A Greedy Best-First Search Algorithm for Accurate Functional Brain Mapping

Nima Asadi (nima.asadi@temple.edu)
Department of Computer and Information Science, Temple University
1925 N 12th St, Philadelphia, PA 19122

Yin Wang (mirrorneuronwang@gmail.com)
Department of Psychology, Temple University
1701 N 13th St, Philadelphia, PA 19122

Ingird Olson (iolson@temple.com)
Department of Psychology, Temple University
1701 N 13th St, Philadelphia, PA 19122

Zoran Obradovic (zoran.obradovic@temple.edu)
Department of Computer and Information Science, Temple University
1925 N 12th St, Philadelphia, PA 19122

Abstract

Detecting the most relevant brain regions for explaining the distinction between conditions is one of the most sought after goals in cognitive neuroimaging research. A popular approach for achieving this goal is the multivariate pattern analysis (MVPA) which is commonly conducted through the searchlight procedure due to its advantages such as being intuitive and flexible with regards to search space size. However, the searchlight approach suffers from a number of limitations that lead to misidentification of truly informative voxels or clusters of voxels which in turn results in imprecise information maps. These limitations mainly stem from the fact that the information value of the search spheres are assigned to the voxel at the center of them, as well as the requirement of manual assignment of searchlight radius. This issue becomes more severe when larger searchlight radius values are selected which makes truly informative voxels less likely to be identified. In this paper we propose a data-driven algorithm for creating the information map of the brain while alleviating the above mentioned issues.

Keywords: functional magnetic resonance imaging (fMRI); Multivariate pattern analysis (MVPA); Greedy Algorithm; Brain Mapping; Searchlight analysis

Introduction

In the common form of fMRI, the blood-oxygen-level dependent (BOLD) contrast is extracted as the response signal in order to measure neural activity in the brain (Logothetis, Pauls, Augath, Trinath, & Oeltermann, 2001). Measurement of this response signal over time forms a time course for each voxel whose dimensions depend on the spatial resolution of the imaging machine. Popular approaches for analyzing the fMRI time courses can be broken down into two main categories: multi-voxel pattern analysis, also known as MVPA, which aims to detect patterns among conditions observed in multiple voxels, and voxel-wise univariate analysis (Norman, Polyn, Detre, & Haxby, 2006; Groppe, Urbach, & Kutas, 2011). Unlike univariate analyses, MVPA approaches are designed to allow researchers to test how dispersed patterns of BOLD activation across multiple voxels relate to experimental conditions. One approach in multi-voxel scheme is to compare and analyze spatially averaged (smoothed) measured BOLD activations across the entire regions of interest. Advantages of this approach include increases in the signal to noise ratio as well as the consistency of the analysis among subjects can be noted. However, spatial smoothing leads to significant loss of information about the patterns of activation within the regions of interest. This information includes the activities and dynamics within subregions which can provide valuable insight into their relation with different mental states. This issue becomes more complex when dealing with regions with higher numbers of voxels. Therefore, in order to capture such information, it is necessary to consider the BOLD activity in smaller spherical subsets. Moreover, in the absence of clear knowledge about the relevant brain regions, or when the neural pattern is too distributed that precise a priori ROIs cannot be drawn, this confirmative approach is not appropriate. The question of identifying relevant subregions with regards to specific conditions is however not a new question. Perhaps the most commonly employed approach for such applications is the searchlight method proposed by Nikolaus Kriegeskorte (Kriegeskorte, Goebel, & Bandettini, 2006; Kriegeskorte & Bandettini, 2007), which given the dimensions of a sphere window, performs a search across a brain region to find the most informative set of neighboring voxels. In this multivariate approach, spatial patterns of activity within the search window are compared between two groups using supervised machine learning approaches or statistical discriminant analysis (Haxby et al., 2001). The search sphere (searchlights) is centered on every voxel, i.e. the derived separability value for each voxel is derived from the discrimination score of its surrounding searchlight, not the voxel individually. Advantages of searchlight analysis include its data-driven nature, its ability in performing whole-brain search without the need to specify brain regions, and its high interpretability. However, searchlight analysis suffers from critical limitations that can lead to erroneous detection of informative subregions (see (Etzel, Zacks, & Braver, 2013)). One major issue with this procedure is that it can declare a subregions with a few highly-informative voxels as informative. This becomes a more serious issue as the search radius becomes larger. Also, choosing an appropriate search radius is essential in this approach, which depends on the shape and size of the region being searched. However, finding the most discriminative subregion through a brute-force search over all possible search radius values is difficult due to its run time complexity, specially when being applied in whole-brain analysis. Aside from this issue, the shape of the searchlight can limit the detection of the subregions with the highest discrimination power due to the fact that searchlight spheres are commonly in the shape of a circle, or square, which forces subregions with irregular shapes to fall between multiple searchlight positions. This issue can partially be relieved through assigning the searchlight sphere as small as possible, at the expense of overfitting.

In order to tackle the mentioned issues, we propose a new algorithm for discovering the most discriminative subregions to exaplain the disparities in neural activities among different populations. Through a completely data-driven search, the proposed approach achieves this goal without the requirement to specify the searchlight sphere radius by the user. Through empirical results on a synthetic dataset we show that this technique significantly increases the speed of the suggested approach. In the next section, we explain the suggested methodology in more detail.

Approach

The purpose of this work is to create the information map of brain (or regions of brain) with regards to a certain condition, e.g. neurological disorder, age, mental state, etc. In other word, given two (or more) populations, the goal is to discover the level at which the BOLD activation of voxel clusters differs between the groups. In this study, we name these difference levels of each voxel/cluster their discriminant score.

Therefore, the discriminant score of each cluster of voxels is the extent to which their neural activity can be distinguished between the groups. Note that voxel clusters are groups of neighboring voxels whose size can span from one voxel to the entire search space. The input to the proposed approach is the average BOLD value of the voxels of two populations, and the output is an information map of the brain where each voxel/cluster is assigned a discriminant score.

There are several approaches to calculate the discriminant score of a voxel, such as discriminant analysis (McLachlan, 2004; Venables & Ripley, 2013) and cluster analysis (Rousseeuw, 1987; Kaufman & Rousseeuw, 2009). The purpose of both of these analyses is to provide a measure of how separable clusters/groups of data are. While several measures have been proposed for cluster analysis, the most general measure is the Davies-Bouldin index (Davies & Bouldin, 1979) which calculates the ratio of average distance between the data points belonging to each cluster (group of data) and their centroid to the distance between the centroids of the clusters. This can be formulated as follows:

$$DB = \frac{1}{n}\sum_{i=1}^{n} max \frac{(dist_i + dist_j)}{d(c_i, c_j)} \qquad (1)$$

where $dist_i$ and $dist_j$ are the average distance of the points from the centroids in each cluster i and j, and c_i and c_j are the centroids of the clusters I and j. A low Davies-Bouldin score means a low intra-cluster distances (high intra-cluster similarity) and high inter-cluster distances (low inter-cluster similarity), which provides a high discriminant score (inverse Davies-Bouldin score).

To create the information map of the search space (whole brain or a region of interest) we propose a greedy best-first search algorithm that traverses every voxel through a heuristic and outputs the optimal discriminant score of the voxel clusters. In other word, a greedy best-first algorithm picks the "best" voxel according to a heuristic, which in our case is the discriminant score of the voxels (Dechter & Pearl, 1985; DeVore & Temlyakov, 1996).

The proposed algorithm goes as follows: Starting from a random voxel V_s, after measuring its discriminant score through one of the mentioned measures, the algorithm calculates the discriminant score for all of its immediate neighbors. Then, V_s is combined with each neighboring voxel (a pair of two voxels at a time), and the discriminant score of their combination is measured. In the next step, the algorithm searches among those neighboring voxels to V_s whose combination with V_s provides a higher discriminant score than both V_s and its neighbor, and picks the one whose combination with V_s gives the highest discriminant score. This neighboring voxel is the new V_s, and the same steps as above are performed recursively. If we reach a point where there are no immediate neighbors to V_s whose combination with Vs provides a higher discriminant score, the algorithm randomly selects a previously untraversed voxel and performs the same criteria starting from the new voxel. The algorithm continues this search until all voxels in the search space are assigned a discriminant score value.

The idea behind this method is to maximize the discriminant score of a group of neighboring voxels. To achieve this goal, when visiting each new voxel, its discriminant score is compared with the cases where it is paired with each of its neighbors one by one, and the best neighbor is selected, combined with it to create a cluster, and the discriminant score of the both of the is set to be the score of the pair(cluster). The algorithm then performs the same criteria with the newly selected voxel.

An example demonstration of this algorithm is provided in Figure 1. where each cube represents a voxel in the search space. The algorithm starts from a single random voxel (the dark red voxel)in the top plot, and continues traversing the search space based on the discriminant score heuristic as the next plots show. The transparent voxels represent the neigh-

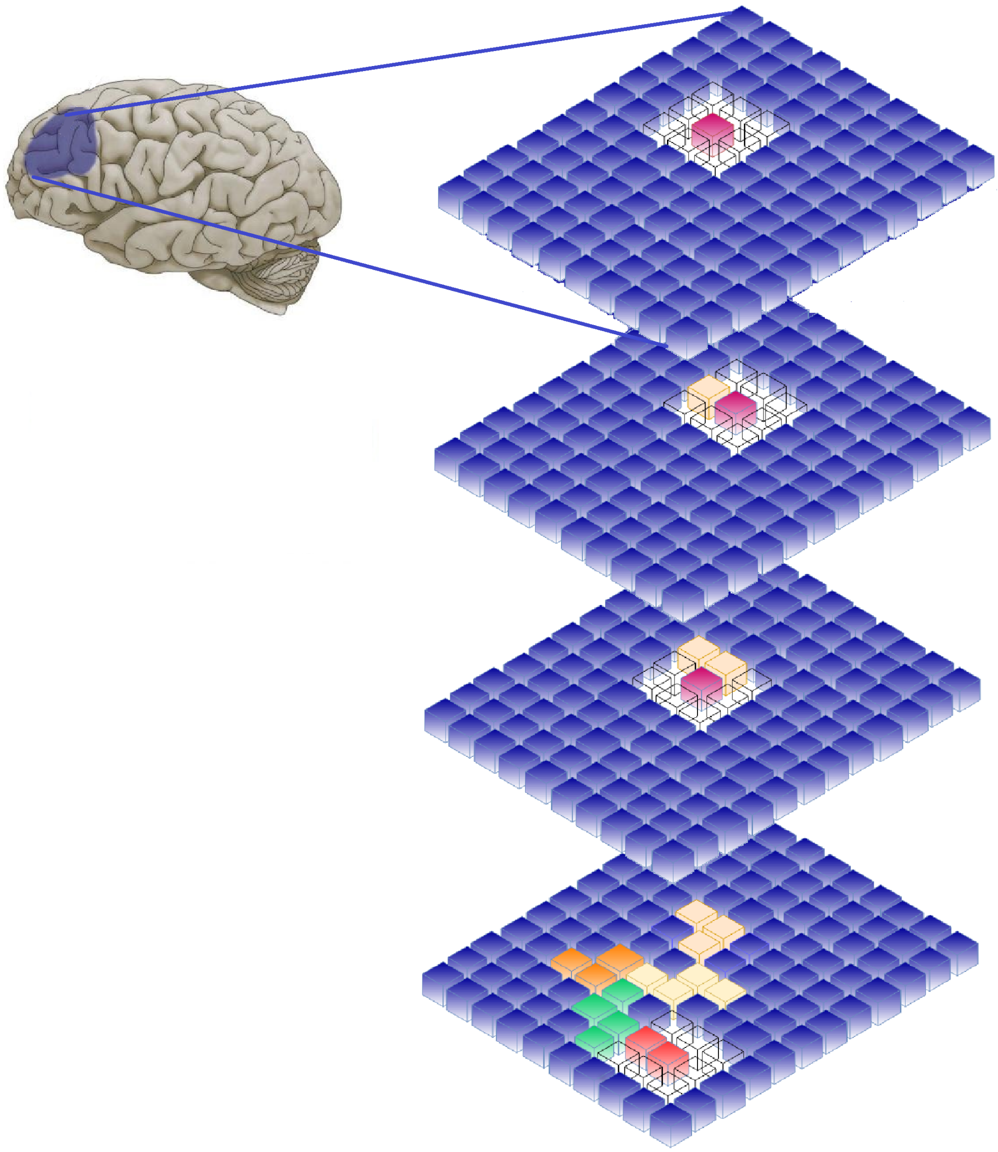

Figure 1: Steps of the path traversed by the proposed algorithm.

bors of the voxel being analyzed in each step. Note that in reality, the search occurs in three dimensional space, but for simplicity of demonstration we only show a flat representation of the search space. To distinguish the clusters from each other, each cluster is depicted with a different color. When the algorithm does not find any neighboring voxels whose combination with the voxel being traversed gives a higher discriminant score than both of them, it starts from a new random neighboring voxel, i.e. a new cluster (color) starts to take shape. As mentioned in the description of the algorithm, all of the voxels in each cluster have similar discriminant score, which is the score of the entire cluster. As mentioned previously, the algorithm stops the search when the entire search space is traversed, resulting in a discriminant score (information) map.

Results

In order to evaluate the reliability of the suggested approach, we performed a classification between two conditions using the most informative subregions (e.g. subregions with highest discriminant score) derived through the suggested algorithm. For this purpose, we created synthetic data of average fMRI time courses for two conditions. These time courses were generated based on values extracted independently from a Gaussian distribution for 10,000 voxels. Noisy values are added to the signal with the constraints of realistic degree of correlation between adjacent voxels. We then applied both the searchlight procedure and the proposed approach to generate two separate information maps of the synthetic search space. After the information maps were created, we selected the voxels with top 100 discriminant scores from both information maps as the feature vector and trained a logistic regression classifier with 5 fold train-test split. For obtaining the information map based on the searchlight approach, we assigned 10 different search radius values r between 1 to 10, and selected the one that provided the best accuracy, which was $r = 3$ voxels The classification results are provided in Figure 2. As it can be seen, the proposed algorithm improves the area under the ROC curve significantly. Moreover, the performance of the proposed algorithm does not depend on tuning the search radius. Further improvement of the results is possible through more advanced and accurate measurements such as nonlinear cluster analysis. Also, performing experiments on biological data to further verify the results is considered for future work.

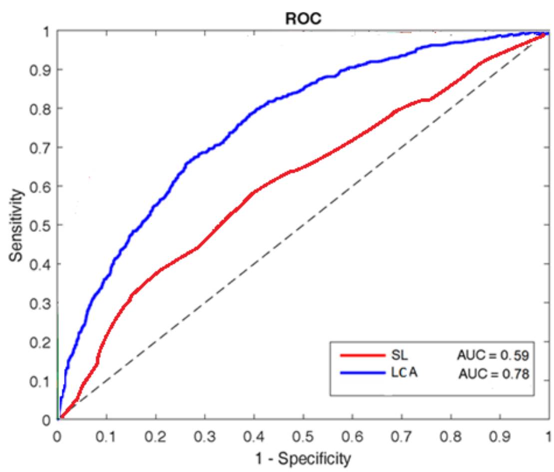

Figure 2: Prediction accuracy using the top 50 voxel/regions derived from the searchlight method (SL) and the proposed algorithm with linear cluster analysis (LCA).

Conclusion

Multivariate pattern analysis (MVPA) is a popular approach for detecting differences between conditions based on the neural activity of distributed voxels as measured by fMRI. A popular technique for such analysis is the searchlight procedure due to its ability to perform whole-brain analysis, as well as its intuitiveness. However, the searchlight approach suffers from a number of limitations that can lead to misidentification of regions as informative (false positives), or failure to detect truly informative voxels or clusters (false negatives). Moreover, determining the optimal search radius is challenging and commonly carried out through trial-and-error, potentially leading to over-fitting.

In this paper we suggested a new MVPA approach to address the mentioned limitations based on a data-driven neighborhood best-first search algorithm. In other word, the proposed method not only removes the requirement of assigning a search radius, it also provides a more precise information map of the search space. The experimental results on a synthetic dataset display higher accuracy of the proposed algorithm compared to the searchlight technique.

This method is useful for basic research on the brain as well as the diagnosis of psychiatric disorders. Future work will include incorporation of non-linear discriminant analysis and further experiments of this algorithm on biological datasets with known ground truth.

References

Davies, D. L., & Bouldin, D. W. (1979). A cluster separation measure. *IEEE transactions on pattern analysis and machine intelligence*(2), 224–227.

Dechter, R., & Pearl, J. (1985). Generalized best-first search strategies and the optimality of a. *Journal of the ACM (JACM)*, *32*(3), 505–536.

DeVore, R. A., & Temlyakov, V. N. (1996). Some remarks on greedy algorithms. *Advances in computational Mathematics*, *5*(1), 173–187.

Etzel, J. A., Zacks, J. M., & Braver, T. S. (2013). Searchlight analysis: promise, pitfalls, and potential. *Neuroimage*, *78*, 261–269.

Groppe, D. M., Urbach, T. P., & Kutas, M. (2011). Mass univariate analysis of event-related brain potentials/fields i: A critical tutorial review. *Psychophysiology*, *48*(12), 1711–1725.

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*(5539), 2425–2430.

Kaufman, L., & Rousseeuw, P. J. (2009). *Finding groups in data: an introduction to cluster analysis* (Vol. 344). John Wiley & Sons.

Kriegeskorte, N., & Bandettini, P. (2007). Combining the tools: activation-and information-based fmri analysis. *Neuroimage*, *38*(4), 666–668.

Kriegeskorte, N., Goebel, R., & Bandettini, P. (2006). Information-based functional brain mapping. *Proceedings of the National academy of Sciences of the United States of America*, *103*(10), 3863–3868.

Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., & Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fmri signal. *Nature*, *412*(6843), 150.

McLachlan, G. (2004). *Discriminant analysis and statistical pattern recognition* (Vol. 544). John Wiley & Sons.

Norman, K. A., Polyn, S. M., Detre, G. J., & Haxby, J. V. (2006). Beyond mind-reading: multi-voxel pattern analysis of fmri data. *Trends in cognitive sciences*, *10*(9), 424–430.

Rousseeuw, P. J. (1987). Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, *20*, 53–65.

Venables, W. N., & Ripley, B. D. (2013). *Modern applied statistics with s-plus*. Springer Science & Business Media.