Tools of the Trade

Region of interest analysis for fMRI

Russell A. Poldrack

Department of Psychology, Department of Psychiatry and Biobehavioral Sciences, and Brain Research Institute, University of California Los Angeles, Los Angeles, CA, USA

A common approach to the analysis of fMRI data involves the extraction of signal from specified regions of interest (or ROI's). Three approaches to ROI analysis are described, and the strengths and assumptions of each method are outlined.

INTRODUCTION

The first step in fMRI statistics is almost invariably to create a thresholded statistical map, showing which regions are 'activated' above some particular threshold. For simple comparisons, it is often sufficient to stop there and simply report the results of the thresholded map. However, there are often reasons to look further into particular regions of interest (ROIs). In this brief review, I will outline several approaches to ROI analysis and discuss the advantages and drawbacks of each.

Why ROIs?

There are three main reasons that one might want to perform an ROI analysis, which have very different justifications and make very different assumptions. The first reason to perform an ROI analysis is simply to explore one's data. In complex designs, such as factorial designs with multiple levels, it can often be difficult to discern the pattern of activity across conditions from an overall map. It is often useful to see the signal in areas of interest plotted for each condition or plotted against other variables of interest. The second reason is to control for Type I error by limiting the number of statistical tests to a few ROIs. The third reason is to limit testing to a region that is functionally defined on the basis of some other information, such as a separate 'localizer' scan or condition. I discuss each of these approaches in turn.

ROIs for exploration

The most theoretically agnostic use of ROI analysis is to simply explore the underlying signal behind a whole-brain voxel-wise analysis. In this case, one is not making any claims regarding the un-biasedness of the analysis; rather, the goal is to depict the pattern of signal across conditions. As noted earlier, this is particularly useful in complex designs such as factorial designs, where there are multiple conditions. The most common approach for exploratory ROI analysis is to create small ROIs (usually spheres) at the peaks of activation clusters; in the case of large clusters, it can be useful to create ROIs for additional local maxima in order to explore multiple regions within the cluster. To ensure that the sphere only contains voxels that were truly activated, these spheres are often masked with the thresholded activation map. It is critical to keep in mind that one cannot make any conclusions about the statistical significance of tests on the resulting data if the selection of the ROI was based on the same contrast; because the region was chosen for its significance in the whole-brain analysis, it will necessarily show a significant result in any ROI analysis. Nonetheless, this approach can be very useful for exploring patterns of activity across conditions.

Although ROI analysis is most often considered for analysis of activations, it can sometimes be equally useful for determining the reasons for lack of activation. As an example, a recent correlational analysis between fMRI and behavioral data in our laboratory failed to uncover any activation in several regions of prior interest, much to our surprise. We performed an ROI analysis on several regions (using small spheres placed in anatomical ROIs) and quickly saw the reason for this lack of activation: Whereas the group as a whole showed a striking correlation, one subject was an extreme outlier who suppressed the correlation. Analysis using robust statistics (Wager *et al.*, 2005) can be particularly useful in cases like this because they can reduce the effects of outliers.

ROIs for statistical control

Another approach is to pre-specify a set of anatomical ROIs, and then to perform statistics across these regions. This is generally done to reduce the severity of correction for

multiple tests; instead of correcting for the large number of voxels in the brain, once can correct only for a small number of ROIs. I would hasten to note that there has not been a systematic study of the power of this approach in comparison to standard methods for whole-brain error control, such as Gaussian random field theory (Worsley et al., 1992), randomization tests (Nichols and Holmes, 2002), or false discovery rate corrections (Genovese et al., 2002).

One difficult question that arises with regard to this approach is how to measure signal within the ROI. It is generally the case that regions specified in this approach will be relatively large (e.g. the entire superior temporal gyrus), and that even if the region is significantly active, this activation may only occur in a small proportion of voxels in the ROI. This suggests that simply averaging across the entire region could swamp the signal from this small number of voxels with noise from the remaining non-activated voxels. One approach that has often been used is to threshold the statistical map and count the number of activated voxels in each region. However, this approach can be very sensitive to the specific threshold. Further, voxel counts have been shown to be an unreliable measure of activation compared to direct measures of signal change (Cohen and DuBois, 1999). These difficulties in defining in the signal suggest that this approach may not be optimal. This approach also relies upon the assumption that the specified regions are functionally homogenous; if, for example, there exist areas of both activation and deactivation within the region, these may cancel each other out.

There is another approach by which ROIs can be used for statistical control, which involves the restriction of voxel-wise analyses to a set of ROIs and then controlling for multiple comparisons only in those voxels, for example, using Gaussian random field theory for small volumes (Worsley et al., 1996) as implemented in SPM, or using randomization tests restricted to the ROIs. This approach does not suffer from the same problems of signal specification as the foregoing methods, but it also does not provide the same level of insight into patterns of activation as measures involving quantification across conditions.

ROIs for functional specification

Another use for ROIs is to examine activity within a set of voxels that are functionally coherent, in order to examine their sensitivity to some other manipulation. This approach is most prevalent in fMRI studies of visual processing, where localizers are used to identify functionally specific regions such as retinotopically organized regions in early visual cortex (Engel et al., 1994) or motion-sensitive voxels in area MT (Tootell et al., 1995). In these cases, the ROIs are generally determined using a 'localizer' scan that is separate from the scan of primary interest. However, there is no requirement that the localizer be performed as part of a separate scan. Friston et al. (2006) have argued that it is

better to embed the localizer in factorial design that also includes the comparison of interest; this allows one not only to use the localizer to examine effects on some other comparison of interest (since they are orthogonal in the factorial design), but also to examine interactions between the localizer manipulation and other conditions. There is an ongoing debate regarding the usefulness of such functional localizer approaches (Friston et al., 2006; Saxe et al., 2006); although they may be useful in some cases, it is important to avoid reifying regions based on their functional activation alone, given the potential importance of distributed representations (Haxby et al., 2001) and the difficulty in ascribing a single function to a region based on fMRI data (Poldrack, 2006).

Defining ROIs

ROIs can be defined either in terms of structural or functional features. Structural ROIs are generally defined based on macroanatomy, such as gyral anatomy. In many cases, the best practice is to define such ROIs for each subject based on their own anatomy, since there can be substantial variability between individuals in macroscopic anatomy. Recent developments in automated anatomical labeling offer the promise of highly reliable labeling of cortical and subcortical structures in individual anatomical images with a minimum of manual intervention (Fischl et al., 2004), though it will remain important to confirm these results against the actual anatomy. One common practice that requires extreme caution is the use of ROIs based on single-subject anatomical atlases, such as the AAL atlas (Tzourio-Mazoyer et al., 2002) or the Talairach atlas (Talairach and Tournoux, 1988); because of the inability of spatial normalization to perfectly match brains across individuals, there will be substantial lack of overlap between any group of subjects and these atlases (Nieto-Castanon et al., 2003). If it is necessary to use atlas-based ROIs (i.e. ROIs not derived from one's own subjects) then the best practice is to use ROIs based on probabilistic atlases of macrosopic anatomy (Hammers et al., 2003; Shattuck et al., 2006) or probabilistic atlases of Brodmann's areas which are available as part of the SPM Anatomy Toolbox (Eickhoff et al., 2006).

Functional ROIs are generally based on analysis of data from the same individual. One common approach is to use a separate 'localizer' scan to identify voxels in a particular anatomical region that show a particular response (e.g. voxels in the fusiform gyrus that are more responsive to faces than other objects); these voxels are then explored to examine their response to some other manipulation. Alternatively, functional ROIs can be created using orthogonal contrasts in a factorial design (Friston et al., 2006). Exploratory ROIs are often created by placing small spheres at local maxima in the statistical map; this provides a set of ROIs that span the clusters of interest. Because the goal of exploratory ROIs is not statistical control, it is also acceptable to place ROIs in anatomical areas of interest

(using one's best judgment about the placement), particularly for examining null results; however, it is again critical to note that although this kind of analysis can be useful for exploration it must not be used for inference since it is heavily biased.

One additional way that ROIs can be created is based on previous studies. Although one can take the stereotactic coordinates from an activation in single study and place an ROI at that location, it is better practice to derive ROIs from meta-analyses of the domain or task of interest. There are now well-established methods for meta-analysis of functional imaging studies (e.g. Turkeltaub *et al.*, 2002; Wager and Smith, 2003), and these methods can be used to generate ROIs that will be less sensitive to noise than those based on single-study activations.

Tools for ROI analysis

There are a number of tools available for ROI analysis. Several common software packages (e.g. SPM, AFNI, BrainVoyager) include tools for ROI analysis. In addition, the MarsBar tool for SPM (http://marsbar.sourceforge.net/) and the FSL ROI Toolbox (http://spm-toolbox.sourceforge.net) provide tools for performing ROI analyses on data from these packages as well.

There are two ways in which data are commonly extracted for ROI analysis. In *parameter estimate extraction*, one extracts the estimated parameter value (e.g. 'beta' images in SPM, 'pe' images in FSL) for each condition in the statistical model, with zero determined by the implicit baseline (i.e. whatever is not included in the model). This can be particularly useful for understanding contrasts that include a number of conditions (assuming that each condition is modeled separately), though it does not in principle provide any new information other than collapsing across voxels within the region, which may decrease noise.

In *hemodynamic response extraction*, the raw data are interrogated and the entire hemodynamic response to each condition across the ROI is estimated, generally using a finite impulse response model that estimates the response at each timepoint following the stimulus (Dale, 1999). This approach provides a different view of the data by showing the entire estimated response in time (making no assumptions about its shape), rather than the fit of an assumed hemodynamic response. It should be noted that this approach can tend to overfit the data given the large number of parameters (one for each timepoint in the hemodynamic response), and thus one can sometimes see estimated hemodynamic responses that are not physiologically plausible, especially with smaller sample sizes. Approaches using constrained basis sets (Woolrich *et al.*, 2004) may be useful for obtaining better estimates of the underlying hemodynamic response, but these methods have not been integrated with existing ROI anlaysis software.

SUMMARY

There are a number of reasons to perform ROI analyses, each of which involves a particular set of assumptions. In my opinion, the use of exploratory ROI analysis should be standard practice in all functional imaging laboratories. It can provide substantial insight into the nature of activation signals in complex models as well as provide valuable assistance in diagnosing model failures, but it is critical that the investigator not use these exploratory analyses for inference as their results are biased by ROI choice. The use of ROIs for control of Type I error suffers from difficulties in the specification of the signal within large regions, and appears to be waning in the face of advancements in whole-brain error control methods and increased sample sizes for fMRI studies. However, the use of ROIs for small volume correction with voxel-wise statistics is well established. The use of functional localizers is well-established in domains where the underlying function can be unequivocally localized (as in visual retinotopy), but there remains debate about the usefulness of functional localizer approaches more generally. In any case, researchers should take care to avoid reifying ROIs, as functional neuroimaging has not yet established the level of selectivity necessary to label regions based on activation (cf. Poldrack, 2006).

REFERENCES

Cohen, M.S., DuBois, R.M. (1999). Stability, repeatability, and the expression of signal magnitude in functional magnetic resonance imaging. *Journal of Magnetic Resonance Imaging*, 10, 33–40.

Dale, A.M. (1999). Optimal experimental design for event-related fMRI. *Human Brain Mapping*, 8, 109–14.

Eickhoff, S.B., Heim, S., Zilles, K., Amunts, K. (2006). Testing anatomically specified hypotheses in functional imaging using cytoarchitectonic maps. *Neuroimage*, 32, 570–582.

Engel, S.A., Rumelhart, D.E., Wandell, B.A., et al. (1994). fMRI of human visual cortex. *Nature*, 369, 525.

Fischl, B., van der Kouwe, A., Destrieux, C., et al. (2004). Automatically parcellating the human cerebral cortex. *Cerebral Cortex*, 14, 11–22.

Friston, K.J., Rotshtein, P., Geng, J.J., Sterzer, P., Henson, R.N. (2006). A critique of functional localisers. *Neuroimage*, 30, 1077–87.

Genovese, C.R., Lazar, N.A., Nichols, T. (2002). Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage*, 15, 870–8.

Hammers, A., Allom, R., Koepp, M.J., et al. (2003). Three-dimensional maximum probability atlas of the human brain, with particular reference to the temporal lobe. *Hum Brain Mapp*, 19, 224–47.

Haxby, J.V., Gobbini, M.I., Furey, M.L., Ishai, A., Schouten, J.L., Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, 293, 2425–30.

Nichols, T.E., Holmes, A.P. (2002). Nonparametric permutation tests for functional neuroimaging: a primer with examples. *Human Brain Mapping*, 15, 1–25.

Nieto-Castanon, A., Ghosh, S.S., Tourville, J.A., Guenther, F.H. (2003). Region of interest based analysis of functional imaging data. *Neuroimage*, 19, 1303–16.

Poldrack, R.A. (2006). Can cognitive processes be inferred from neuroimaging data? *Trends in Cognitive Science*, 10, 59–63.

Saxe, R., Brett, M., Kanwisher, N. (2006). Divide and conquer: a defense of functional localizers. *Neuroimage*, 30, 1088–96; discussion 1097–89.

Shattuck, D.W., Hojatkashani, C., Mirza, M., et al. (2006). *Construction of a 3D Probabilistic Atlas of Human Brain Structures*, Paper presented at the 12th annual Meeting of the Organization for Human Brain Mapping, Florence, Italy..

Talairach, J., Tournoux, P. (1988). *Co-planar Stereotaxic Atlas of the Human Brain.* Stuttgart: Thieme..

Tootell, R.B., Reppas, J.B., Kwong, K.K., et al. (1995). Functional analysis of human MT and related visual cortical areas using magnetic resonance imaging. *Journal of Neuroscience*, *15*, 3215–30.

Turkeltaub, P.E., Eden, G.F., Jones, K.M., Zeffiro, T.A. (2002). Meta-analysis of the functional neuroanatomy of single-word reading: method and validation. *Neuroimage*, *16*, 765–80.

Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., et al. (2002). Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *Neuroimage*, *15*, 273–89.

Wager, T.D., Smith, E.E. (2003). Neuroimaging studies of working memory: a meta-analysis. *Cognitive and Affective Behavioral Neuroscience*, *3*, 255–74.

Wager, T.D., Keller, M.C., Lacey, S.C., Jonides, J. (2005). Increased sensitivity in neuroimaging analyses using robust regression. *Neuroimage*, *26*, 99–113.

Woolrich, M.W., Behrens, T.E., Smith, S.M. (2004). Constrained linear basis sets for HRF modelling using Variational Bayes. *Neuroimage*, *21*, 1748–61.

Worsley, K.J., Marrett, S., Neelin, P., Evans, A.C. (1992). A three-dimensional statistical analysis for CBF activation studies in human brain. *Journal of Cerebral Blood Flow and Metabolism*, *12*, 900–18.

Worsley, K.J., Marrett, S., Neelin, P., Vandal, A.C., Friston, K.J., Evans, A.C. (1996). A unified statistical approach for determining significant signals in images of cerebral activation. *Human Brain Mapping*, *4*, 58–83.