One-class Machine Learning Approach for fMRI Analysis

D R Hardoon
School of Electronics & Computer Science, University of Southampton, Southampton, UK
L M Manevitz
Department of Computer Science, University of Haifa, Haifa, Israel

Key words to describe this work: fMRI, SVM, Neural Networks, Machine Learning, One-Class

Key Results: One-Class Machine Learning techniques are applied to classify whether a subject is performing a task or not by looking solely at the raw fMRI slices of his brain.

How does the work advance the state-of-the-art?: Attains the ability to decode the 'state of mind' of an individual without the need of learning the non-active state of a task.

Motivation (Problems addressed): Current methodology for identifying active and non-active fMRI requires images acquired during a resting state.

Introduction

Functional magnetic resonance imaging (fMRI) is an imaging technique that can be used in principle to map different sensor, motor and cognitive functions to specific regions in the brain. fMRI allows the carrying out of specific non-invasive studies within a given subject while providing an important insight to the neural of basis brain processes. 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 techniques 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 this work 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 (Japkowicz 1999). 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.

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.

Method

Our primary technique for the one-class approach is the compression neural network method (Cottrell et. al 1988, Japkowicz et. al 1995, Manevitz & Yousef 2000). We apply a design of a feed-forward neural network where in order to accommodate the usage of only positive examples we use a "bottleneck" with assumption that the images are represented in a m dimensional space where we choose a three level network with m inputs, m outputs and k neurons on the hidden level, where $m > k$. Figure 1 gives a graphical example of the bottleneck network. This network is then trained using the standard back-propagation to learn the identity function on the sample example.

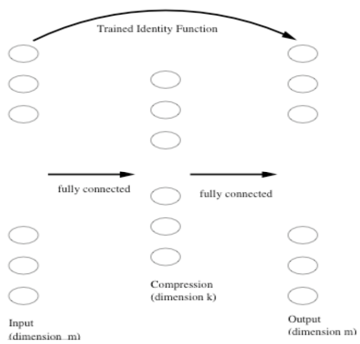

Fig 1. Bottleneck NN Architecture

We compare the compression neural network to the one-class Support Vector Machine (SVM) method (Scholkopf et. al 1999). Under this method, instead of separating positive and negative samples in the kernel feature space, as in standard (two-class) SVM, the origin is the only negative sample and therefore the method separates the positive samples from the origin via using relaxation parameters in SVM.

Experiments

The fRMI scans are of a volunteer flexing their index finger on the right hand inside a MR-scanner while 12 image slices of the brain were obtained from the 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, … flex. Two hundred fMRI scans are taken over this sequence; ten for each action and rest. A split of 80 positive scans for training and 20 positive and 20 negative for testing was used. 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. The NN was used with a 60% compression on the hidden layer, while both SVM classifiers were used in their default setting as set by the OSU-SVM 3.00 package with a linear kernel with C=1 and a radial based (RBF) kernel with $\gamma = 1$, the one-class SVM was used with relaxation parameter $\mu = 0.5$. The entire experiment was rerun in a separate session with the same individual. The two sessions are analyzed in Table 1.

	Session1	Session 2
NN	56.19%±1.26%	58.92%±2.03%
1-SVM	59.18%±1.47%	54.81%±1.18%
2-SVM	68.06%±2.10%	69.56%±4.12%

Table 1. Motor Data: Session 1 & 2 results

Conclusion

We showed that raw fMRI slices can be classified according to user tasks based solely on one-class information. The compression NN and the one-class SVM achieve about the same level of success. Detailed analysis of the errors (not presented here) indicates that the two one-class techniques make distinct errors, which suggests the possibility of combining the techniques in the future for better results. The compression NN by its nature performs a high-level feature extraction. Analysis of these features could lead to feature selection in the original data domain.

Acknowledgments

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

References

T. Mitchell, R. Hutchinson, R. Niculescu, F. Pereira, X. Wang, M. Just, and S. Newman, “Learning to decode cognitive states from brain images,” Machine Learning, vol. 1-2, pp. 145–175, 2004.

N. Japkowicz, “Are we better off without counter examples?” in Proceedings of the First International ICSC Congress on Computational Intelligence Methods and Applications, 1999, pp. 242–248.

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

L. Manevitz and M. Yousef, “Document classification on neural networks using only positive eaxmples,” in Proceedings of SIGIR, 2000, pp. 304–306

G. W. Cottrell, P. Munro, and D. Zipser, “Image compression by back propagation: an example of extensional programming,” Advances in Cognitive Science, vol. 3, 1988.