fMRI Analysis via One-class Machine Learning Techniques

David R. Hardoon
University of Southampton, ECS, ISIS Research Group
Southampton, UK SO17 1BJ. drh@ecs.soton.ac.uk

Larry M. Manevitz
University of Haifa, Department of Computer Science
Haifa, Israel 31905. manevitz@cs.haifa.ac.il

Abstract

We show how one-class compression Neural Networks and one-class SVM can be applied to fMRI data to learn the classification of brain activity associated with a specific motor activity. For comparison purposes, we use two labeled data and see what degree of classification ability is lost compared with the usual two-class SVM.

1 Introduction

Functional magnetic resonance imaging (fMRI) allows the carrying out of specific non-invasive studies within a given subject while providing an important insight to the neural basis of brain processes. Neurons, which are the basic functional unit of the brain, consume a higher level of oxygen when active, hence blood with a higher level of oxygenation is supplied to those active neurons. fMRI makes indirect use of this effect by detecting areas of the brain which have an elevated consumption of oxygen. This effect can be used to identify areas of the brain associated with specific functions.

The current methodology used to identify such regions is to compare, using various mathematical techniques, the elevation of oxygen consumption during a task with that used during a resting state. [Mitchell *et al.*, 2004] applied machine learning techniqes to this problem, when considering the classification of the cognitive state of a human subject. Thus, in order to determine the elevation of oxygen consumption during a task, images acquired during a resting state are required. In order to keep the alternation between activity, a reference time-course is needed, where the resting and active states are embedded. In this paper we further consider the problem of identifying fMRI scans that have only been acquired during the "active" state, i.e. scans acquired during the duration when the human subject has performed the given task. The basic intuitions are that, if available, two-class classification should perform better; although not always. However, as is the case under consideration here, often we have some reasonable sampling of the positive examples; i.e. the distribution of positive examples can be estimated; while the negative examples are either non-existent or episodic; i.e. not necessarily representative.

Obtaining good results under this assumption is known to be quite challenging; nonetheless it is often the most realistic assumption. For the fMRI classification described above, this problem is particularly non-trivial as we expect the data to be of very high dimension and extremely noisy, as the brain concurrently works on many given tasks. It is also quite natural to assume that there is only representative data of the task of interest; and not necessarily representative data of the negation of this task thus making the one-class learning techniques appropriate. In this work, we use two major one-class learning techniques - "bottleneck" or compression neural networks [Manevitz and Yousef, 2001] and a common version of the one-class Support Vector Machine (SVM) [Scholkopf *et al.*, 1999]. We point out that we use the entire brain slice, with no pre-filtering - i.e. the data is the entire slice, labeled with the task.

2 Experiments

The fRMI scans are of a volunteer[1] flexing their index finger on the right hand inside a MR-scanner while 12 image slices of the brain were obtained from a T2*-weighted MR scanner. The time-course reference of the flexing is built from the subject performing a sequence of 20 total actions and rests consisting of rest, flex, rest, $\ldots$ flex. Two hundred fMRI scans are taken over this sequence; ten for each action and rest. The individual fMRI images are dicom[2] format of size 128×128. Each image is labelled as either 1 (active) or -1 (inactive).

Thus, in our data we have 100 positive and 100 negative images for each of the 12 slices. For the bottleneck neural network 80 positive samples were chosen randomly and presented for training and 40 samples, consisting of the remaining 20 positive and 20 random negative samples, were used for testing. This experiment was redone with ten independent random runs. The limitation to 20 negative samples out of a possible 100 was chosen to keep the testing fair between the positive and negative classes. We manually cropped the non-brain background from the scans; resulting in a slightly different input/output size for each slice of about $8,300$ inputs and outputs. The compression percentage arising from the bottleneck was chosen by experimenting with different possible values. A uniform compression of about 60% gave the best results for the hidden layer. The irrelevant (non-brain) image data was cropped for each slice resulting in a slightly

[1] Provided by Ola Friman [Friman, 2003].

[2] For information regarding dicom see http://medical.nema.org/

different input/output size for the network for each slice. The network was trained to the identity using 20 epochs on the above chosen data. Following training the network was used as a classification filter, with an input value being classified as positive if the error level was lower or equal to a threshold chosen heuristically from training and classified as negative. We used the same protocol in a one-class SVM. Additionally, we used the two-class SVM where we randomly selected 160 training images and the remaining 40 for testing. This was also repeated 10 times.

2.1 First Experimental Results

The obtained results are an average over all the slices. Each slice was averaged over 10 repeats where in each repeat a random split of training testing was selected. Both SVM classifiers were used in their default setting as set by the OSU-SVM 3.00 package[3]. In addition, the two-class SVM was used with the unnormalised data as we have experimentally found that with normalised data the overall results were significantly worse. In Table 1 we are able to observe that while the one-class SVM performs better with the RBF kernel, the two-class SVM is better with the linear kernel. In Table 2 we compare the one-class to two-class techniques. As initially expected we are able to observe that the two-class approaches outperform those of the one-class. The one-class SVM is slightly better then the bottleneck compression NN. We further analyse the statistics of the methods i.e. the separation of the classified samples to their true classes. In Table 3[4] we compute and show the statistics of the fMRI images of the Positive samples that were classified as positive, denoted as true-positive, and the positive samples that were classified as negative, denoted as false-negative and the statistics of the negative fMRI images samples that were classified as negative, denoted as true-negative, and those that were classified as positive, denoted as false-positive. We are able to observe in that the compression NN is able to find a higher rate of true-positive fMRI images then the one-class SVM and the two-class methods even though they have obtained an higher overall success rate. Also we observe with regard to the negative samples that the two-class methods perform better then the one-class. This is expected as the one-class methods make no use of the negative samples and eminently will have a lower ability in classifying it. Additional experiments have been performed on visual instead of motor tasks with very similar results.

Table 1: SVM Results (success).

Method	Linear kernel	RBF kernel
One-class SVM	$49.12\% \pm 0.86\%$	$59.18\% \pm 1.47\%$
Two-class SVM	$68.06\% \pm 2.10\%$	$44.70\% \pm 1.12\%$

Table 2: Methods success results.

Method	Result on slices
BN - NN	$56.19\% \pm 1.26\%$
One-class SVM	$59.18\% \pm 1.47\%$
Two-class SVM	$68.06\% \pm 2.10\%$

Table 3: Methods Statistics

Method	True-Positive	False-Negative	std
BN - NN	78.96%	21.04%	$\pm 3.15\%$
One-class SVM	72.83%	27.17%	$\pm 1.98\%$
Two-class SVM	71.55%	28.45%	$\pm 3.21\%$
	True-Negative	False-Positive	
BN - NN	33.42%	66.58%	$\pm 3.45\%$
One-class SVM	39.25%	60.75%	$\pm 3.25\%$
Two-class SVM	65.64%	54.46%	$\pm 3.02\%$

[3] OSU SVMs Toolbox http://www.ece.osu.edu/$\sim$maj/osu_svm/
[4] std stands for Standard Deviation

3 Conclusions

We have found that one class classification can be done, even with the "noisy" data and with the full slices of the brain scan. Comparable results (about 58% accuracy) were obtained under both one-class SVM and Compression-Based NN techniques. For future work we would further investigate automated feature reduction as it might be fruitful, compare scans of the same individual across different fMRI sessions. The work presented here was for one individual. We intend to compare training and classification across individuals.

Acknowledgments

This collaboration was supported by the Caesarea Rothschild Foundation, and by the *HIACS* Reseach Center, the University of Haifa. The first author is funded by the European project LAVA num. IST-2001-34405 and PASCAL network of excellent num. IST-2002-506778. We also thank Ola Friman of Harvard Medical School for his generous sharing of his data.

References

[Friman, 2003] Ola Friman. *Adaptive Analysis of Functional MRI Data*. PhD thesis, Linkoping Studies in Science and Technology, 2003.

[Manevitz and Yousef, 2001] Larry Manevitz and Malik Yousef. One-class svms for document classification. *Journal of Machine Learning Research 2*, pages 139–154, 2001.

[Mitchell *et al.*, 2004] T.M. Mitchell, R. Hutchinson, R.S. Niculescu, F. Pereira, X. Wang, M. Just, and S. Newman. Learning to decode cognitive states from brain images. *Machine Learning*, 1-2:145–175, 2004.

[Scholkopf *et al.*, 1999] B. Scholkopf, J. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson. Estimating the support of a high-dimensional distribution. Technical Report 99-87, Microsoft Research., 1999.