}(v)$. When modeling a contrast of effects at the second stage as described above, we obtain the estimator of the group-level contrast of interest directly from fitting the model. Otherwise, it is common to define a contrast of the group-level effects following estimation, for example $\hat{\theta}(v) = \mathbf{C} \hat{\beta}(v)$. In either case, we target the scientific questions of interest by conducting statistical tests about hypothesized values of the group-level parameters.

For functional neuroimaging analyses, we commonly specify null hypotheses of the form $H_0: \theta(v) = 0$, for example testing the equality of the mean BOLD responses from two experimental tasks. For the cocaine dependence study, we are interested in testing the equality of mean brain activity related to an inhibitory control task following treatment and the mean activity at baseline (prior to treatment). The alternative hypothesis defines a distinct set of values for the parameter (contrast) of interest. For instance, we may specify $H_1: \theta(v) > 0$,

indicating that the mean brain activity (related to response inhibition) following treatment is greater than the mean activity prior to treatment. By specifying multiple contrast vectors $\mathbf{C}$, we can define several parameters that address different study objectives.

Conducting statistical tests about hypotheses usually makes use of three important quantities, namely an estimate of the parameter of interest $\hat{\theta}(v)$, the value of the parameter under the null hypothesis θ_0 , and the standard error of the parameter estimate, denoted $s.e.(\hat{\theta}(v))$. The standard error quantifies the sampling variability about the parameter estimate obtained in the analysis. The next step in conducting tests of hypotheses is to define a test statistic, which combines the three aforementioned quantities. The specific mathematical expression used for the test statistic depends on the statistical model applied to the data, but it has the following basic form

$$t(v) = \frac{\hat{\theta}(v) - \theta_0}{\sqrt{s.e.(\hat{\theta}(v))}}. \quad (5)$$

The test statistic is a random variable that (generally) follows a known probability distribution, e.g., t-distribution, under the null hypothesis. Computing the test statistic at each voxel produces a t-statistic map, and we apply various methods of thresholding to identify voxels producing statistically significant effects or contrasts.

To determine whether the data provide sufficient evidence to reject the null hypothesis, we compare the computed value of the test statistic at each voxel to a specified threshold. Following classical statistical techniques, we utilize a thresholding quantity that represents a percentile from the sampling distribution of the test statistic under the null hypothesis. The selected percentile corresponds to the significance level α specified by the analyst, reflecting the type-I error rate or the probability of rejecting a true null hypothesis. We generally set a small type-I error rate, and obtain the $[100 \times (1-\alpha)]$ percentile. If the observed value of the test statistic exceeds the threshold corresponding to the selected percentile, then the data provide sufficient evidence to reject the null hypothesis in favor of the alternative hypothesis. For more details about estimation and inference for fMRI analysis, we refer to the work of Friston and colleagues [4,24,25], Bullmore et al. [27], and Worsley [28].

One challenge present in the analysis of functional neuroimaging data is that we conduct tests of hypotheses at the voxel level, often resulting in hundreds of thousands of tests. Consequently, we must take additional measures to maintain a reasonable type-I error rate, since it may become inflated due to the large number of tests performed. There is no single measure of type-I error in multiple testing problems. The *familywise error rate* (FWE) is a standard measure, quantifying the chance of any type-I errors occurring. The maximum statistic, e.g., in a t-statistic map, plays a key role in FWE control.

One approach to control the FWE rate is to perform a *Bonferroni*-type correction. The Bonferroni correction simply divides the significance level by the number of tests performed, maintaining a FWE rate at (or below) the nominal significance level. The large number of tests in our context typically requires a corrected significance level that is overly conservative. In practice, we often strike a compromise by specifying a threshold that is more stringent than conventional levels, such as $\alpha=0.05$, but less conservative than the Bonferroni method. The *theory of random fields* provides a framework for an alternative method of adjusting *p*-values that takes into account the non-independence between neighboring voxels. When combined with the GLM, the random field theory (RFT) methods comprise a flexible framework for neuroimaging inference [11]. RFT methods account for dependence in the data, as captured by the *maximum distribution*. These methods approximate the upper tail of the maximum

distribution, the end needed to find small FWE thresholds. For a detailed description of RFT methods, we refer to the work of Worsley and colleagues [15].

A relatively new approach to the multiple comparisons problem is the False Discovery Rate (FDR) [29,30]. The FDR is the expected proportion of erroneous rejections, *among all rejections*. Therefore, FDR controls the *expected proportion* of false positives among voxels exceeding a threshold, unlike Bonferroni or random field methods that control the chance of *any* false positives. Nichols and Hayasaka [31] give a comparative review of FDR and methods for controlling the familywise error rate in functional neuroimaging analyses.

Region of Interest Inferences—In some cases, researchers seek to determine whether there is evidence of task-related changes in brain activity within a particular region, rather than to determine globally where task-related activations occur, leading to ROI analyses. An important aspect of performing ROI analyses involves the selection of the targeted anatomical region, and there are established techniques for doing so [32,33,34]. Following selection of a targeted region, ROI analyses often involve averaging brain activity measurements over all of the voxels in the ROI (for each scan) and then proceeding with statistical analyses with each scan contributing a single measurement of brain activity. There are several advantages of ROI analyses for anatomically-focused hypotheses. First, they substantially reduce the number of voxels, often to a single mean data point as just described. For the mean summary statistic approach, ROI analyses yield brain activity measurements that are less variable than the voxel-level measurements. The reduced number of voxels mitigates the multiple comparisons issue, which in addition to the precision gained by the use of summary statistics, may provide increased sensitivity to detect significant task-related changes in brain activity. Data from an ROI are also simple to visualize, for example allowing plotting of the measured brain activity levels within these regions by experimental condition or by subgroup membership.

The simplicity of the data structure for ROI analyses also enables the application of more sophisticated (and often more computationally-intensive) statistical techniques that may entail assumptions that are better suited for the data. Nieto-Castanon and colleagues [35] presented ROI analysis methods, along with an accompanying SPM toolbox for implementation, based on multivariate analysis methods for fMRI data discussed by Friston et al. [36] and Worsley et al. [37]. We present a statistical model for detecting task-related alterations in measured brain activity that accounts for both spatial correlations between voxels within the ROI and temporal correlations between the multiple scans on each subject [23]. Conducting analyses at the voxel level within the ROI is beneficial for detecting significant activations in a relatively small area within the ROI that may not be detected when averaging data across the entire region. The spatio-temporal ROI model [23] also borrows strength to increase statistical power by incorporating information from voxels exhibiting similar task-related functionality. A major limitation of all ROI analyses is that activations outside an ROI are not detectable, so one should apply ROI analyses only for anatomically-focused hypotheses.

Nonparametric methods—Complexities in fMRI and PET data may support the use of nonparametric methods for finding appropriate thresholds for statistical inferences. Nonparametric methods are more flexible than parametric approaches, since they do not require assumptions regarding underlying probability distributions, and we are not limited to test statistics with known null-distributions. *Permutation testing*, introduced into the functional neuroimaging literature by Holmes et.al. [16], is a well-established method for nonparametric testing that readily accounts for the multiple comparisons problem. Nichols and Holmes [38] provide details on conducting permutation tests. Bullmore et al. [39] present procedures, called *wavestrapping*, for resampling wavelet coefficients of 1-D time series, 2- or 3-D spatial maps, and 4-D spatiotemporal processes, and they propose wavelet-based methods for multiple hypothesis testing.

STATISTICAL METHODS FOR FUNCTIONAL CONNECTIVITY

Functional connectivity has been defined as the temporal correlation between spatially remote neurophysiological events [40]. Various statistical methods exist for examining functional connectivity, ranging from descriptive approaches to inferential procedures, and we highlight several methods in this section.

Clustering

Cluster analysis is a data-driven method that can assist in identifying voxels exhibiting similar patterns of task-related brain activity. Neuroimaging applications currently utilize several clustering algorithms, including the K -means approach [41,42,43,44], fuzzy clustering [45, 46,47,48], hierarchical clustering methods [49,50,51,42,52], a hybrid hierarchical K -means approach [53], and dynamical cluster analysis (DCA) [54]. The performances of clustering methods are largely influenced by characteristics of the data. Many algorithms are capable of detecting well-separated clusters. In other cases, however, the performance of a given method depends on certain cluster characteristics such as compactness, shape, size, and heterogeneity. As a result of our simulation study comparing various clustering algorithms [49], we advocate the use of Ward's and the beta-flexible hierarchical clustering methods as well as K -means and fuzzy clustering procedures when the number of clusters is known with a reasonable degree of certainty. A multiple classification approach also may permit likelihood-based comparisons across clustering solutions produced by various algorithms, allowing the analyst to detect the algorithm that provides the *best* clustering solution for the data [50].

We often apply a clustering algorithm to a summary statistic of the scans, such as the effects from a linear regression model that represent stimulus-specific averages across subjects. The primary objective of neuroimaging clustering is to identify collections of voxels (clusters) that exhibit similar brain activity patterns and that reveal distinct neural response patterns between clusters. One important feature of the detected clusters is that they may consist of non-contiguous voxels, offering the potential of identifying associations that may exist between anatomically distant voxels, as illustrated in Figure 3. We provide a brief exposition of hierarchical and K -means procedures. We use G to denote the total number of clusters, despite the standard use of the name K -means (suggesting K as the total number of clusters).

Clustering algorithms use measures of distance to determine dissimilarity between voxels (or clusters). The K -means algorithm begins by specifying the number of clusters G , randomly initializing the voxels to the G groups, and calculating the group centroids (means). The algorithm iteratively reassigns voxels to the cluster with the nearest centroid, based on the distance calculations representing the dissimilarity in brain function, and recalculates the cluster means. The algorithm ceases when no reassignments occur. The K -means algorithm is common in neuroimaging applications, in part because the computations are fast and efficient. However, a disadvantage of the K -means approach is that it requires prior specification of the number of clusters (G), and often there is no scientific basis for setting G in functional neuroimaging applications.

Agglomerative hierarchical clustering begins with each voxel representing a separate cluster and proceeds with successive merges until all clusters unite. Beginning with V clusters, we calculate a $V \times V$ distance matrix, with each element representing the dissimilarity in brain function between the voxels corresponding to the associated row and column of the matrix. Hierarchical clustering then proceeds as follows: (1) Identify the smallest distance and combine the corresponding clusters, reducing the total number of clusters by one after each iteration; (2) compute an updated distance matrix by removing rows and columns associated with the two merged clusters and adding a row and column containing updated distances between the new cluster and all other clusters; (3) repeat steps until all clusters unite. Specific hierarchical

clustering algorithms use different updating functions to measure the distances between clusters (For review, see [49]).

Independent Component Analysis

Independent component analysis (ICA) is becoming increasingly popular for analyzing functional neuroimaging data [55,56,57]. Neuroimaging signals measuring brain activity consist of different sources of variability, such as machine artifacts, physiological pulsation, and the true brain activity changes induced by experimental conditions. Most conventional analysis tools such as GLM make specific assumptions about the relationship between the measured brain activity and the experimental stimuli. These methods pre-specify a design matrix that, combined with a vector of parameters, represents the expected neural responses associated with various experimental stimuli. The validity of the inferred spatial activation patterns depends on the validity of the assumed spatio-temporal signal characteristics. Structured noise from sources other than the experimental conditions may bias the parameter estimates, inflate the residual errors, and reduce the statistical significance. As an alternative, ICA does not require any prior knowledge regarding the spatio-temporal structure underlying the measured brain activity. ICA identifies independently distributed spatial patterns that are associated with different time-courses and hence offers a useful exploratory technique to investigate functional connectivity. In particular, ICA has been very useful for studying functional connectivity in resting state data, where there are no clear task-related activations [56,58].

Single subject ICA—The basic goal of ICA is to express the observed 3-D brain images as linear combinations of statistically independent latent component signals. There are two major types of ICA models. The *classical ICA* procedure decomposes the observed data into the product of a matrix representing statistically independent spatial source signals and a matrix containing the time series associated with these source signals. A second approach, referred to as *noisy ICA* or *probabilistic ICA*, assumes the observed data is composed not only of the ICA decomposition but also of noise [59,60]. Our following introduction focuses on probabilistic ICA and includes some discussion about classical ICA.

The probabilistic ICA model is formulated as a generative linear latent model,

$$\mathbf{Y}_i(\mathbf{v}) = \mathbf{A} \mathbf{s}_i(\mathbf{v}) + \varepsilon_i(\mathbf{v}), \quad (6)$$

where $\mathbf{Y}(\mathbf{v})$ is the time series vector observed at voxel $\mathbf{v}$, $\mathbf{s}_i(\mathbf{v})$ is an unknown vector of source signals associated with voxel $\mathbf{v}$, and $\mathbf{A}$ is an unknown matrix called the *mixing matrix* because it mixes the independent source signals to generate the observed data. Each element in $\mathbf{s}_i(\mathbf{v})$ represents the presence and the “amplitude” of a particular source in the observed data $\mathbf{Y}_i(\mathbf{v})$. The number of sources, i.e., the number of columns in $\mathbf{A}$, is smaller than the number of scans or time points. In model (6), each column of $\mathbf{A}$ represents a latent time series associated with a specific spatial source. In probabilistic ICA, the model also includes an error term $\varepsilon_i(\mathbf{v})$, representing variability in the data that is not explained by the sources. The source $\mathbf{s}_i(\mathbf{v})$ is non-Gaussian, while the error is assumed to follow a Gaussian distribution. Figure 4 depicts (a) the thresholded map for a spatial source signal along with its (b) associated time series for a subject in the cocaine-dependence study.

The primary distinction between probabilistic and classical ICA is that the classical approach does not include the Gaussian noise component. Therefore, classical ICA attempts to characterize the data completely by the source signals, which are combined by the mixing matrix. The classical noise-free model precludes any test of significance and thresholding techniques for the source estimates. Issues sometimes arise with classical ICA since any valid spatial variation, such as a slight difference between the left and right hemisphere or spatial

difference in the background noise level, may be treated as spatial source signals. Consequently, clusters of voxels that show similar temporal patterns may split into different spatial maps. The ICA model in (6) is similar to the standard GLM, where the mixing matrix $\mathbf{A}$ corresponds to the design matrix and the source signal $s_i(v)$ corresponds to the unknown parameter vector. However, unlike GLM, ICA does not require pre-specification of the mixing matrix prior to model fitting, but instead estimates the matrix from the observed data, and the spatial source signals in ICA have the additional constraint of being statistically independent. For details regarding the estimation of probabilistic ICA and other ICA models, readers can refer to Beckman and Smith [60].

Group ICA—Scientific objectives in functional neuroimaging studies often target conclusions that generalize to groups of individuals, so we now discuss methods for performing ICA on multi-subject functional neuroimaging data. We present two approaches for performing ICA on multi-subject imaging data. The first approach involves concatenating the subjects' data along the time dimension [57], and then applying a standard 2-D ICA to derive a set of “group” independent components and aggregated mixing matrix. The aggregated mixing and unmixing matrices both partition into sub-matrices corresponding to each subject. We can reconstruct the ICA spatial maps for individual subjects by multiplying a subject's portion of the unmixing matrix with the observed data from the corresponding subject. We can also reconstruct the subject-specific latent time series by partitioning the aggregated mixing matrix.

A second multi-subject approach is the tensor probabilistic ICA [60], extending the single subject probabilistic ICA in (6). While single-subject probabilistic ICA represents the data by different spatial processes along with their associated temporal dynamics, the tensor probabilistic ICA for multiple subjects adds an additional subject dimension to the spatial and temporal domains. Tensor probabilistic ICA decomposes imaging data from multiple subjects into a sum of products of spatial source signals, corresponding latent temporal dynamics, and subject-specific loadings associated with each of the spatio-temporal processes. Some other works in group ICA include Lukic et al. [67] and Svensen et al. [62], and available software tools to implement ICA include Melodic in FSL [Oxford Centre for Functional Magnetic Resonance Imaging of the Brain (FMRIB) Software Library] and GIFT (Group ICA of fMRI Toolbox).

Bayesian Hierarchical Modeling

One may also implement parametric model-based analyses to quantify functional connectivity and provide statistical evidence regarding the likelihood of functional connectivity between brain regions. Patel et al. [63] presented one such approach, formulating a Bayesian hierarchical model. For every scan, the method initially classifies each voxel as either active or inactive. The model utilizes data that reflect joint activation between any pair of voxels, say v_1 and v_2 . Specifically, combining across all subjects and all scans, z_1 is the number of times that both voxels v_1 and v_2 are active, z_2 is the number of times that voxel v_1 is active but v_2 is inactive, z_3 is the number of times that voxel v_1 is inactive but v_2 is active, and z_4 is the number of times that both voxels v_1 and v_2 are inactive. Our Bayesian model for the joint activation of each pair of brain voxels assumes a multinomial likelihood for $z_1, \dots, z_4$ with corresponding parameters $\theta_1, \dots, \theta_4$. By assuming an appropriate prior distribution for the parameters $\theta_1, \dots, \theta_4$, we can derive an explicit posterior distribution for $\theta_1, \dots, \theta_4$, conditional on the data.

Patel and colleagues [63] developed a measure of functional connectivity that is based on the relative difference between the marginal probability that voxel v_1 is active and the probability that voxel v_1 is active conditional on elevated activity in voxel v_2 . Larger differences between these conditional and marginal probabilities reflect voxel pairs exhibiting stronger functional connections. For a pair of functionally connected voxels, Patel et al. [63] also defined a

measure, called *ascendancy*, which determines a hierarchical relationship between the functionally connected voxels. Voxel v_1 is ascendant to v_2 when the marginal activation probability of v_1 is larger than that of v_2 . The posterior probability distribution derived from the Bayesian hierarchical model yields a measure of the degree of functional connectivity between a voxel pair and the degree of ascendancy of one voxel relative to the other.

Spatio-temporal modeling

Our spatial model [22] establishes a unified framework to address objectives for activation studies and provide information on functional connectivity in the brain. The preliminary data-driven cluster analysis defines neural processing clusters with similar patterns of activity, but generally does not quantify or test the degree of association between the intracluster voxels. The cluster-specific model-based estimates of the spatial dependence parameters, ρ_g from model (4), determine the extent of functional connectivity between the voxels within each cluster. In addition, one may conduct classical hypothesis testing (using a Wald or likelihood ratio test) for each neural processing cluster to evaluate the null hypothesis of no functional connectivity between voxels within the cluster versus the one-sided alternative of positive functional connectivity.

Nonparametric Wavelet-Based Methods

Correlations between voxels may arise due to neurophysiological influences as well as to background or non-neurophysiological sources, such those induced by the scanner and by image preprocessing. When evaluating the functional connectivity between two brain regions, it is important to account for the background spatial correlation inherent in neuroimaging data. Breakspear and colleagues [64] developed a wavelet-based nonparametric technique for testing the null hypothesis that the correlations between two selected brain regions are typical of the data set and not unique to the regions of interest. This was achieved through spatio-temporal resampling of the data in the wavelet domain. Patel and colleagues [65] extended this work to a 4-D wavelet-based nonparametric approach for determining whether the functional connectivity observed in an experiment is significantly greater than the background correlation. Specifically, they presented a spatio-temporal wavelet packet resampling method that generates surrogate data that preserves only the average background correlation within an axial slice, across axial slices, and through each voxel time series, while excluding the specific correlations due to true functional relationships.

Extensions Causal Associations

Structural Equation Modeling—For functional neuroimaging data, *effective connectivity* addresses the influence that one neuronal system exerts upon others [66] and how this varies with the experimental context. *Structural equation modeling* (SEM) is a well-established statistical technique that has applications to effective connectivity analyses. SEM focuses on the covariance structure that reflects associations between the variables. Parameter estimation in SEM minimizes differences between the observed *covariances* and those implied by a path (or structural) model. The parameters of the SEM represent the connection strengths between the brain activity measurements in different regions and correspond to measures of effective connectivity. Applying SEM to PET data corresponding to object and spatial vision tasks, McIntosh and Gonzales-Lima [67] demonstrated the dissociation between ventral and dorsal visual pathways. Grafton et al. [68] used SEM to assess the effect of pallidotomy on effective connectivity in the motor system of Parkinson patients. SEM also has been used to describe changes in effective connectivity associated with memory tasks [69,70]. Several statistical software packages implement SEM including, but not limited to, SAS and SPSS.

Dynamic Causal Modeling—Dynamic causal modeling (DCM) is a general method to estimate effective connectivity from neuroimaging data in a Bayesian framework [71,72]. The aim of DCM is to estimate parameters at the neuronal level such that the modeled BOLD signals are maximally similar to the observed BOLD signals. DCM regards the brain as a deterministic nonlinear dynamic system that receives inputs and that produces outputs [72]. DCM parameterizes *effective connectivity* in terms of coupling, representing the influence of one brain region on another. DCM attempts to estimate coupling parameters by measuring the responses to perturbations in the specified system [72]. Additional details on DCM may be found in articles by Mechelli and colleagues [73], Penny et al. [74], and Friston [75].

Partial least-squares—*Partial Least Squares (PLS)* regression generalizes features from principal component analysis (PCA) and multiple regression analysis. PLS regression looks for orthogonal factors (called *latent variables*) that perform a decomposition of both neural responses $\mathbf{Y}$ and experimental variables (predictors) $\mathbf{X}$ simultaneously, such that these factors explain as much as possible of the covariance between $\mathbf{X}$ and $\mathbf{Y}$. The term ‘*partial*’ refers to the computation of the optimal least-squares fit to a *part* of a correlation or covariance matrix [76], specifically the *cross-block* correlations between some set of experimental variables and the associated neural responses. PLS measures cross-block correlations and creates a new set of variables optimized for maximum covariance using the fewest possible dimensions [69]. It is ideally suited for data that have highly correlated dependent measures, such as neuroimaging data [77-79].

STATISTICAL PREDICTION TECHNIQUES

fMRI and PET play important roles in defining the neural basis of illness and of risk factors for diseases such as major psychiatric disorders [80,81,82,83]. To increase the translational significance of functional neuroimaging techniques to clinical practice, an important area of research involves methods for predicting disease progression and treatment response. Numerous studies have linked observed neural processing characteristics from fMRI and PET with disease development or with clinical responses to treatment [84,85,86,87,88]. For example, several authors have established associations between treatment response and pre- to post-treatment changes in measured brain activity [89,90,91,92,93,94,95].

The potential clinical insights gained from evaluating both baseline and post-treatment scans are offset in practice due to the unavailability of post-treatment scans at the time that a clinician makes treatment decisions for a particular patient. This pragmatic shortcoming signals the utility of developing a statistical framework to predict treatment-related brain responses, which can then be evaluated along with the baseline scans and a patient’s history to inform clinical decision making. Guo and colleagues [96] proposed a predictive statistical model that uses a patient’s pre-treatment scans, coupled with relevant patient characteristics, to predict the patient’s brain activity following a specified treatment regimen. The predicted post-treatment brain activity maps, along with the observed baseline brain scans, provided a physician with objective and clinically relevant information that he or she may incorporate into the treatment selection process. The predictive model of Guo represents an important first step in attempting to aid treatment decisions by using functional neuroimaging data and provides a foundation upon which future research, including that on predicting clinical symptom responses, can build.

Guo et al. [96] established their predictive framework using a Bayesian hierarchical model constructed on the pre- and post-treatment scans of designated training data. The first level of the hierarchy models within-subject temporal activation effects, while the second level modeled the subject-specific effects in terms of population parameters. The predicted post-treatment maps follow from the conditional probability distribution of the post-treatment maps given the pre-treatment maps. They illustrated their method using PET data from a study of

working memory in patients with schizophrenia. In this application, they predicted each patient's distributed patterns of brain activity corresponding to varying working memory load levels, following a 12-week treatment regimen with either risperidone (2-6 mg/day) or olanzapine (10-20 mg/day). The predicted post-treatment maps from the Bayesian hierarchical model were quite accurate when compared to the observed post-treatment maps (Figure 5). To implement the predictive framework in practice, a clinician obtains data from a patient prior to initiating a new treatment and applies the predictive algorithm to the baseline data to predict patterns of post-treatment brain activity under various treatment alternatives. These predicted maps can then contribute to clinical decision-making regarding treatment options. One benefit of the Bayesian framework is that it allows updating of the predictive model after obtaining both pre- and post-treatment scans on more individuals.

SOFTWARE

The massive amounts of data collected in functional neuroimaging studies pose challenges for implementing statistical analyses. Fortunately, substantial advances have been made in the development of software for processing and for statistical analyses of neuroimaging data. Table 1 shows software tools that are helpful for implementing several of the statistical techniques described in this article. Some of the software tools provide essentially comprehensive environments for processing and analyzing functional neuroimaging data, while others are more specialized, implementing very specific analyses. Popular software tools in the neuroimaging community include SPM, which is based on the MATLAB software packaged developed by MathWorks; Analysis of Functional NeuroImages (AFNI), which is based on the C programming language; and FSL. Numerous other tools exist for implementing statistical analyses of functional neuroimaging data, but we do not attempt to provide an exhaustive list here.

DISCUSSION

Neuroimaging statistics is an emerging area that has helped to foster the widespread use of functional neuroimaging technology for research and clinical applications. Statistics plays a pivotal role in functional neuroimaging research in its interplay with other fields, such as neuroscience and imaging physics. Using probabilistic arguments and modeling techniques, statistics makes valuable contributions concerning methods for the design and conduct of neuroimaging experiments and tools for objectively quantifying and measuring the strength of scientific evidence provided by the data. The massive data structures and complex patterns of correlations pose challenges for many conventional statistical methods. These (and other) complications give rise to the development of new statistical methods and the adaptation of existing approaches for applications to functional neuroimaging data. The advancements in neuroimaging statistical methodologies provide effective means for processing, exploring, analyzing, and visualizing fMRI and PET data.

Applications of functional neuroimaging statistical methods have provided us with a better understanding of the neural basis of cognitions, emotions, behaviors, and the pathophysiology of psychiatric and neurologic disorders. We have developed methods that enable the conduct and analyses of activation studies that attempt to identify brain regions revealing task-related changes in measured activity. In addition, we have developed methods to analyze data from functional connectivity studies that seek to determine which brain regions show similar patterns of activation when performing an experimental task. Statistical prediction for functional neuroimaging data is a nascent area that stands to have important clinical applications. The recent development of predictive models for fMRI and PET data demonstrate the feasibility of providing accurate predictions of brain activity maps following treatment. We ultimately envisage the development of statistical algorithms that use functional neuroimaging markers

to predict a patient's clinical response (e.g., depression response) to treatment. Methodology for statistical prediction based on functional neuroimaging data represents an important area for future research, and preliminary work in this area provides a promising outlook for the potential utility of functional neuroimaging data in a clinical setting.

Acknowledgements

The authors thank Dr. Clinton Kilts and Mr. Timothy Ely for their input on the fMRI and PET data examples. We also thank Ms. Angela Thomas for her assistance with manuscript preparation.

Funded by: National Institutes of Health; Grant number: K25-MH65473 and R01 MH079251-01A1

LITERATURE CITED

1. Mosconi L, De Santi S, Rusinek H, et al. Magnetic resonance and PET studies in the early diagnosis of Alzheimer's disease. *Expert Rev Neurother* 2004;4(5):831–49. [PubMed: 15853510]
2. Kalivas PW, Volkow ND. The neural basis of addiction: A pathology of motivation and choice. *American Journal of Psychiatry* 2005;162(8):1403–13. [PubMed: 16055761]
3. Friston KJ, Williams S, Howard R, et al. Movement-related effects in fMRI time-series. *Magnetic Resonance in Medicine* 1996;35:346–55. [PubMed: 8699946]
4. Friston, KJ. Introduction: experimental design and statistical parametric mapping. In: Frackowiak, et al., editors. *Human Brain Function*. 2. 2003.
5. Wang J, Wang Z, Geoffrey K, et al. To smooth or not to smooth? ROC analysis of perfusion fMRI data. *Magnetic Resonance Imaging* 2005;23(1):75–81. [PubMed: 15733791]
6. Buxton, RB. Principles and techniques. Cambridge: Cambridge University Press; 2002. Introduction to functional magnetic resonance imaging.
7. Worsley KJ, Liao CH, Aston J, et al. A general statistical analysis for fMRI data. *NeuroImage* 2002;15:1–15. [PubMed: 11771969]
8. Lange N, Zeger SL. Non-linear fourier time series analysis for human brain mapping by functional magnetic resonance imaging (with discussion). *Applied Statistics, Journal of the Royal Statistical Society, Series C* 1997;46(1):1–29.
9. Macey PM, Macey KE, Kumar R, et al. A method for removal of global effects from fMRI time series. *NeuroImage* 2004;22(1):360–66. [PubMed: 15110027]
10. Beckmann CF, Jenkinson M, Smith SM. General multilevel linear modeling for group analysis in FMRI. *NeuroImage* 2003;20:1052–63. [PubMed: 14568475]
11. Worsley KJ, Friston KJ. Analysis of fMRI time-series -Revisited again. *NeuroImage* 1995;2:173–81. [PubMed: 9343600]
12. Bowman FD, Kilts C. Modeling intra-subject correlation among repeated scans in positron emission tomography (PET) neuroimaging data. *Human Brain Mapping* 2003;20(2):59–70. [PubMed: 14505332]
13. Penny, W.; Holmes, A. Random-Effects Analysis (Chapter 42). Frackowiak, RSJ.; Friston, KJ.; Frith, CD.; Dolan, RJ., editors. Elsevier Academic Press; New York: 2004. p. 843-50.
14. Worsley KJ. Local maxima and the expected euler characteristic of excursion sets of, F, and t fields. *Advances in Applied Probability* 1994;26:13–42.
15. Worsley KJ, Evans AC, Marrett S, et al. A three-dimensional statistical analysis for rCBF activation studies in the human brain. *Journal of Cerebral Blood Flow and Metabolism* 1992;12:900–18. [PubMed: 1400644]
16. Holmes AP, Blair RC, Watson JDG, et al. Non-parametric analysis of statistic images from functional mapping experiments. *Journal of Cerebral Blood Flow and Metabolism* 1996;16:7–22. [PubMed: 8530558]
17. Raz J, Turetsky B. Wavelet ANOVA and fMRI [3813]. *Proceedings of SPIE: Wavelet Applications in Signal and Image Processing VII* 1999:561–70.
18. Katanoda K, Matsuda Y, Sugishita M. A Spatio-temporal regression model for the analysis of functional MRI data. *NeuroImage* 2002;17:1415–28. [PubMed: 12414281]

19. Gössl C, Auer DP, Fahrmeir L. Bayesian spatiotemporal inference in functional magnetic resonance imaging. *Biometrics* 2001;57:554–62. [PubMed: 11414583]
20. Woolrich MW, Jenkinson M, Brady JM, et al. Fully Bayesian spatio-temporal modeling of fMRI data. *IEEE Trans Med Imaging* 2004b;23:213–31. [PubMed: 14964566]
21. Penny WD, Trujillo-Barreto NJ, Friston KJ. Bayesian fMRI time series analysis with spatial priors. *NeuroImage* 2005;24:350–62. [PubMed: 15627578]
22. Bowman FD. Spatiotemporal modeling of localized brain activity. *Biostatistics* 2005;6(4):558–75. [PubMed: 15843592]
23. Bowman FD. Spatio-temporal models for region of interest analyses of functional neuroimaging data. *Journal of the American Statistical Association*. 2007accepted
24. Friston KJ, Jezzard P, Turner R. Analysis of functional MRI time series. *Human Brain Mapping* 1994;1:153–71.
25. Friston KJ, Holmes AP, Worsley KJ, et al. Statistical parametric maps in functional imaging: A general linear approach. *Human Brain Mapping* 1994a;2:189–210.
26. Brodmann K. Vergleichende Lokalisationslehre der Großhirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues. Leipzig: Barth JA. 1909
27. Bullmore E, Brammer M, Williams SC, et al. Statistical methods of estimation and inference for functional MR image analysis. *Magnetic Resonance in Medicine* 1996;35(2):261–77. [PubMed: 8622592]
28. Worsley KJ. An overview of some new developments in the statistical analysis of PET and fMRI data. *Human Brain Mapping* 1997;5(4):254–58.
29. Benjamini Y, Hochberg Y. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J Roy Stat Soc Ser B* 1995;57:289–300.
30. Genovese CR, Lazar NA, Nichols T. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage* 2002;15(4):870–8. [PubMed: 11906227]
31. Nichols T, Hayasaka S. Controlling the family wise error rate in functional neuroimaging: a comparative review. *Statistical Methods in Medical Research* 2003;12:419–46. [PubMed: 14599004]
32. Maldjian JA, Laurienti PJ, Kraft RA, et al. An automated method for neuroanatomic and cytoarchitectonic atlas-based interrogation of fMRI data sets. *NeuroImage* 2003;19:1233–39. [PubMed: 12880848]
33. Tzourio-Mazoyer N, Landeau B, Papathanassiou D, et al. Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *NeuroImage* 2002;15:273–89. [PubMed: 11771995]
34. Hammers A, Koepp MJ, Free SL, et al. Implementation and application of a brain template for multiple volumes of interest. *Human Brain Mapping* 2002;15:165–74. [PubMed: 11835607]
35. Nieto-Castanon A, Ghosh SS, Tourville JA, et al. Region of interest based analysis of functional imaging data. *NeuroImage* 2003;19:1303–16. [PubMed: 12948689]
36. Friston KJ, Frith CD, Frackowiak RS, et al. Characterizing dynamic brain responses with fMRI: a multivariate approach. *Neuroimage* 1995;2(2):166–72. [PubMed: 9343599]
37. Worsley KJ, Poline JB, Friston KJ, Evans AC. Characterizing the response of PET and fMRI data using multivariate linear models. *Neuroimage* 1997;6(4):305–319. [PubMed: 9417973]
38. Nichols TE, Holmes AP. Nonparametric permutation tests for functional neuroimaging: a primer with examples. *Human Brain Mapping* 2001;15:1–25. [PubMed: 11747097]
39. Bullmore E, Fadili M, Breakspear M, et al. Wavelets and statistical analysis of functional magnetic resonance images of the human brain. *Statistical Methods in Medical Research* 2003;12:375–99. [PubMed: 14599002]
40. Friston KJ, Frith C, Liddle P, et al. Functional connectivity: The principal component analysis of large data sets. *J Cereb Blood Flow Metab* 1993;13:5–14. [PubMed: 8417010]
41. Balslev D, Nielsen FA, Frutiger SA, et al. Cluster analysis of activity-time series in motor learning. *Human Brain Mapping* 2002;15:135–45. [PubMed: 11835604]
42. Goutte C, Toft P, Rostrup E, et al. On clustering fMRI time series. *NeuroImage* 1999;9:298–310. [PubMed: 10075900]

43. Goutte C, Nielsen FÅ, Liptrot MG, et al. Feature-space clustering for fMRI meta-analysis. *Human Brain Mapping* 2001;13(3):165–83. [PubMed: 11376501]
44. MacQueen, J. Some methods for classification and analysis of multivariate observations. In: Le Cam, LM.; Neyman, J., editors. *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*. Berkeley CA: University of California Press; 1967. p. 281-97.
45. Baumgartner R, Ryner L, Richter W, et al. Comparison of two exploratory data analysis methods for fMRI: fuzzy clustering vs. principal component analysis. *Magnetic Resonance Imaging* 2000;18:89–94. [PubMed: 10642106]
46. Fadili MJ, Ruan S, Bloyet D, et al. A multistep unsupervised fuzzy clustering analysis of fMRI time series. *Human Brain Mapping* 2000;10:160–78. [PubMed: 10949054]
47. Fadili MJ, Ruan S, Bloyet D, et al. On the number of clusters and the fuzziness index for unsupervised FCA application to BOLD fMRI time series. *Med Image Anal* 2001;5:55–67. [PubMed: 11231177]
48. Sommer, FT.; Wichert, A. *Exploratory analysis and data modeling in functional neuroimaging*. Cambridge, MA: The MIT Press; 2003. p. 17-42.
49. Bowman FD, Patel R, Lu C. Methods for detecting functional classifications in neuroimaging data. *Human Brain Mapping* 2004;23(2):109–19. [PubMed: 15340933]
50. Bowman FD, Patel R. Identifying spatial relationships in neural processing using a multiple classification approach. *NeuroImage* 2004;23:260–68. [PubMed: 15325373]
51. Cordes D, Haughton V, Carew JD, et al. Hierarchical clustering to measure connectivity in fMRI resting-state data. *Magnetic Resonance Imaging* 2002;20:305–17. [PubMed: 12165349]
52. Stanberry L, Nandy R, Cordes D. Cluster analysis of fMRI data using dendrogram sharpening. *Hum Brain Mapp* 2003;20:201–19. [PubMed: 14673804]
53. Filzmoser P, Baumgartner R, Moser E. A hierarchical clustering method for analyzing functional MR images. *Magn Reson Imaging* 1999;17:817–26. [PubMed: 10402588]
54. Baune A, Sommer FT, Erb M, et al. Dynamical cluster analysis of cortical fMRI activation. *NeuroImage* 1999;9:477–89. [PubMed: 10329287]
55. McKeown MJ, Makeig S, Brown GG, et al. Analysis of fMRI data by blind separation into independent spatial components. *Human Brain Mapping* 1998;6(3):160–88. [PubMed: 9673671]
56. Beckmann CF, DeLuca M, Devlin JT, et al. Investigations into resting-state connectivity using independent component analysis. *Phil Trans R Soc B* 2005a;360:1001–13. [PubMed: 16087444]
57. Calhoun V, Adali T, Pearlson G, et al. A method for making group inferences from functional MRI data using independent component analysis. *Human Brain Mapping* 2001;14:140–51. [PubMed: 11559959]
58. Damoiseaux JS, Rombouts SA, Barkhof F, et al. Consistent resting-state networks across healthy subjects. *Proc Natl Acad Sci U S A* 2006;103(37):13848–53. [PubMed: 16945915]
59. Hyvärinen, A.; Karhunen, J.; Oja, E. *Independent component analysis*. New York: Wiley; 2001. p. 293-304.
60. Beckmann CF, Smith SM. Probabilistic independent component analysis for functional magnetic resonance imaging. *IEEE Trans on Medical Imaging* 2004;23(2):137–52.
61. Lukic AS, Wernick MN, Hansen LK, et al. A spatially robust ICA algorithm for multiple fMRI data sets. *IEEE International Symposium on Biomedical Imaging Proceedings* 2002:839–42.
62. Svensén M, Kruggel F, Benali H. ICA of fMRI group study data. *NeuroImage* 2002;16:551–63. [PubMed: 12169242]
63. Patel R, Bowman FD, Rilling JK. A Bayesian approach to determining connectivity of the human brain. *Human Brain Mapping* 2006a;27:267–76. [PubMed: 16092131]
64. Breakspear M, Brammer MJ, Bullmore ET, et al. Spatiotemporal wavelet resampling for functional neuroimaging data. *Human Brain Mapping* 2004;23(1):1–25. [PubMed: 15281138]
65. Patel R, Van De Ville D, Bowman FD. Determining significant connectivity by 4D spatiotemporal wavelet packet resampling of functional neuroimaging data. *NeuroImage* 2006;31:1142–55. [PubMed: 16546405]
66. Friston KJ. Functional and effective connectivity in neuroimaging: A synthesis. *Human Brain Mapping* 1994b;2(12):56–78.

67. McIntosh AR, Gonzales-Lima F. Structural equation modelling and its application to network analysis in functional brain imaging. *Human Brain Mapping* 1994;2:2–22.
68. Grafton ST, Sutton J, Couldwell W, et al. Network analysis of motor system connectivity in Parkinson's disease: modulation of thalamocortical interactions after pallidotomy. *Human Brain Mapping* 1994;2:45–55.
69. McIntosh AR, Bookstein FL, Haxby JV, et al. Spatial pattern analysis of functional brain images using partial least squares. *Neuroimage* 1996;3:143–57. [PubMed: 9345485]
70. Nyberg L, McIntosh AR, Cabeza R, et al. Network analysis of positron emission tomography regional cerebral blood flow data: ensemble inhibition during episodic memory retrieval. *J Neurosci* 1996;16:3753–59. [PubMed: 8642418]
71. Friston KJ. Bayesian estimation of dynamic systems: an application to fMR. *NeuroImage* 2002;16:513–30. [PubMed: 12030834]
72. Friston KJ, Harrison L, Penny W. Dynamic causal modeling. *NeuroImage* 2003;19(4):1273–1302. [PubMed: 12948688]
73. Mechelli A, Price CJ, Noppeney U, et al. A dynamic causal modeling study on category effects: Bottom-up or top-down mediation? *Journal of Cognitive Neuroscience* 2003;15:925–34. [PubMed: 14628754]
74. Penny WD, Stephan KE, Mechelli A, et al. Comparing dynamic causal models. *NeuroImage* 2004;22(3):1157–72. [PubMed: 15219588]
75. Friston KJ. Models of brain function in neuroimaging. *Annual Review of Psychology* 2005;56:57–87.
76. Wold, H. Soft modeling: the basic design and some extensions. In: Wold, KGJaH, editor. *Systems under indirect observation: Causality– Structure–Prediction, Part II*. North-Holland Publishing Company: Amsterdam; 1982. p. 1-54.
77. Hasegawa K, Kimura T, Miyashita Y, et al. Nonlinear partial least squares modeling of phenyl alkylamines with the monoamine oxidase inhibitory activities. *J Chem Inf Comput Sci* 1996;36:1025–29. [PubMed: 8831142]
78. Nestor PG, O'Donnell BF, McCarley RW, et al. A new statistical method for testing hypotheses of neuropsychological/MRI relationships in schizophrenia: Partial least squares analysis. *Schizophr Res* 2002;53:57–66. [PubMed: 11728838]
79. Marti'nez-Montes E, Valde's-Sosa PA, Miwakeichi F, et al. Concurrent EEG/fMRI analysis by multiway partial least squares. *NeuroImage* 2004;22:1023–34. [PubMed: 15219575]
80. Machulda MM, Ward HA, Borowskit B, et al. fMRI in the diagnosis of early Alzheimer's disease. *Neuroimage* 2001;13:704.
81. Eckert T, Barnes A, Dhawan V, et al. FDG PET in the differential diagnosis of parkinsonian disorders. *Neuroimage* 2005;26:912–21. [PubMed: 15955501]
82. Hirao K, Ohnishi T, Hirata Y, et al. The prediction of rapid conversion to Alzheimer's disease in mild cognitive impairment using regional cerebral blood flow SPECT. *Neuroimage* 2005;28:1014–21. [PubMed: 16129627]
83. Whalley HC, Simonotto E, Moorhead W, et al. Functional imaging as a predictor of schizophrenia. *Biol Psychiatry*. 2006in press
84. Bruder GE, Stewart JW, Mayberg HS. Developing predictors of clinical response to antidepressants: neurophysiological and neuropsychological findings. *Biol Psychiatry* 2002;51:1S–198S.
85. Rauch SL, Dougherty DD, Cosgrove GR, et al. Cerebral metabolic correlates as potential predictors of response to anterior cingulotomy for obsessive compulsive disorder. *Biol Psychiatry* 2001;50:659–67. [PubMed: 11704072]
86. Brannan SK, Mayberg HS, McGinnis S, et al. Cingulate metabolism predicts treatment response: a replication. *Biol Psychiatry* 2000;47:107S.
87. Mayberg HS, Brannan SK, Mahurin RK, et al. Cingulate function in depression: A potential predictor of treatment response. *NeuroReport* 1997;8:1057–61. [PubMed: 9141092]
88. Mayberg HS. Modulating dysfunctional limbic-cortical circuits in depression: towards development of brain-based algorithms for diagnosis and optimised treatment. *Br Med Bull* 2003;65:193–207. [PubMed: 12697626]

89. Moresco RM, Colombo C, Fazio F, et al. Effects of fluvoxamine treatment on the in vivo binding of [F-18] FESP in drug naive depressed patients: a PET study. *Neuroimage* 2000;12:452–65. [PubMed: 10988039]
90. Brody AL, Saxena S, Silverman DH, et al. Brain metabolic changes in major depressive disorder from pre- to post-treatment with paroxetine. *Psychiatry Research: Neuroimaging* 1999;91:127–39.
91. Lee JS, Kim BN, Kang E, et al. Regional cerebral blood flow in children with attention deficit hyperactivity disorder: comparison before and after methylphenidate treatment. *Hum Brain Mapp* 2005;24:157–64. [PubMed: 15486990]
92. Goldapple K, Segal Z, Garson C, et al. Modulation of cortical-limbic pathways in major depression - Treatment-specific effects of cognitive behavior therapy. *Arch Gen Psychiatry* 2004;61:34–41. [PubMed: 14706942]
93. Kennedy SH, Evans KR, Kruger S, et al. Changes in regional brain glucose metabolism measured with positron emission tomography after paroxetine treatment of major depression. *Am J Psychiatry* 2001;158:899–905. [PubMed: 11384897]
94. Mayberg HS, Brannan SK, Tekell JL, et al. Regional metabolic effects of fluoxetine in major depression: Serial changes and relationship to clinical response. *Biol Psychiatry* 2000;48:830–43. [PubMed: 11063978]
95. Mayberg HS, Brannan SK, Silva JA, et al. Antidepressant drug response is not the placebo effect: An FDG PET study. *Biol Psychiatry* 2001;49(525)
96. Guo, Y.; Bowman, FD.; Kilts, C. Department of Biostatistics Technical Documents Series. Emory University; Atlanta: 2006. Predicting the brain response to treatment using a Bayesian Hierarchical model.

Figure 1. Voxel time series from a cocaine addict in the inhibitory control study. The figure depicts the varying brain activity (BOLD responses) across 170 scans (blue) and indicates the onsets of the “stop” cue (red) signaling to the subject to refrain from responding to the “go” cue.

Figure 2.

Schematic illustration indicates two-stage statistical modeling procedure and voxel-specific statistical inferences for activation studies. The first-stage model evaluates the vector of BOLD responses in terms of subject-specific effects. The second-stage model combines the subject-specific effects to estimate the associated group-level or population parameters. Hypothesis testing is performed on a contrast of the group-level effects. Computing the test statistic at each voxel produces a t-statistic map. Highlighted voxels in the activation map are voxels producing statistically significant contrasts.

Figure 3.

Three clusters based on the mean brain activity of cocaine addicts prior to treatment in inhibitory control study. Each cluster contains voxels that exhibit similar patterns of brain activity associated with correct inhibitions of a prepotent response. The orange cluster falls within the left and right superior temporal gyrus, the left insula, and right middle frontal gyrus. The yellow cluster includes voxels within the anterior and posterior cingulate and the precuneus. The red cluster consists primarily of voxels in the posterior cingulate and the precuneus.

Figure 4. Results based on a probabilistic independent component analysis of a cocaine addict's fMRI data from the inhibitory control study. This figure depicts (a) a thresholded spatial map for an independent component and (b) the estimated time series associated with this component. Red and blue regions in the spatial map represent voxels with significant positive and negative representations on this component."

Figure 5.

Predicted (a) and observed (b) post-treatment brain activity (regional cerebral blood flow) measurements for four schizophrenia patients, corresponding to a working memory task. The axial slice shown is 6mm below the anterior commissure. Notable differences exist between the predicted post-treatment distributed patterns of brain activity for these patients, and there is a high degree of concordance between the observed and predicted maps for each patient.

Table 1

List of available software tools

Statistical Technique	Software
GLM	SPM, FSL, AFNI
GLMM	MATLAB [†]
Spatial Autoregressive Model	MATLAB [†]
Clustering	MATLAB, HybridClust ^{††}
ICA	FSL and SPM [*]
Functional Connectivity Bayesian Hierarchical Model	MATLAB ^{††}
DCM	SPM
PLS	SPM ^{**}
Prediction	MATLAB [†]

[†] Specialized MATLAB functions developed by authors.

^{††} Set of MATLAB-based functions developed by R. Patel.

^{*} Must have the Group ICA of fMRI Toolbox (GIFT) add-on to SPM.

^{**} Must have the Multivariate Methods for fMRI (MM) toolbox add-on to SPM. Note that this table is not intended to provide a comprehensive list of available software, but instead to provide readers with some direction on how to implement some of the statistical techniques described in this article.