

Etzel, J. A., Cole, M.W., Zacks, J.M. Braver, T.S. "Enhanced decoding of task representations under reward motivation conditions". Poster presentation at the 2012 Society for Neuroscience (SfN) Meeting.

A major function assumed to be mediated by the brain frontoparietal cognitive control network is the representation of task sets (or goals). The neural coding of task set information has typically been investigated by studies employing multi-tasking or task-switching paradigms. In prior work, we observed enhanced task-switching performance and increased activation in the cognitive control network under conditions of high motivation (i.e., availability of performance-contingent reward incentives; Savine and Braver, 2010).

In the current study we examined the hypothesis that motivation-related influences on brain activity and performance occur via improved encoding and representation of task set information. A multivariate pattern analysis (MVPA) was used to test whether task information could be decoded more successfully on reward incentive compared to randomly intermixed non-incentive trials. 18 participants took part in a two-session fMRI experiment involving cued task-switching between face and word classification tasks, with each trial consisting of an advance cue indicating which task was to be performed on the upcoming bivalent target stimulus (face with word superimposed).

A whole-brain searchlight analysis of the first session (in which no incentives were available) identified clusters of voxels that could reliably classify task identity (face or word). This analysis found significant classification within a widespread set of regions, including prefrontal and parietal components of the cognitive control network, plus extrastriate visual cortex. In a subsequent analysis, each identified cluster (treated as a ROI) was trained to classify task identity on the first session, and then tested on the second session, providing an unbiased assessment of task decoding accuracy. Successful classification was observed for this second session on incentive trials, with significantly higher accuracy than on non-incentive trials. Consistent effects were observed across the full set of regions; of the individual ROIs, the incentive effects were expressed most strongly in left extrastriate visual cortex (as determined via permutation and sensitivity testing). Together, these results indicate that enhanced task-switching performance on incentive trials might arise from more effective task coding and representation. More broadly, they highlight the utility of reward motivation manipulations for understanding the nature of task representation in the human brain.

Etzel, J.A., Cole M.W., Braver T.S. "Looking Outside the Searchlight," 2011. Poster presentation at the NIPS workshop on Machine Learning and Interpretation in Neuroimaging (MLINI 2011).

Searchlight analysis (information mapping) with pattern classifiers is a popular method of fMRI analysis often interpreted as localizing informative voxel clusters. Applicability and interpretation is limited, however, by its dependency on searchlight radius, assumption that information is present at all spatial scales, and susceptibility to overfitting. These problems are demonstrated in a dataset in which, contrary to common expectation, voxels identified as informative, as a group, do not clearly contain more information than those not so identified.

Etzel, J. A., Cole, M. W., Braver, T. S. "A Strategy for Characterizing the Spatial Distribution of Information Within a Region of Interest," 2011. Poster presentation at the 2011 Organization for Human Brain Mapping (HBM) Meeting.

Introduction

Many multivariate analyses of fMRI data follow a region-of-interest (ROI) approach: a group of voxels is identified using anatomical (or other) criteria, then the classification accuracy of this group of voxels is determined (Etzel, et al. 2009). Significant classification accuracy is interpreted as indicating that the pattern of activity in the voxels making up the ROI consistently varies with the stimulus types.

For some hypotheses we wish to go beyond the level of the ROI and identify which voxels within the ROI are contributing to the successful classification. Several methods for identifying locations with information below the level of the ROI exist, such as multivariate searchlights (Kriegeskorte, et al. 2006) and mapping support vector machine weights (e.g. Lee, et al. 2010). These methods can be quite useful, but do not characterize all types of patterns nor distinguish certain information distributions. The method described in this work aims to characterize which voxels contain information in a well-classifying ROI.

As an analogy for the spatial arrangements of interest here, consider a dessert of molded gelatin. Taken as a whole, the dish has a clear flavor (the ROI has information; it can classify the stimuli). But is the flavor of the dish due to rare, but strongly flavored pineapple chunks in unflavored gelatin (a subset of the ROI, possibly spatially discontinuous, with high information while the rest of the ROI contains minimal signal), or is the gelatin flavored, without fruit at all (the information is spread throughout the ROI with minimal spatial variation)? Or perhaps a combination, such that chunks of pineapple are scattered in pineapple gelatin (areas of high information in an area with weaker, but still detectable, information)?

Methods

We assume that a large ROI (on the order of a thousand voxels) capable of classifying the stimuli has been identified using anatomical or other independent criteria. We also assume that the classification was performed in each subject individually, followed by combining results across subjects.

First a searchlight analysis (with radius of one or two voxels) is performed within the target ROI, resulting in an accuracy for each voxel (the accuracy of that voxel's searchlight). The significance level is then determined for each voxel by an across-subjects t-test.

At this point voxels surviving the statistical threshold can be plotted, generally interpreted as areas of the ROI with information, at the spatial scale of the searchlight, across subjects. While informative, we may also ask whether information is present at other spatial scales, or if the voxels not identified by the searchlights are also capable of classification.

Using the searchlight results as an index of information at a small spatial scale, a series of statistical thresholds are chosen. At each threshold the voxels with searchlight significance greater than the threshold are identified. These voxels, with their searchlights, are combined into a subset. The rest of the voxels in the ROI are combined into a second subset. Examining the accuracy, size, and location of these subsets allows us to determine if the voxels leading to the ROI's classification are concentrated in

discrete clusters (the pineapple), and if the voxels outside of these cluster (the gelatin) contain information.

Results and Conclusions

Activity patterns within ROIs from several fMRI and simulated datasets were investigated using this strategy, demonstrating its usefulness and implementation. Exploring the spatial organization of information within a well-classifying ROI has the potential to increase our understanding of the results of ROI-based multivariate analyses, as well as the structure of fMRI data in general. The method described in this work is a strategy for beginning such exploration.

References

Etzel JA, Gazzola V, Keysers C. 2009. An introduction to anatomical ROI-based fMRI classification analysis. *Brain Research* 1282:114-125.

Kriegeskorte N, Goebel R, Bandettini P. 2006. Information-based functional brain mapping. *PNAS* 103(10):3863-3868.

Lee S, Halder S, Kübler A, Birbaumer N, Sitaram R. 2010. Effective functional mapping of fMRI data with support-vector machines. *Human Brain Mapping* 31(10):1502-1511.