# Pittsburgh Brain Activity Interpretation Competition 2006

### Methods Description

### The Power of Linear Methods

## Project Abstract

To learn to predict complex input stimuli from large-scale patterns of brain activation in the EBC fMRI competition we combined simple linear methods, augmented by ridge regression, with voxel selection and both temporal averaging of voxel time series and across-subject averaging of the predicted stimuli. Our approach distinguishes itself by it's simplicity and, while somewhat basic, may ultimately be included as part of a more complex effort. In addition, the very noisy nature of the fMRI signal and the extreme undersampling of the joint probability distribution of stimuli and voxel activation patterns (especially in this very interesting case of natural stimuli) make complex methods hard to learn and prone to overtraining.

## Introduction

Perhaps the most fundamental and challenging pursuit in science is the attempt to understand the complex functioning of the human brain. How do we analyze and process sensory information to produce a coherent picture of the world around us? How do we use this picture to make decisions and to act? Where in the brain do we find language or logic? Ultimately these questions lead to the biggest question of all, the nature of consciousness itself. While these questions have occupied researchers for generations, it is only comparatively recently that advances in imaging technology have provided a quantitative window Into the spatiotemporal dynamics of large-scale cortical processing. While the authors of this submission are theoretical physicists by training, a substantial part of our interest in the EBC competition was as an introduction to the benefits and limitations in the fMRI signal analysis, the applicability of various computational techniques and their impact on our understanding of cognitive processes.

Our approach was centered upon linear models and our goal was for a best overall prediction. While it is highly likely that substantial information is present in nonlinear interactions, linear methods are very robust and computationally efficient. These last two concerns are especially relevant in fMRI signal analysis with tens of thousands of voxel time series and a much smaller number of samples (typically a thousand or so). At the very least, our results (typically 0.45 from the EBC score for predictions of the third movie) illustrate how far you can get with a simple approach

## Method

Our methods were fairly simple. We imported the raw functional DICOM data into AFNI. The temporal offset of each slice within a volume was corrected using the ‘tshift’ option in AFNI’s 3dvolreg script. This script was also used to motion corrected by aligning all volumes from one subject to a single “template” volume from that subject. Next, AFNI’s 3dAutomask script was used to automatically create a brain mask so that only voxels with intensities likely to have come from inside the brain would be included in subsequent analysis. The data was then imported into Matlab using the Princeton MVPA toolbox and Ziad Saad’s AFNI-Matlab library. Finally, the data were then z-scored so that each voxel’s time-course had zero mean an unit standard deviation. After a GLM analysis we typically had about 40,000 voxels for each subject.
Forty thousand voxels is still way too many to deal with so as another preprocessing step we selected (in order sorted by maximum linear correlation with the stimulus) only those voxels that were highly correlated with a particular stimulus. While the number of selected voxels varied, we typically used a few hundred. To reduce noise in the voxel time series, all selected voxels were passed through a simple linear boxcar temporal filter with widths ranging from 3 to 7 timepoints. The selected voxels were then linearly combined with weights that were choosen to minimize the mean error between stimulus and the weighted sum of voxels (basic linear regression) over the training movie and with a ridge penalty. With optimal weights we then made a prediction for the test movie. This process was carried out for each stimulus (1-13 of the basic features) and for each subject. While the EBC competition did not say explicitly that all subjects watched the shows in the same temporal order we found such a high degree of correlation between a feature from one subject on any given movie and the same feature from any other subject on the same movie that we assumed this to be true. Indeed, averaging each predicted feature across subjects proved to be a very useful technique. Predictions were made for all possible cross-validation pairings of the movies. For each stimulus, the small number of method parameters (ridge penalty, temporal filter width and number of selected voxels) was optimized for the best correlation across subjects and training. Occasionally we found it useful to spatially average a voxel with it’s nearest neighbors in three-dimensional space. This was particularly useful for food, arousal and music. All the analysis beyond the original preprocessing was done in MATLAB.

## Results and Discussion

Typical results during the training runs are summarized in the table included below. The correlations are higher than the test submission presumably due to a substantial amount of overtraining. The linear methods worked well and averaging (in time, space and across subjects) almost always worked. Many other good ideas did not. As the stimulus space is highly structured (in fact a principal components analysis (pca) of the full 30-dimensional feature vector reveals a much lower dimensionality of only about 15) we tried predicting pca components of the features instead of the features themselves. We also tried a variety of nonlinear methods but with no substantial improvement. Perhaps most disappointing was the fact that nothing that we already knew about the brain seemed to help in our analysis. Our method is most notable for it’s simplicity and the judicious use of averaging filters. Perhaps this is a reflection of the noise in the data or merely a reflection of our ignorance about what is important in the brain signal. We spent most of our time looking for stable and understandable results. We believe that the linear predictors now tell us. In the future, and with these stable results in hand we are now much more emboldened to look for subtle and nonlinear effects.

| feature | s1m2test | s1m1test | s2m2test | s2m1test | s3m2test | s3m1test | mean |
|---|---|---|---|---|---|---|---|
| MeanCor | 0.550 | 0.548 | 0.653 | 0.582 | 0.585 | 0.570 | 0.581 |
| Amusement | 0.309 | 0.512 | 0.507 | 0.542 | 0.628 | 0.664 | 0.527 |
| Attention | 0.269 | 0.331 | 0.713 | 0.432 | 0.273 | 0.498 | 0.419 |
| Arousal | 0.144 | 0.354 | 0.517 | 0.227 | 0.483 | 0.500 | 0.371 |
| Body Parts | 0.543 | 0.534 | 0.646 | 0.659 | 0.571 | 0.716 | 0.610 |
| Environmental Sounds | 0.447 | 0.432 | 0.692 | 0.592 | 0.511 | 0.604 | 0.546 |
| Faces | 0.701 | 0.664 | 0.748 | 0.717 | 0.732 | 0.634 | 0.699 |
| Food | 0.557 | 0.583 | 0.551 | 0.569 | 0.597 | 0.552 | 0.568 |
| Language | 0.801 | 0.769 | 0.786 | 0.688 | 0.800 | 0.630 | 0.746 |
| Laughter | 0.645 | 0.548 | 0.665 | 0.664 | 0.668 | 0.669 | 0.643 |
| Motion | 0.724 | 0.692 | 0.685 | 0.735 | 0.692 | 0.698 | 0.704 |
| Music | 0.738 | 0.607 | 0.720 | 0.589 | 0.693 | 0.602 | 0.658 |
| Sadness | 0.673 | 0.418 | 0.647 | 0.412 | 0.394 | -0.033 | 0.418 |
| Tools | 0.597 | 0.683 | 0.617 | 0.738 | 0.559 | 0.681 | 0.646 |