SUBSPACE MODELS FOR FUNCTIONAL MRI DATA ANALYSIS

Ola Friman

Laboratory of Mathematics in Imaging, Department of Radiology
Brigham and Women's Hospital, Harvard Medical School

ABSTRACT

The models used for analyzing functional MRI (fMRI) data have profound impact on the detection of active brain areas. In this paper temporal and spatial linear subspace models for fMRI analysis are reviewed. General principles of how such subspaces should be constructed in order to obtain optimal detection performance are discussed and it is shown that customarily employed subspace models can be significantly improved upon.

1. INTRODUCTION

Functional MRI data consist of a temporal sequence of brain image volumes in which a dynamic pattern, referred to as the Blood Oxygen Dependent (BOLD) response, has been induced. As the BOLD response is widely accepted as an indirect indicator of brain activity, a fundamental question in fMRI analysis is: *Which voxels contain a BOLD response?* In order to answer this question, first consider each voxel in the spatio-temporal data set as an isolated entity. The time series of intensity values for each voxel can be decomposed as

$$\text{Voxel time series} = \underbrace{\text{BOLD response}}_{\mathbf{X}\boldsymbol{\beta}} + \underbrace{\text{Drift}}_{\mathbf{S}\boldsymbol{\theta}} + \text{Noise}. \quad (1)$$

To attribute variance in the observed voxel time series to each of these components, models of the BOLD response, drift and noise must be specified. A theoretically and computationally tractable class of models of the BOLD response and drifts is linear subspace models. Such models are extensively used in fMRI analysis. In Eq. 1, the basis functions spanning these model subspaces are gathered in the columns of the matrices $\mathbf{X}$ and $\mathbf{S}$. Signal variance in the observed time series is then explained by determining appropriate parameter vectors (or coordinates in the subspace relative to the bases in $\mathbf{X}$ and $\mathbf{S}$) $\boldsymbol{\beta}$ and $\boldsymbol{\theta}$. Once the signal variance has been distributed among the components in Eq. 1, a scalar test statistic λ, indicating the presence of a BOLD response in the time series, can be calculated. When λ exceeds a suitably chosen threshold we decide that there is a BOLD response present and declare the voxel active. Due to the noise component this decision will however be made under some amount of uncertainty. There is a risk that λ exceeds the threshold even if there is no BOLD response in the time series and vice versa. Figure 1 illustrates this situation. The shaded areas indicate the probability of making a decision error. To give an accurate answer to the question stated above, we would like to have a detection situation where the distributions in Fig. 1 are maximally separated, i.e. we want the probability of making an error to be as small as possible.

In this paper we discuss various linear subspace models for functional MRI data analysis and their impact on the distributions

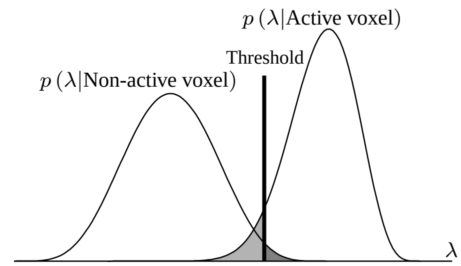

Fig. 1. Distributions of a test statistic λ, given the absence and presence of a BOLD response in the observed voxel time series.

in Fig. 1. We point to the intrinsic cone geometry in many linear subspace models and discuss how this geometry can be exploited for improving detection performance. In addition to the temporal subspace models of the BOLD response and drift components, spatial subspace models of local brain activity patterns are also considered.

2. SUBSPACE MODELS

In this section temporal and spatial subspace models for fMRI data analysis are discussed. To form a complete detection method, the models are to be used together with a detector, for example the General Linear Model (GLM) [1], Canonical Correlation Analysis (CCA) [2] or Generalized Likelihood Ratio (GLR) [3] detectors. These detectors have the same optimality properties, given the correctness of the chosen models. It should therefore be stressed that it is the model choices alone that dictate the detection performance.

2.1. Subspace models of the BOLD response

Linear subspace models of the BOLD response are widely used in fMRI analysis. An example is the truncated Fourier subspace, which is spanned by a set of sine and cosine functions. The Taylor subspace, spanned by a basic BOLD response shape and a number of derivatives with respect to the variables controlling the shape of this basic response, is another example. Figure 2 illustrates a blocked design model shape $x(t)$ and its temporal derivative $\dot{x}(t)$. In the two-dimensional subspace spanned by these functions we find (approximate) temporally shifted versions of the basic BOLD shape:

$$x(t+\tau) \approx x(t) + \tau\dot{x}(t). \quad (2)$$

This two-dimensional subspace model is hence able to account for small unknown delays in the BOLD response. Obviously, high-dimensional subspaces better account for variations in the BOLD response and therefore provide a more sensitive detection. In terms

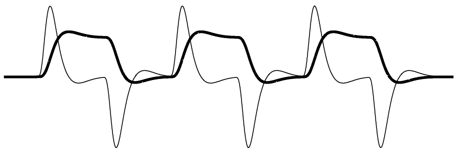

Fig. 2. A putative BOLD response model shape $x(t)$ (bold line) and its temporal derivative $\dot{x}(t)$.

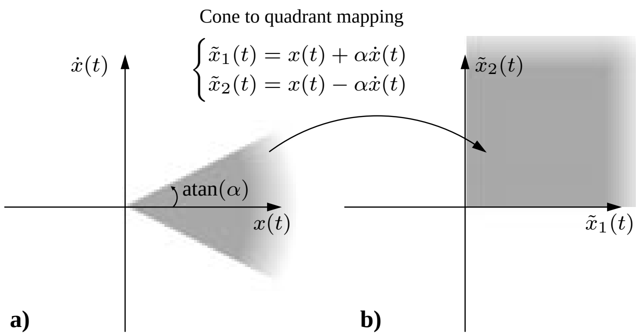

Fig. 3. The cone geometry exemplified in the Taylor subspace case. In a), plausible BOLD responses are located in a cone. This cone can be mapped to the first quadrant in a new basis, as shown in b).

of the distributions in Fig. 1, the higher the dimensionality of the subspace the farther to the right the "active" distribution will be located. However, with higher dimensionality the "non-active" distribution becomes broader since the risk of finding spurious fits to noise increases (also known as overfitting). Optimal detection performance is thus obtained with a low-dimensional subspace which captures the most important variations in the BOLD response. A general approach for finding such subspaces is to produce a large set of plausible BOLD response shapes and then find a set of basis functions that are able to reconstruct these shapes as well as possible, e.g. in a least square error sense [4, 2, 5].

In all subspace models discussed above, except for the truncated Fourier model, there is an intrinsic cone geometry that can be exploited to improve detection performance. To exemplify this cone geometry, consider the Taylor subspace basis functions in Fig. 2 and Eq. 2. The Taylor approximation in Eq. 2 holds reasonably well for small temporal shifts, which in the subspace correspond to points within a cone centered around the main basis function $x(t)$, as illustrated in Fig. 3a. Points outside this cone correspond to model shapes we do not accept as realistic. The degrees of freedom in the model can thus be reduced, without loosing detection sensitivity, by restricting the admissible BOLD response shapes to those within the cone. Methods for handling such constraints are however required. In [6] a GLR detector for this particular problem is derived. An alternative way to proceed is to introduce a new basis, constructed so that the cone is mapped to the first quadrant in this new basis, see Fig. 3. Detection can subsequently be carried out with the GLM approach with non-negativity constraints imposed on the regression weights [7]. A corresponding non-negative constrained CCA is presented in [8].

Enforcing the cone constraints is a way of incorporating prior

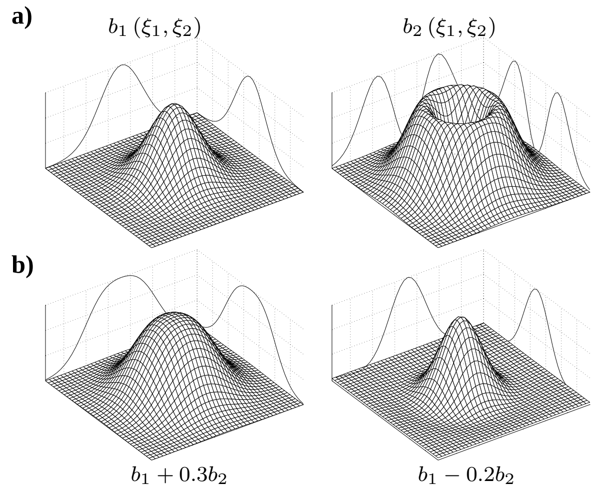

Fig. 4. a) A Gaussian at its derivative with respect to the width parameter. Note that the extension to 3D Gaussian functions is trivial. b) Two spatial shapes/filters found in the subspace spanned by the functions in a).

knowledge into the detection. Compared to a Bayesian approach, where knowledge enters as prior distributions on the parameters in the model, the cone approach is computationally orders of magnitudes faster due to the existence of closed form solutions.

2.2. Subspace models of local activity patterns

Even though brain activity is inferred by examining the temporal behavior of voxel intensities, detection performance is significantly improved if the spatial dimensions of the data are exploited. Since active brain areas extend over several voxels, a BOLD response can be detected with greater sensitivity if we first perform a spatial averaging. The averaging should however be matched to the local activity pattern so that only time series containing the BOLD response are averaged. To this end we need a model of plausible shapes of local activity patterns. It is tractable to use subspace models also for this purpose. Figure 4a illustrates two basis functions spanning such a subspace. The left basis function in Fig. 4a is a Gaussian smoothing kernel, and the right basis function is the derivative of this Gaussian with respect to the parameter determining its width. By a Taylor approximation, these two basis functions span a two-dimensional subspace containing approximate Gaussian shapes of different widths, see Fig. 4b.

Now assume that we wish to examine a certain voxel in the fMRI data set. We center the spatial basis functions over this voxel and average the time series in its neighborhood weighted according to the basis functions. Thus, to each spatial basis function there is an associated time series. These associated time series are to be linearly combined so as to extract a BOLD response, defined by a BOLD response model, as accurately as possible. In regression terminology, the time series associated with the spatial basis functions act as independent variables while the BOLD response model acts as dependent variable(s). Since all operations are linear and the order in which they are applied therefore is interchangeable, finding the combination of associated time series that best fits the BOLD response model can be interpreted as finding the spatial filter that best extracts a BOLD response. Note that the traditional

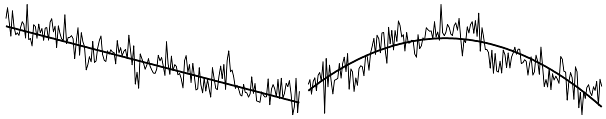

Fig. 5. Two examples of fMRI time series containing drifts, and the resulting fits of a polynomial drift subspace model.

Gaussian pre-smoothing of the fMRI images is equivalent to using a one-dimensional subspace spanned by the single Gaussian filter kernel used for the smoothing. However, by using more than one spatial basis function we can find better matches to local activity patterns and thus obtain better detection sensitivity. As for the modeling of the BOLD response discussed in the previous section, with more degrees of freedom comes an increased risk of constructing something similar to a BOLD response out of pure noise time series. Hence, the non-active distribution in Fig. 1 becomes broader when the subspace dimensionality increases. To alleviate this problem it is possible to exploit a cone geometry, exactly in the same manner as shown in Fig. 3. Other subspace models of local activity patterns and examples of the cone geometry are presented in [2].

2.3. Subspace models of drift components

Slowly varying temporal drifts are abundant in fMRI data, see Fig. 5. These drifts are not interesting per se, but they must still be accounted for in order to correctly distribute signal variance among the components in Eq. 1. Linear subspaces spanned by a set of polynomial or Discrete Cosine Transform (DCT) basis functions are commonly employed drift models. Subspaces spanned by large scale wavelets have also been proposed [9]. However, since the origin of these drifts are still poorly understood and we have very little *a priori* knowledge about them, it is difficult to specify a linear subspace model valid for every fMRI data set. As before, the model subspace should be of as low dimensionality as possible but still be able to account for the drifts. Unmodeled drift components end up as additional noise and we may loose detection sensitivity. Using an excess of basis functions for modeling the drifts results in inflated variances of parameter estimates, i.e. extra uncertainty regarding to which of the components in Eq. 1 the observed variance should be attributed. As a consequence the distributions in Fig. 1 become wider and the detection performance suboptimal. In both cases the estimated noise autocorrelation structure will be unnecessary complex, either because drift components remain in the data or because autocorrelation is induced when low frequency basis functions are removed. This has for example lead to the widespread belief that fMRI noise has a $1/f$ spectral characteristic.

With the above discussion in mind, the goal is to find a compact drift model subspace adapted to the specific data set at hand. Such a subspace can for example be estimated with a Maximum Likelihood (ML) approach; see [3] for a general description of the technique and [10] for an fMRI analysis application. The ML subspace model is spanned by the eigenvectors of the data covariance matrix obtained when each time series has been orthogonalized to the BOLD response model. Hence, the most variable components in the residual data are chosen as basis functions for the drift model subspace. In a similar manner, subspaces spanned by maximally autocorrelated temporal components found in the specific data set were proposed as drift models in [2]. Since drifts are well characterized by both high autocorrelation and energy, these latter subspace models capture drifts with greater efficacy than do subspaces spanned by polynomial or DCT basis functions.

3. RESULTS

3.1. Spatial subspace model and cone constraints

The established approach for detecting active voxels in fMRI data employs a multi-dimensional subspace model of the BOLD response and a one-dimensional subspace model of local activity patterns (in the framework of the GLM [1]). The improved detection performance offered by a multi-dimensional spatial subspace model together with cone constraints is here demonstrated. A synthetic putative BOLD response shape, similar to the one shown in Fig. 1, with slightly different delays in different voxels, was embedded in Gaussian shaped patterns of different sizes in a real fMRI data set acquired during a resting condition. Detection was then carried out with a Taylor subspace model of the BOLD response (Fig. 2) and a pre-smoothing with a Gaussian filter (i.e. a one-dimensional spatial model). A similar detection, but with a two-dimensional spatial subspace (Fig. 4) and cone constraints enforced, was also performed. For the record, the CCA detector was used in both cases because it can handle multi-dimensional temporal and spatial subspace models simultaneously [2], whereas the GLM detector only handles one multi-dimensional subspace. Again it should however be stressed that it is the model choice and not the detector choice that affects the result; the GLM detector is just a special case of the CCA detector. The upper panel in Fig. 6 shows the resulting maps of canonical correlations and the lower panel the result after a thresholding operation. The multi-dimensional spatial model and the cone constraints clearly provide a better detection.

3.2. Drift model

It is somewhat more difficult to evaluate drift models since we know very little about both the nature of the drifts (and other physiological phenomena) and the noise. However, with a correct drift model it is reasonable to hypothesize that the noise should be approximately temporally white and spatially uniform, at least within the brain. Under these assumptions the performance of different drift models can be compared. Polynomial, DCT, ML and maximum autocorrelation subspace models of different dimensionalities were fitted and removed from an fMRI data set acquired during a resting condition. The ML and maximum autocorrelation subspaces were adapted to this specific data set (not shown due to lack of space). The residual noise time series were then examined from autocorrelation, variance and the homogeneity of the variance across voxels (spatial heteroscedasticy) points of view, see Fig. 7. The residual noise autocorrelation was assessed with the Durbin-Watson statistic, which attains the value 2 when there is no autocorrelation. The average Durbin-Watson statistic over all within-brain voxels is plotted in Fig. 7a. Similarly, the average residual noise variance is plotted in Fig. 7b. Since the maximum autocorrelation and ML subspace models are tailored for removing autocorrelation and variance respectively, it is not surprising that these models perform best in these respects. Finally, in Fig. 7c the spatial homogeneity of the noise variance is examined. The conclusion drawn from the results in Fig. 7 (together with examinations carried out using other fMRI data sets) is that drift models

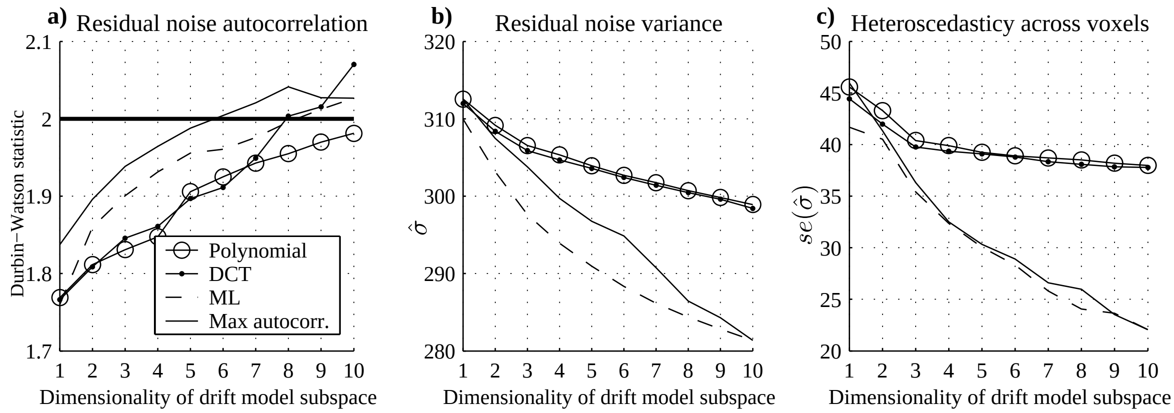

Fig. 7. a) The mean of the Durbin-Watson statistic over all within-brain voxels for different dimensionalities of the drift subspace model. The Durbin-Watson statistic attains the value 2 for white noise. b) The average residual noise variance. c) The standard error of the residual noise variance across voxels.

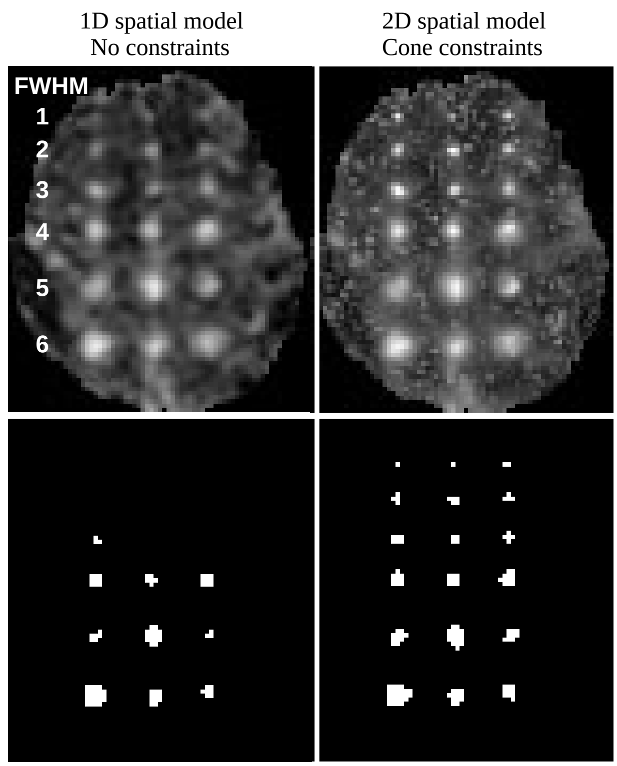

Fig. 6. Upper panel: Canonical correlation maps obtained with different subspace models. The size of the embedded activity patterns, given to the left, is in voxel units. The spatial subspace model was in both cases based on a Gaussian kernel with a width of 4 voxel FWHM. Bottom panel: The canonical correlation maps thresholded so that no truly non-active voxels are declared active.

adapted to the specific data set at hand more effectively captures drifts than do the customarily employed and more *ad hoc* polynomial and DCT subspace models.

4. REFERENCES

[1] K.J. Friston, P. Jezzard, and R. Turner, "Analysis of functional MRI time-series," *Human Brain Mapping*, vol. 1, pp. 153–171, 1994.

[2] O. Friman, *Adaptive Analysis of Functional MRI Data*, Ph.D. thesis, Linköping University, Sweden, 2003, Dissertation No 836, ISBN 91-7373-699-6.

[3] L. Scharf and B. Friedlander, "Matched subspace detectors," *IEEE Transactions on Signal Processing*, vol. 42, no. 8, pp. 2146–2157, 1994.

[4] C. Liao, K. Worsley, J. Poline, J. Aston, G. Duncan, and A. Evans, "Estimating the delay of the fMRI response," *NeuroImage*, vol. 16, no. 3, pp. 593–606, 2002.

[5] G. Hossein-Zadeh, B. Ardekani, and H. Soltanian-Zadeh, "A signal subspace approach for modeling the hemodynamic response function in fMRI," *Magnetic Resonance Imaging*, vol. 21, no. 8, pp. 835–843, 2003.

[6] S. Ramprashad, T. Parks, and R. Shenoy, "Signal modeling and detection using cone classes," *IEEE Transactions on Signal Processing*, vol. 44, no. 2, pp. 329–338, 1996.

[7] Å. Björck, *Numerical Methods for Least Squares Problems*, SIAM, Society for Industrial and Applied Mathematics, 1996.

[8] S. Das and P.K. Sen, "Restricted canonical correlations," *Linear Algebra and its Applications*, vol. 210, pp. 29–47, 1994.

[9] F. Meyer, "Wavelet-based estimation of a semiparametric generalized linear model of fMRI time-series," *IEEE Transactions on Medical Imaging*, vol. 22, no. 3, pp. 315–322, 2003.

[10] B.A. Ardekani, J. Kershaw, K. Kashikura, and I. Kanno, "Activation detection in functional MRI using subspace modeling and maximum likelihood estimation," *IEEE Transactions on Medical Imaging*, vol. 18, no. 2, pp. 101–114, 1999.