

A SPATIO-TEMPORAL SUPPORT VECTOR MACHINE SEARCHLIGHT FOR fMRI ANALYSIS

A. Ravishankar Rao, Rahul Garg, Guillermo A. Cecchi

Computational Biology Center
IBM T.J. Watson Research Center
Yorktown Heights, NY 10598, USA

ABSTRACT

We apply support vector machines (SVMs) in the context of fMRI analysis, in order to identify brain regions that are predictive of the experimental conditions. For the spatial SVM, we use the data within local 3D windows, called a searchlight, to train an SVM classifier to distinguish different experimental protocol conditions. Brain regions with high classification accuracy are identified as being implicated in the experimental task. Similarly for the temporal SVM, we use temporal sequences for every voxel to train a classifier.

A major technical challenge is the higher computational overhead associated with SVMs. We overcome this by using parallel programming techniques based on MPI (message passing interface) that achieve load balancing.

We report results on two separate datasets used previously in the literature. The SVM searchlight produces results comparable to the GLM for the spatial domain. In the temporal domain, the SVM searchlight was applied to a publicly available dementia dataset, and identified prominent novel regions such as the frontal cortex and pre-motor cortex which did not appear in the earlier study.

Index Terms— fMRI imaging, analysis, classification, support vector machine, machine learning, high-performance computing

1. INTRODUCTION

Research investigating the analysis of fMRI images has grown dramatically in the past decade. This field is dominated by techniques using the general linear model, known as the GLM approach. This approach is useful and has resulted in many insights into brain function. However, it is limited in that it is a *linear* technique, and hence may fail to capture many brain processes which are inherently non-linear. Hence an important research direction is to explore the advantages that alternate techniques that employ non-linear methods can provide. In this paper we focus on one such technique based on the support-vector machine (SVM) formalism.

SVMs have been successfully applied in a number of domains including biometrics, image retrieval and medical imaging. We propose the use of local SVMs operating in a

“searchlight” manner, in that a local brain neighborhood is used to provide training data in a suitably low-dimensional space. We train a supervised classifier based on local fMRI activity, where the training labels are derived from the experimental protocol. We ask the question: which brain areas can predict the experimental protocol, given the spatio-temporal activity within that area? This approach agrees well with the current thinking on computational in cortical networks in the brain, wherein there is a dense lateral network that facilitates computation within local brain regions [1].

The fMRI image undergoes a domain decomposition step, which operates either in the spatial domain, temporal domain, or a combination of spatial and temporal domains. This step creates the inputs that are used to train the SVM. For the spatial domain decomposition, we use a sliding cubic window of size $N \times N \times N$ voxels. For a temporal domain decomposition, we use non-overlapping temporal windows at each voxel of length T time steps. For a spatio-temporal decomposition, we use a sliding window of $N \times N \times N \times T$ samples.

A significant challenge in applying the spatial domain decomposition is that it creates a heavy computational overhead, as a classifier has to be built for each voxel in the brain. We describe a method to overcome this challenge by applying high-performance computing through the use of parallel programming techniques. Specifically, we use OPEN-MPI, a language that implements a message-passing interface.

The main contributions of this paper are to show that the SVM methodology can be used to process fMRI images advantageously over the existing GLM approach, and to provide efficient computational methods for doing so.

2. BACKGROUND

Kriegeskorte *et al.* [2] used a spherical searchlight approach to process fMRI images. However, their approach used linear multivariate analysis, in contrast to the non-linear SVM approach proposed by us.

Haynes *et al.* [3] extend the approach in [2] by applying an SVM within each searchlight location. They use the data within a searchlight to distinguish between two experimen-

tal conditions. They deploy an averaging process within the searchlight to reduce the dimensionality of the problem. Similarly Mourao-Miranda *et al.* [4] use pre-processing via PCA or singular value decomposition to perform dimensionality reduction before applying an SVM. In contrast, we do not employ such an averaging, and use the original data within each searchlight. This permits our method to be more sensitive to the brain activity in each voxel. As pointed out by Kriegeskorte, [2], fine-scale patterns of activity may contain neuroscientifically relevant information.

Though this increases the dimensionality of the classification problem, we present effective parallel programming techniques to ameliorate this burden. Furthermore, our results show that the regions of high SVM classification accuracy are not noisy, and do not have a salt-and-pepper appearance, which is one of the processing outcomes to avoid [2].

The approach in our current paper is similar to that used by Xiao *et al.* [5]. The main difference is that Xiao *et al.* use an information measure based on the norm of the SVM weights after training. In contrast, in the current paper, we use the cross-validated classification accuracy. Furthermore, the method in Xiao *et al.* was used on 2-D images, whereas we have extended the method to analyze 3-D images in the current paper.

2.1. The support vector machine

We provide a brief description of SVMs [6]. Given a two-class data set $\{\vec{y}_i, c_i\}$, where \vec{y} represent the data and $c = \{-1, 1\}$ the classes, the objective of the SVM is to minimize the norm of the weight vector associated with the separating hyperplane, constrained to produce a good separation between the classes. It solves the following optimization:

$$\min: \|\vec{w}\| + \lambda \sum_i \epsilon_i, \text{ subj. to: } c_i(\vec{w} \cdot \vec{y}_i - b) \geq 1 - \epsilon_i \quad (1)$$

where ϵ is a misclassification or slack variable, λ a constraint on misclassifications, and b is an offset constant that centers the data.

3. METHODS

We use spatial and temporal decompositions as follows.

3.1. Spatial decomposition

In the spatial technique, we take a cubic window centered at a voxel located at (x, y, z) , as shown in Fig. 1. The size of the window is $N \times N \times N$ for the sake of convenience, though other window geometries can be used. The data within the searchlight is then vectorized to form a vector of size $N^3 \times 1$, denoted by V . At each time instant, we can create a set of labels that are related to the experimental protocol or subject responses. For instance, we may observe the brain activity following a stimulus presentation, say S_1 . Then the label associated with the vector V is S_1 . The labels L can be drawn

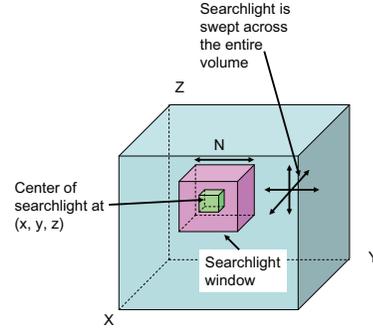


Fig. 1. This figure shows spatial decomposition, where a searchlight window is swept across the entire brain volume.

from the following set $L \in \{S_1, S_2, S_3\}$ for a three-stimulus experiment. The pair $\{V, L\}$ constitutes a training instance. For a given searchlight location, we can obtain multiple such pairs, at different time instants, which creates a set of data that can be used to train a supervised classifier such as the SVM. This searchlight window is then swept across the entire volume of the fMRI scan, and we can obtain the classification accuracy at each voxel.

We use cross-validation to obtain an accurate estimation of classification accuracy. This involves splitting the data into N folds, such that the data in one fold is used for testing, and the data in the remaining folds is used for training.

The intuitive interpretation of this paradigm is that a high classification accuracy implies that the local brain activity within the searchlight is sufficient to predict the stimulus condition. Typically, brain regions that are actively involved in a given task will produce high classification accuracies.

3.2. Temporal decomposition

In the temporal decomposition, we create a training vector for each voxel as shown in Figure 2. For the sake of specificity, we use the experimental protocol in the paper by Buckner [7], where there are two types of stimulus conditions, indicated by labels S_1 and S_2 . S_1 refers to a one-trial condition with a single visual stimulus, and S_2 refers to a two-trial condition with two consecutive visual stimuli. The duration of each condition is 8 time steps as shown. We form a vector, V of size 8×1 . Each vector V is associated with a label $L \in \{S_1, S_2\}$. We form pairs $\{V, L\}$ for each voxel as the training input to the classifier. We create 14 such labeled pairs.

We can also use a combined spatio-temporal decomposition, though we do not present results with such an approach due to space limitations.

3.3. Computational Challenges

Since we are using N -fold cross-validation, we have to train the SVM classifier N times at each voxel in the fMRI scan. This poses a heavy computational burden when the dimensionality of the dataset is high, and using a single CPU to

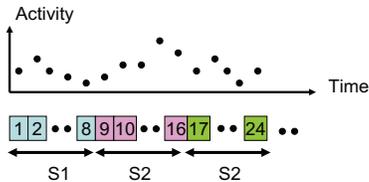


Fig. 2. This figure illustrates the temporal domain decomposition at each voxel.

perform this is not feasible. We have developed a solution to this problem using high-performance computing. The problem can be decomposed and solved independently per searchlight location. This allows the computation to be parallelized. We have performed this parallelization using MPI (Message Passing Interface standard).

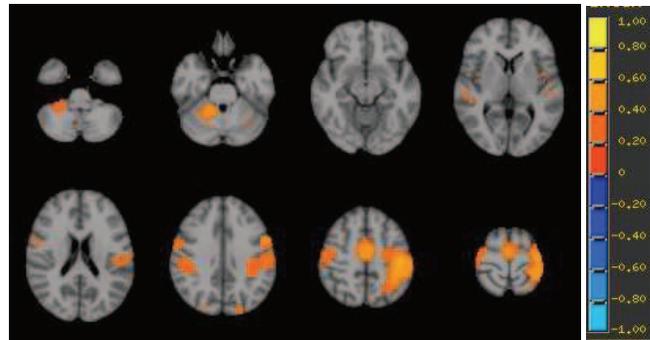
The main issue is load-balancing, and we solve this problem by using a work-crew scheduling algorithm presented in [8]. This involves a master that keeps assigning remaining work to crew members who are available.

We implemented our algorithm on an IBM Blue Gene/L supercomputer, using 1024 processors, and were able to run the SVM spatial searchlight for a single subject in the fMRI dataset cited in [9] in approximately 1 hour. This computation would have taken more than a month on a single CPU machine. Other platforms can also be used, such as clusters. The code is very portable, and MPI is available for use on multiple platforms through implementations such as OPEN-MPI.

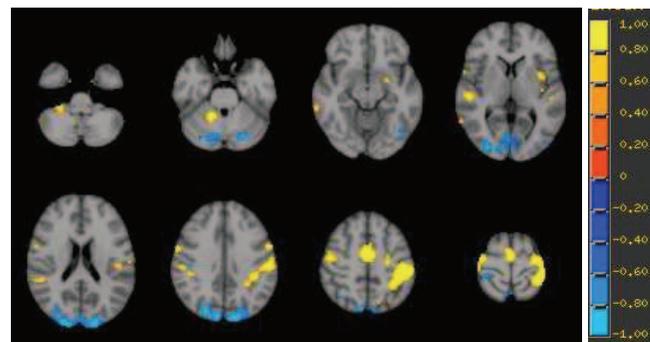
4. RESULTS AND DISCUSSION

In all the results shown, we used the same kernel size and the default parameter settings in the SVM package¹. Figure 3 shows the result of applying the spatial SVM searchlight to finger-tapping fMRI data used in a prior study [9]. Here, subjects were cued with three types of stimuli: auditory, a small visual stimulus, or a large visual stimulus. Following the stimulus, they were instructed to tap their fingers. This was followed by a period of rest. We used a spatial window of size $3 \times 3 \times 3$ voxels to create training vectors, $V(x, y, z)$ where V is of dimensionality 27×1 . Here x , y and z denote the coordinates of the center of the searchlight window. The supervised class labels, L correspond to the periods of finger tapping, denoted by S_1 and rest, denoted by S_2 . Thus, the label $L \in \{S_1, S_2\}$ for each time instant. This spatial window is moved across the entire brain. We perform 5-fold cross-validation to determine the classification accuracy at each searchlight location. The resulting map of classification accuracies for the first subject is shown in Fig. 3(A). Since the chance accuracy is 50%, accuracies higher than this can be considered meaningful. We use a threshold of 65% classification accuracy to depict the results in Figure 3. This

¹LIBSVM, <http://www.csie.ntu.edu.tw/~cjlin/libsvm>



(A)



(B)

Fig. 3. (A) The SVM prediction accuracy for the finger-tapping dataset [9]. Here we show the classification accuracy for the first subject. In the color scale, the value 1 represents a classification accuracy of 100%. (B) The GLM activation map for the same subject is shown. In the color scale, the value 1 represents a z-score of 10, and the value -1 represents a z-score of -10, which denotes de-activation.

threshold can be related to a p -value through the binomial distribution, as shown in [10].

As this result shows, the auditory, motor and visual cortices are implicated in the task, which agrees with the stimulus and response present in the experimental design. For the sake of comparison, we show the GLM activation map in Figure 3(B). There is a good overlap between the two maps. Furthermore, the SVM is able to capture both activation and de-activation present in the GLM map, such as in the de-activated visual areas of Figure 3(B).

Next, we show the result of applying the temporal SVM searchlight on a dementia dataset [7] which is available on the website *fmridc.org*. In this experiment, subjects were presented with two types of visual stimuli: either a single stimulus, denoted by S_1 , or two successive stimuli S_2 . Subjects were asked to press a button each time they saw the stimulus. We use 8 time points following the presentation of the stimulus to create the training vectors $V(x, y, z)$. The dimensionality of V is 8×1 . The supervised class labels are $L \in \{S_1, S_2\}$.

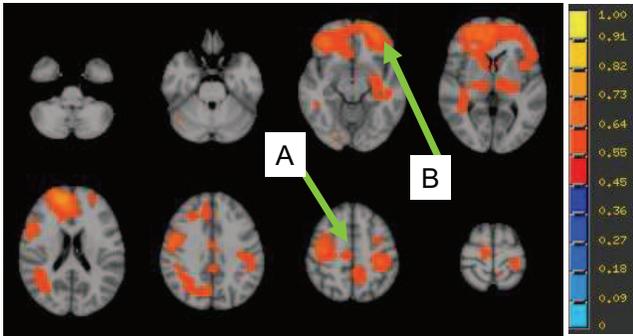


Fig. 4. SVM classification accuracy map for a single subject. Note the areas in the pre-motor cortex, and the frontal cortex. In the color scale, the value 1 represents a classification accuracy of 100%.

We use 5-fold cross validation to determine the accuracy of an SVM classifier at each voxel in the brain. The resulting accuracy is shown in the following figures. The chance accuracy is 50%, so accuracies higher than this are meaningful.

In Fig. 4 we show the accuracy map for the first subject. We observe significant accuracy in the pre-motor area, Brodmann Area 6, indicated by the label “A”. We also observe regions of high accuracy in the frontal cortex, indicated by the label “B”.

The identification of the pre-motor and frontal cortex are certainly reasonable, as one would expect executive function to be involved in the task, and to distinguish the two types of stimuli.

The significant aspect of the findings in Fig. 4 is that regions such as the pre-motor and frontal cortex were not identified in the original study by Buckner *et al.* [7]. Rather, they identified the visual and motor cortices as showing heightened activity. This implies that by using the SVM, we are able to generate alternate, plausible hypotheses for the brain regions that are implicated in a given task. This should provide a richer set of alternate mechanisms to neuroscientists in their quest to understand brain function.

One of the drawbacks of using the SVM searchlight approach in this paper is that it reduces the effective resolution of the active regions identified by the size of the searchlight window. This is because any voxel within the window that is highly predictive will improve the classification accuracy of the whole window. However, given the noisy nature of most fMRI measurements, such smoothing is typically employed.

5. CONCLUSION

In this paper we have presented a flexible approach for fMRI image analysis, called the spatio-temporal SVM searchlight. The bulk of the methodology to analyze fMRI images is based on GLM. In this paper we present an method using SVMs, which provides a mechanism to use non-linear kernels. This constitutes a powerful method with the ability to probe non-

linear relationships in the data.

We overcome computational hurdles through the use of parallel-processing and high-performance computing techniques. We have obtained encouraging initial results of applying this method to two different sets of experimental data involving different protocols. These results need to be extended and obtained for larger studies such as cross-subject studies, which will establish a role for the SVM searchlight.

Acknowledgements

We are grateful to the reviewers for their comments.

6. REFERENCES

- [1] D. D. Stettler and *et al.*, “Lateral connectivity and contextual interactions in macaque primary visual cortex,” *Neuron*, vol. 36, no. 4, pp. 739–750, 2002.
- [2] Nikolaus Kriegeskorte, Rainer Goebel, and Peter Bandettini, “Information-based functional brain mapping,” *PNAS*, vol. 103, no. 10, pp. 3863–3868, 2006.
- [3] J. Haynes and *et al.*, “Reading hidden intentions in the human brain,” *Current Biology*, vol. 17, pp. 323–328, 2007.
- [4] Mourao-Miranda and *et al.*, “Dynamic discrimination analysis: a spatial-temporal svm,” *Neuroimage*, vol. 36, no. 1, pp. 88–99, 2007.
- [5] Y. Xiao, A. R. Rao, G.A. Cecchi, and E. Kaplan, “Improved mapping of information distribution across the cortical surface with the support vector machine,” *Neural Networks*, vol. 21, no. 2-3, pp. 341–348, 2008.
- [6] Vladimir N. Vapnik, *Statistical Learning Theory*, Wiley-Interscience, 1998.
- [7] R. Buckner and *et al.*, “Functional brain imaging of young, nondemented, and demented older adults,” *J. Cog. Neurosci.*, vol. 12:Suppl. 2, pp. 24–34, 2000.
- [8] William Gropp, Ewing Lusk, and Anthony Skjellum, *Using MPI*, MIT Press, 1999.
- [9] V. M. Eguíluz and *et al.*, “Scale-free brain functional networks,” *Physical Review Letters*, vol. 94, pp. 018102, 2005.
- [10] F. Pereira, T. Mitchell, and M. Botvinick, “Machine learning classifiers and fMRI: A tutorial overview,” *NeuroImage*, vol. 45, no. 1, Supplement 1, pp. S199 – S209, 2009, Mathematics in Brain Imaging.