

Pittsburgh Brain Activity Interpretation Competition 2006

Methods Description

Brain Reading with Selective Training Data, Transduction, and Behavior Based Group Analysis

Project Abstract

Here we report our approach to the “Pittsburgh Brain Activity Interpretation Competition 2006” in which the goal was to predict the responses of three subjects who were watching video scenes during a functional magnetic resonance imaging (fMRI) study. Our work uses the continuous valued output from support vector machine classification (SVC) models to estimate a heterogeneous set of temporal “features” as rated by three subjects in response to audio/video stimuli. The primary challenge to this approach was to provide the SVC with appropriate class labels for each feature. As part of our approach, we attempted to fully utilize the unique properties of the data and the parameters of the competition, which included the existence of interdependencies between the feature ratings, a high degree of consistency in the ratings across subjects, and access to the target data (season 3) images. Primary conclusions from our work indicate that 1) the quality of the labels used for training the SVCs is more important than both the quantity of training data and image data variable reduction/selection, and 2) human behavior (in this case feature rating responses to multimedia stimuli) can be much more consistent across individuals than brain structure or cognitive strategy and the behavior’s temporal nature makes it inherently much easier to “align” than spatial brain data. We are excited by the potential of this work to provide a new approach to conceptualizing group analyses in the field of neuroimaging.

Introduction

Determining relationships that allow us to generalize beyond our direct experience is the primary goal of scientific exploration as well as a critical aspect of daily life. Often it is the availability and quality of observational data that directly impacts our ability to establish such relationships. In terms of understanding how the brain works, the development of functional neuroimaging (fMRI) coupled with ongoing improvements in MR systems and the continuing advances in the interrelated disciplines of statistics, data mining, machine learning, and computer vision has recently opened a new frontier for exploration. These exciting technologies are now allowing us to draw meaningful conclusions from vast amounts of data.

To date, there have been relatively few studies focused on the methodology of “brain reading” from fMRI data (the use of fMRI to predict brain states or mental representations related to spatial and temporal brain activation patterns). Fundamental methodological studies of neuroimaging classification include (Kjems, et al. 2002; Kustra and Strother 2001; LaConte, et al. 2003; LaConte, et al. 2005b; Mitchell, et al. 2004; Mourao-Miranda, et al. 2005; Shaw, et al. 2003; Strother, et al. 2004; Strother, et al. 2002). As argued by (Cox and Savoy, 2003), fMRI is well posed to study issues related to brain states based on its non-invasiveness and spatio-temporal (multivariate) qualities. Much interest in this topic was catalyzed by the desire to evaluate the evidence for a localized vs. distributed coding scheme for the (high-order)

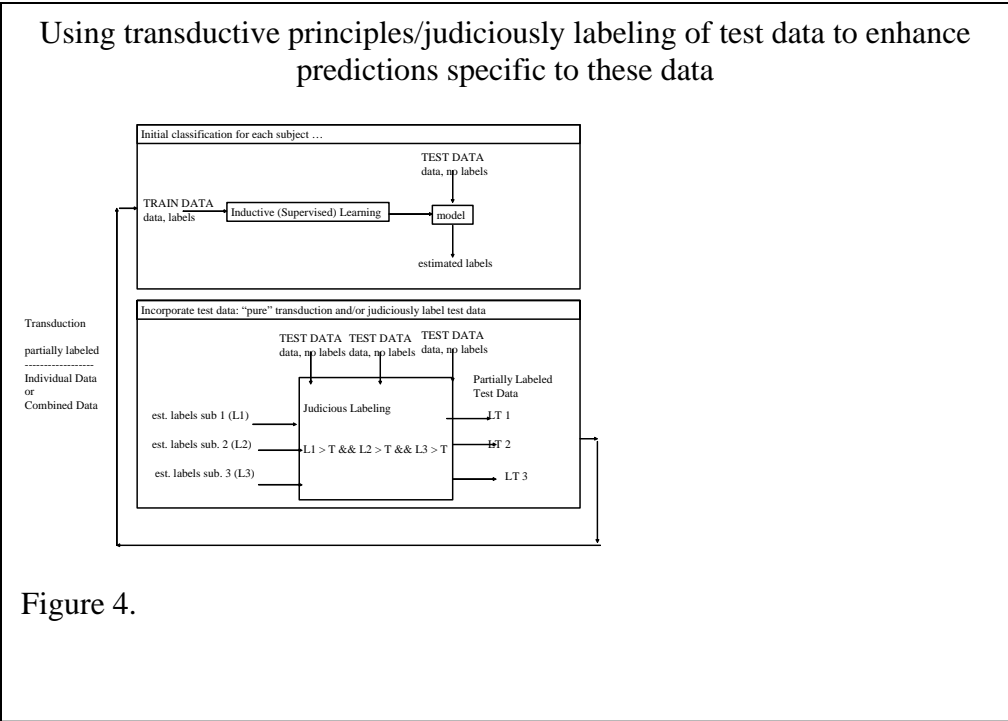
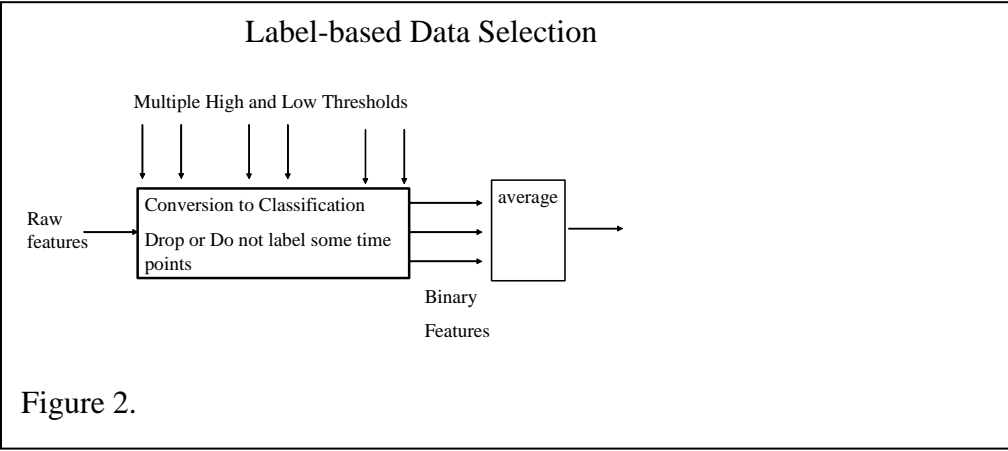
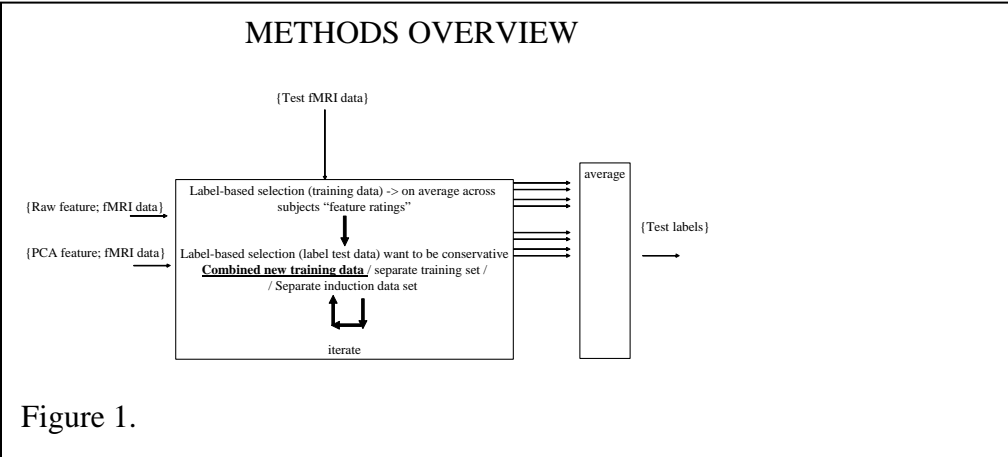
extrastriate visual cortex (Cox and Savoy, 2003; Downing et al., 2001; Hanson et al., 2004; Haxby et al., 2001; Ishai et al., 1999; Kanwisher et al., 1997; O'Toole et al., 2005). Consequently, there has been an amazing surge in cognitive neuroscientific interest in classification of brain states from fMRI data, with examples ranging from lie detection (Davatzikos et al., 2005), to unconsciously perceived sensory stimuli (Haynes and Rees, 2005) to behavioral choices in the context of emotional (fear) perception (Pessoa and Padmala, 2005), to memory of pictures (Polyn et al., 2005).

This report outlines our approach to meeting the brain reading goals of the “Pittsburgh Brain Activity Interpretation Competition 2006,” and assumes the reader is already familiar with its context. Briefly, the data competition consisted of data generated from 3 subjects who participated in 3 fMRI runs corresponding to clips from seasons 1, 2, and 3 of the television situation comedy, “Home Improvement” (multiple clips from the season’s episodes were shown during a run). After the fMRI session, the subjects provided ratings for several “features” associated with the clips. Our descriptions use the terminology of the competition, mixed with terminology from the SVM literature and fMRI “brain reading” literature. Thus “features” and the subjects’ “feature ratings” refer to the 29 categories that describe (some of) the dynamics in the T.V. episode clips and the subjects’ ratings about the nature of these dynamics, respectively. These feature ratings were used to “label” the fMRI data (not used here, but an equivalent concept in the fMRI literature that uses the generalized linear model is that the feature ratings could serve as covariates in a design matrix). To avoid conflicting uses of the word feature, we use ‘variable’, ‘voxel’, or ‘pixel’ to refer to the spatial image locations that are repeatedly measured over the course of an imaging run.

Method

The key ideas that we hoped to capitalize on was a two-stage strategy of using classification to localize important time regions for each feature and then apply regression to accurately estimate feature ratings. From the onset, another important concept was that of variable selection – obtaining the “best” subset of the high dimensional brain image data to perform inductive learning. If time permitted, we would then proceed to using multiple classification and regression models and explore ideas such as consensus, bagging, and boosting. No positive results from any of these ideas will be reported here – time did not permit, or no convincingly good results were easily obtained.

Much more successful were our remaining ideas, which include providing the SVC with the highest possible quality of the labeled data. That is, in which the relationship between the behavioral labels and brain images were most representative or meaningful. In addition, the feature rating labels, themselves, had interrelationships that could be exploited by methods such as systems identification, canonical correlation, or principal components analysis. Moreover these ratings were largely consistent across the three subjects. Finally, the competition included access to season three’s unlabeled imaging data. Since the goal of the competition was to predict these specific labels (as opposed to being able to predict any arbitrary data), we felt that concepts such as transductive learning [Vapnik] would provide another means of incorporating as much prior knowledge as possible into the modeling exercise.



Using SVD (Principal Components Analysis) to effectively utilized relationships between features

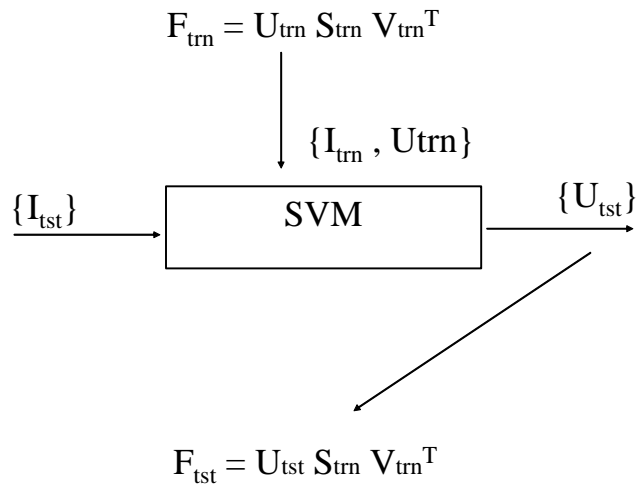


Figure 5.

We used the C-based SVMlight software (Joachims 1999) for classification, modifying it to handle fMRI image data. The data we used was the spatially normalized data and the convolved, downsampled feature ratings.

Results and Discussion

Feature Labels had very different properties, making appropriate labeling a challenge

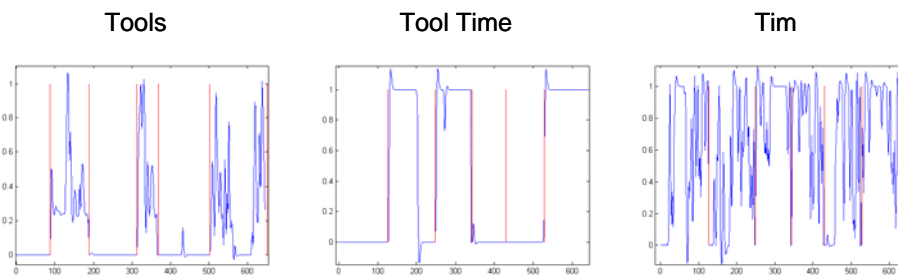


Figure 6.

Correlation structure of the features for each subject (seasons 1 and 2)

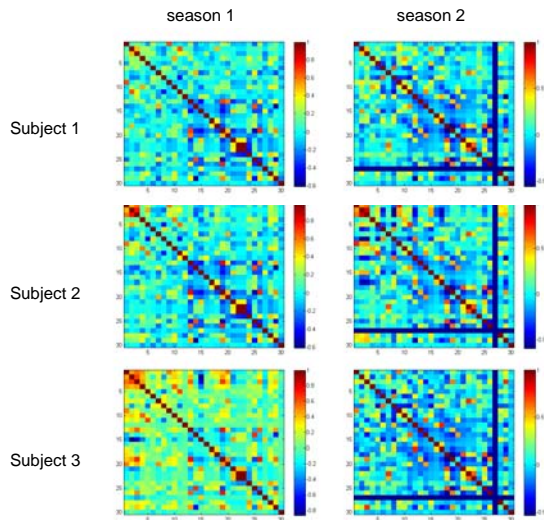


Figure 7. Note: Wilson is a missing feature in season 2

First 2 principal components account for a lot of the variance in the feature ratings and are very correlated between seasons 1 and 2

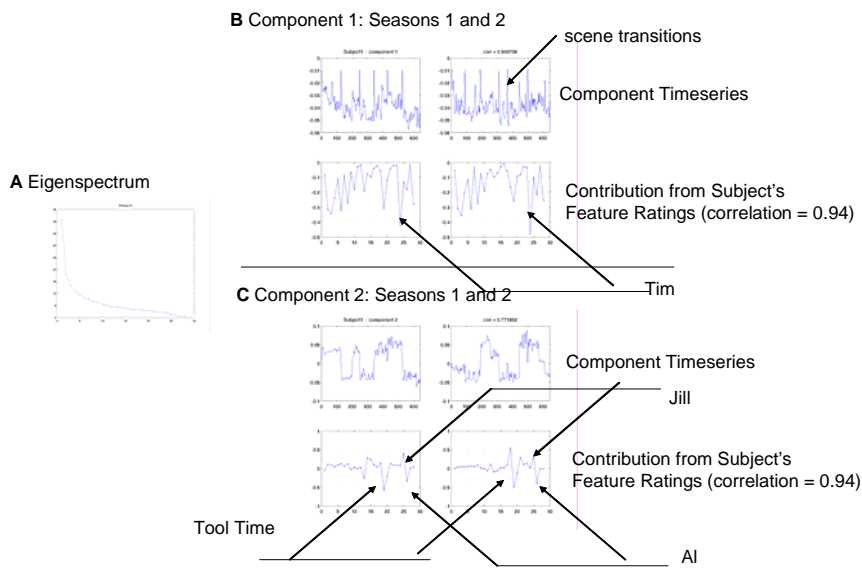


Figure 8.

References

- Cherkassky V, Mulier F. 1998. Learning From Data: Concepts, Theory, and Methods. New York: John Wiley & Sons, Inc.
- Cox DD, Savoy RL. (2003): Functional magnetic resonance imaging (fMRI) "brain reading": detecting and classifying distributed patterns of fMRI activity in human visual cortex. *NeuroImage* 19(2):261-270.
- Davatzikos C, Ruparel K, Fan Y, Shen DG, Acharyya M, Loughhead JW, Gur RC, Langleben DD. (2005): Classifying spatial patterns of brain activity with machine learning methods: Application to lie detection. *Neuroimage* 28(3):663-668.
- Downing P, Liu J, Kanwisher N. (2001): Testing cognitive models of visual attention with fMRI and MEG. *Neuropsychologia* 39(12):1329-1342.
- Hanson SJ, Matsuka T, Haxby JV. (2004): Combinatorial codes in ventral temporal lobe for object recognition: Haxby (2001) revisited: is there a "face" area? *NeuroImage* 23(1):156-166.
- Hastie T, Tibshirani R, Friedman J. 2001. The Elements of Statistical Learning: Data Mining, Inference, and Prediction. New York: Springer-Verlag.
- Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P. (2001): Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293(5539):2425-2430.
- Haynes JD, Rees G. (2005): Predicting the stream of consciousness from activity in human visual cortex. *Current Biology* 15(14):1301-1307.
- Ishai A, Ungerleider LG, Martin A, Schouten HL, Haxby JV. (1999): Distributed representation of objects in the human ventral visual pathway. *Proceedings of the National Academy of Sciences of the United States of America* 96(16):9379-9384.
- Joachims T. 1999. Making large-scale SVM learning practical. In: Scholkopf B, Burges C, Smola A, editors. *Advances in Kernel Methods - Support Vector Learning*. Cambridge, Mass: MIT Press.
- Kanwisher N, McDermott J, Chun MM. (1997): The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience* 17(11):4302-4311.
- Kjems U, Hansen LK, Anderson J, Frutiger S, Muley S, Sidtis J, Rottenberg D, Strother SC. (2002): The quantitative evaluation of functional neuroimaging experiments: Mutual information learning curves. *NeuroImage* 15(4):772-786.
- Kustra R, Strother S. (2001): Penalized discriminant analysis of [O-15]-water PET brain images with prediction error selection of smoothness and regularization hyperparameters. *IEEE Trans. Biomed. Eng.* 20(5):376-387.
- LaConte S, Anderson J, Muley S, Ashe J, Frutiger S, Rehm K, Hansen LK, Yacoub E, Hu X, Rottenberg D and others. (2003): The evaluation of preprocessing choices in single-subject BOLD fMRI using NPAIRS performance metrics. *Neuroimage* 18(1):10-27.
- LaConte S, Strother S, Cherkassky V, Anderson J, Hu XP. (2005c): Support vector machines for temporal classification of block design fMRI data. *NeuroImage* 26(2):317-329.
- Mitchell TM, Hutchinson R, Niculescu RS, Pereira F, Wang XR, Just M, Newman S. (2004): Learning to decode cognitive states from brain images. *Mach. Learn.* 57(1-2):145-175.

- Mourao-Miranda J, Bokde ALW, Born C, Hampel H, Stetter M. (2005): Classifying brain states and determining the discriminating activation patterns: Support Vector Machine on functional MRI data. *NeuroImage* 28:980-995.
- O'Toole AJ, Jiang F, Abdi H, Haxby JV. (2005): Partially distributed representations of objects and faces in ventral temporal cortex. *Journal of Cognitive Neuroscience* 17(4):580-590.
- Pessoa L, Padmala S. (2005): Quantitative prediction of perceptual decisions during near-threshold fear detection. *Proceedings of the National Academy of Sciences of the United States of America* 102(15):5612-5617.
- Polyn SM, Natu VS, Cohen JD, Norman KA. (2005): Category-specific cortical activity precedes retrieval during memory search. *Science* 310(5756):1963-1966.
- Shaw ME, Strother SC, Gavrilescu M, Podzebenko K, Waites A, Watson J, Anderson J, Jackson G, Egan G. (2003): Evaluating subject specific preprocessing choices in multisubject fMRI data sets using data-driven performance metrics. *NeuroImage* 19(3):988-1001.
- Strother S, LaConte S, Hansen LK, Anderson J, Zhang J, Pulapura S, Rottenberg D. (2004): Optimizing the fMRI data-processing pipeline using prediction and reproducibility performance metrics: I. A preliminary group analysis. *NeuroImage* 23:S196-S207.
- Strother SC, Anderson J, Hansen LK, Kjems U, Kustra R, Sidtis J, Frutiger S, Muley S, LaConte S, Rottenberg D. (2002): The quantitative evaluation of functional neuroimaging experiments: The NPAIRS data analysis framework. *NeuroImage* 15(4):747-771.
- Vapnik V. 1995. *The Nature of Statistical Learning Theory*. New York: Springer Verlag.