Sparse PCA a new method for unsupervised analyses of fMRI data

K. Sjöstrand[1,2], T. E. Lund[2], K. H. Madsen[1,2], R. Larsen[1]

[1]Informatics and Mathematical Modelling, Technical University of Denmark, Copenhagen, Denmark, [2]Danish Research Centre for Magnetic Resonance, Copenhagen University Hospital, Hvidovre, Denmark

Introduction

Exploratory analysis of functional MRI data aims at revealing known as well as unknown properties in a data-driven manner devoid of hypotheses on the time course of the hemodynamic response. This uncommitted approach usually precedes confirmatory modeling and may point to unexpected results that otherwise would be lost. Common approaches include clustering methods, principal component analysis (PCA) and in particular independent component analysis (ICA). ICA assumes that the measured activity patterns consist of linear combinations of a set of statistically independent source signals. Under favorable circumstances, one of more of these signals describe activation patterns, while others model noise and other nuisance factors. This work introduces a competing method for fMRI analysis known as sparse principal component analysis (SPCA). We argue that SPCA is less committed than ICA and show that similar results, with better suppression of noise, are obtained.

Methods

Standard PCA derives a set of variables by forming linear combinations of the original variables. The new variables are orthonormal and describe the main sources of variation in the data set. The projected data vectors are known as principal components (PCs) and are uncorrelated. The transformation can be written $\boldsymbol{Z}=\boldsymbol{XB}$ where $\boldsymbol{X}$ is the (n by p) data matrix, the columns of $\boldsymbol{Z}$ are the PCs, and $\boldsymbol{B}$ is the orthonormal loading matrix. Sparse PCA (SPCA) aims at approximating the properties of regular PCA while keeping the number of non-zero loadings small, that is, each derived variable is a linear combination of a small number of original variables. Implementing this in a principled manner has proven difficult, but a recent method [1] has shown great promise. The method poses regular PCA as a regression problem, and adds a constraint on the sum of absolute values for each loading vector. The constraint, known from the LASSO [2] regression technique, drives some loadings to exactly zero, while the others are adjusted to approximate the properties of PCA.

Figure 1: Synthetic signal without noise at $t=200$

Results

Two data sets with similar properties were used for validation, one synthetic and one in-vivo. Both data sets are related to a retinotopic study where a subject was visually stimulated using an 8 Hz reversing checkerboard expanding ring. Each expansion lasted 30 s. In-vivo data consisted of 381 volumes of 40 slices (64 by 64 matrix) acquired on a 3T Siemens Magnetom Trio scanner using a gradient echo EPI sequence. The volumes were rigid-body aligned using SPM2. A single axial slice sectioning the visual cortex was chosen for the analysis. The data was pre-filtered for effects of low-frequency drift, residual movement and physiological noise [3]. The synthetic data set was set to produce a pattern with similar properties as the expected retinotopic activation in the visual cortex. It consists of an area where intensities are governed by the function $sin(2\pi t/30+\delta)$ where t follows time and delta follows the lateral spatial position, see Figure 1. Gaussian noise was added with 12.5 times the amplitude of the signal.

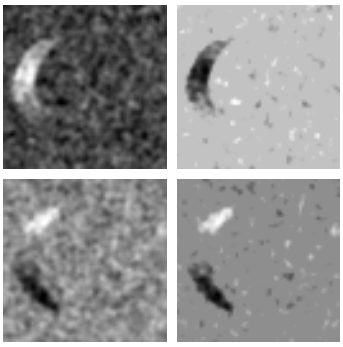

Figure 2: Independent component 1 & 2 (left) versus sparse component 1 & 2 (right).

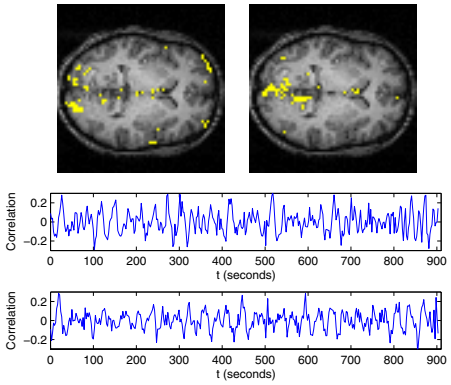

Figure 3: Thresholded activation of independent component 2 & 5 with their corresponding time courses.

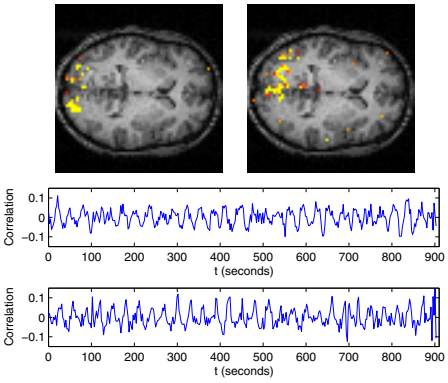

Figure 4: Sparse components 4 & 5. No thresholding of activation was performed. Time courses show clear paradigm correlations.

The synthetic data can be described using a linear model with two components since $sin(2\pi t/30+\delta)=cos(\delta)sin(2\pi t/30)+sin(\delta)cos(2\pi t/30)$. Figure 2 shows that both ICA and SPCA were able to identify this property to a large extent. The main difference between the methods lies in the noise suppression performed by SPCA. The ability to reconstruct the signal without noise was assessed for both methods, resulting in sum-of-squared (SSQ) differences of 613.8 and 311.7 for ICA and SPCA respectively. If the signal area was considered alone, the SSQ differences were 107.4 for ICA and 171.3 for SPCA. Visual assessment of the experiments on in-vivo data show slightly stronger paradigm response for SPCA. The resulting activation patterns seem to be more coherent and limited using SPCA.

Discussion

The SPCA method comes with a number of advantages over ICA. Following its resemblance with regular PCA, components have a natural ordering according to their variance. Moreover, since the derived loading vectors are strictly sparse, no thresholding is necessary to limit the spatial extent of the BOLD response. Instead, each loading vector can be individually controlled ranging from no sparseness, resulting in a regular PCA, to the all zeroes vector. Thresholding of the independent components is potentially misleading, and their statistical properties will no longer be preserved. SPCA is arguably a less committed approach than ICA since it uses second-order statistics, whereas ICA uses higher orders to achieve independence. The addition of a sparse prior to PCA is a sensible choice for fMRI analyses, since there are clinical and empirical evidence that the BOLD response has limited extents in both time and space. The SSQ differences show that SPCA did better at separating the noise from the signal, while ICA managed to model the actual signal more precisely. We conclude that SPCA and ICA has similar performance, but note that SPCA is more flexible, less committed and easier to interpret.

References

[1] H. Zou, T. Hastie, and R. Tibshirani, "Sparse principal component analysis", technical report, Statistics Department, Stanford University, 2004.

[2] R. Tibshirani, "Regression shrinkage and selection via the lasso", *Journal of the Royal Statistical Society - Series B Methodological* 58(1), pp. 267-288, 1996.

[3] T. E. Lund et al. "Non-white noise in fMRI: Does modelling have an impact? *NeuroImage in press*