Reprint from NeuroImage 19:3 2003, p 837-845

Adaptive Analysis of fMRI Data

Ola Friman, Magnus Borga, Peter Lundberg and Hans Knutsson
Linköping University, Sweden

ABSTRACT

This article introduces novel and fundamental improvements of fMRI data analysis. Central is a technique termed constrained Canonical Correlation Analysis, which can be viewed as a natural extension and generalization of the popular General Linear Model method. The concept of spatial basis filters is presented and shown to be a very successful way of adaptively filtering the fMRI data. A general method for designing suitable hemodynamic response models is also proposed and incorporated into the constrained canonical correlation approach. Results that demonstrate how each of these parts significantly improves the detection of brain activity, with a computation time well within limits for practical use, are provided.

I. INTRODUCTION

This article introduces fMRI data analysis tools which improve the detection of brain activity on a fundamental level. The analysis approach should be easily assimilated by researches acquainted with the widely used General Linear Model (GLM) method, although terminology and notations are somewhat different. The employed analysis technique is Canonical Correlation Analysis (CCA) and an extension termed constrained CCA. CCA operates on two multidimensional data sets and is for this reason more general than the GLM, where one side is univariate (a voxel time series in the fMRI analysis context). CCA can therefore be seen as residing at the top of the hierarchy of regression methods.

At the core of this article is the improvement of detection performance. Two key concepts are sensitivity and specificity. Sensitivity is the ability to correctly identify the active voxels in the fMRI data while specificity is the ability to identify the non-active voxels. The aim is to maximize these quantities simultaneously by producing statistical parameter maps which serve as good bases for the ensuing classification problem, where voxels are declared as either active or non-active. While the classification problem is not touched upon in this article, issues such as models for the hemodynamic response and spatial filtering of the data are addressed.

The hemodynamic response evoked by a stimulus has in recent years been subject to much research and several models for its origin have been proposed (Buxton et al., 1998; Friston et al., 2000). For brain activity detection purposes, it is sufficient to capture the basic shape of the response and simplified models such as the reference timecourse indicating stimulus presentations convolved with a suitable impulse response is widely used. In the above models, the parameters describing the shape of the response enter nonlinearly and are therefore here referred to as nonlinear models. A number of reports observe variations in the shape of the hemodynamic response (Aguirre et al., 1998; Glover, 1999), for example between different subjects and even between different brain areas within the same subject. Hence, it is important for an analysis method to allow for such variations in order to maximize detection sensitivity. When using a nonlinear model it is not trivial to find the parameter values that best fit the model to an observed hemodynamic response. The variations should preferably be captured in a linear subspace spanned by a set of temporal basis functions. Parameters can then conveniently be algebraically determined in a least square error sense as in the GLM. An example of a linear model is the convolutive model (with fixed parameters in the impulse response) augmented with the temporal derivative. This model can account for small temporal delays in the response. Temporal delay is however not likely to be the only variation in the hemodynamic response. Variations in response width, overshoot and undershoot are also conceivable. By adding more temporal basis functions it is possible to capture these variations too. However, each additional basis function deteriorates the specificity and eventually the temporal basis set will be useless for discriminating active and non-active brain voxels. Put in another way, what we gain in sensitivity by adding more basis functions we may loose in specificity. Here we introduce a general method for generating compact linear subspaces with minimal loss of specificity.

Currently, most methods for detecting brain activity in fMRI consider a single voxel and the corresponding time series at the time. However, what makes the fMRI data analysis special and different compared to traditional time series analysis is the spatial context in which the time series exist. The spatial dimensions have however been sparsely exploited in fMRI analysis. Most commonly a plain Gaussian smoothing of the images is applied prior to the calculation of the statistical parameter maps. Such simplistic treatment of the data often introduce unnecessary blurring and does not make full use of the available spatial information. By generalizing the GLM into CCA, more sophisticated spatial filtering is achieved.

The outline of this article is as follows. First, a review of CCA is given. Then a general method for designing linear subspace models of the hemodynamic response is presented. Next, spatial filters suitable for fMRI analysis are introduced. Finally, these models and filters are used in a constrained CCA framework, resulting in an adaptive analysis of fMRI data.

II. THEORY

The methods employed in this article are based on *Principal Component Analysis* (PCA) and *Canonical Correlation Analysis* (Hotelling, 1936; Anderson, 1984). We assume some familiarity with PCA while an overview of CCA and constrained CCA are provided.

As previously stated, CCA operates on two multivariate random variables, here denoted $\mathbf{x} = [x_1, \ldots, x_m]^T$ and $\mathbf{y} = [y_1, \ldots, y_n]^T$. The question is how to make linear combinations of the variables

in $\mathbf{x}$ and $\mathbf{y}$,

$$x = w_{x_1} x_1 + \ldots w_{x_m} x_m = \mathbf{w}_x^T \mathbf{x}, \quad (1)$$

$$y = w_{y_1} y_1 + \ldots w_{y_n} y_n = \mathbf{w}_y^T \mathbf{y}, \quad (2)$$

so that the resulting scalar variables x and y correlate maximally. First it can be noticed that if $\mathbf{x}$ and $\mathbf{y}$ are one-dimensional, the ordinary Pearson correlation is obtained. If either $\mathbf{x}$ *or* $\mathbf{y}$ is one-dimensional we essentially have the GLM approach. Inserting Eq. 1 and Eq. 2 into the definition of correlation ρ between x and y yields

$$\rho = \frac{\mathbf{w}_x^T \mathbf{C}_{xy} \mathbf{w}_y}{\sqrt{\left(\mathbf{w}_x^T \mathbf{C}_{xx} \mathbf{w}_x\right)\left(\mathbf{w}_y^T \mathbf{C}_{yy} \mathbf{w}_y\right)}}, \quad (3)$$

where $\mathbf{C}_{xx}$ and $\mathbf{C}_{yy}$ contain the covariances between the variables within the $\mathbf{x}$ and $\mathbf{y}$ datasets respectively, while $\mathbf{C}_{xy}$ contains the between sets covariances. To find the maximum obtainable correlation we set the partial derivatives of Eq. 3 with respect to the weights $\mathbf{w}_x$ and $\mathbf{w}_y$ to zero and arrive in the following eigenvalue problems,

$$\mathbf{C}_{xx}^{-1} \mathbf{C}_{xy} \mathbf{C}_{yy}^{-1} \mathbf{C}_{yx}\, \mathbf{w}_x = \rho^2\, \mathbf{w}_x, \quad (4)$$

$$\mathbf{C}_{yy}^{-1} \mathbf{C}_{yx} \mathbf{C}_{xx}^{-1} \mathbf{C}_{xy}\, \mathbf{w}_y = \rho^2\, \mathbf{w}_y. \quad (5)$$

Thus, we find the optimal weight or regression vectors $\mathbf{w}_x$ and $\mathbf{w}_y$ as the eigenvectors belonging to the largest eigenvalues of the matrices $\mathbf{C}_{xx}^{-1} \mathbf{C}_{xy} \mathbf{C}_{yy}^{-1} \mathbf{C}_{yx}$ and $\mathbf{C}_{yy}^{-1} \mathbf{C}_{yx} \mathbf{C}_{xx}^{-1} \mathbf{C}_{xy}$. The eigenvalues of these matrices are equal and the largest is the squared maximum correlation. In practice, the covariance matrices must of course be replaced by their estimates.

In general, the regression weights in $\mathbf{w}_x$ and $\mathbf{w}_y$ will adopt both positive and negative values. Das and Sen (Das and Sen, 1994) have presented a procedure for constrained CCA, where the weights are restricted to be non-negative. If we for simplicity take $\mathbf{C}_{xy}$ to be the correlation matrix, we note that

$$\rho^2_{unconstrained} \geq \rho^2_{constrained} \geq \max_{i,j} \left([\mathbf{C}_{xy}]_{i,j} \right)^2. \quad (6)$$

Of course, if the unconstrained solution has only positive components in $\mathbf{w}_x$ and $\mathbf{w}_y$, this is the solution also to the constrained problem. At the other end, the smallest squared correlation we are guaranteed to obtain is given by the largest correlation between any of the input variables in $\mathbf{x}$ and $\mathbf{y}$, i.e. the largest element in $\mathbf{C}_{xy}$. In (Das and Sen, 1994) it is shown that the constrained solution equals an unconstrained solution to a modified CCA problem where one or several variables in $\mathbf{x}$ and $\mathbf{y}$ have been excluded. All possible deletions must be tested in order to find the solution, which gives the constrained CCA the unpleasant property of growing exponentially in the number of dimensions of $\mathbf{x}$ and $\mathbf{y}$. In the current work we use small dimensionalities of $\mathbf{x}$ and $\mathbf{y}$ and the constrained CCA solution is found by solving a limited number of small eigenvalue problems. Two important properties of the constrained CCA are that we are guaranteed to find the *global* optimum and that we can find this optimum algebraically, i.e. no iterative numerical search is required. It is shown how this constrained version of CCA significantly improves detection performance in fMRI analysis.

III. METHODS

Temporal Basis Functions

As stated already, for analysis purposes it is desirable to use a linear model of the hemodynamic response, i.e. a set of temporal basis functions of which a hemodynamic response shape can be constructed. This is mainly due to computational issues since closed form solutions can be found. Therefore, inherently nonlinear physiologically derived models of the hemodynamic response are not well suited for analysis aiming at detecting brain activity in the fMRI data. The question is how to find a linear model that captures the important variations in the hemodynamic response with as few temporal basis functions as possible. Assume that we have a nonlinear hemodynamic response model, for example a convolutive model or the Balloon model (Buxton et al., 1998). For each parameter in such a model there is a certain range of values that result in physiologically realistic shapes of the response. By varying the parameters randomly or systematically within these ranges, a large number of plausible response shapes $h_i(t)$ can be produced. We would like to find a set of temporal basis functions that can reconstruct these simulated responses as accurately as possible with as few basis functions as possible. This is exactly what a principal component analysis applied to the simulated responses achieves. Given the large number of responses generated by the nonlinear model, PCA finds "eigen"-timecourses $e_k(t)$ that minimize the mean square error $E\left[\sum_t \left(h\left(t\right) - \hat{h}\left(t\right)\right)^2\right]$ for a partial expansion to P components:

$$\hat{h}(t) = m(t) + \sum_{k=1}^{P} w_k e_k(t), \quad (7)$$

where $m(t)$ denotes the average simulated response. The "eigen"-timecourses $e_k(t)$ model the most significant deviations from the mean response. The P "eigen"-timecourses together with the average timecourse are thus good candidate basis functions for our linear subspace model of the hemodynamic response. For further reference, the notation introduced in the Theory section is adopted, $y_1(t) = m(t), y_2(t) = e_1(t), \ldots, y_n(t) = e_P(t)$, and the final response $y(t)$ is constructed as

$$y(t) = w_{y_1} y_1(t) + \ldots + w_{y_n} y_n(t). \quad (8)$$

P should be chosen small in order to keep a reasonable specificity. Here we chose $P = 1$, implying that two temporal basis functions will be used: the average response and the first principal "eigen"-response.

Even though we now have a compact linear subspace that captures the most important hemodynamic response variations, unrealistic responses can still be constructed, e.g. by setting $w_{y_1} = 0$ in Eq. 8 and thereby ignoring the average response. To counteract such constructions we utilize the possibility to find a *constrained* solution where all regression coefficients w_{y_i} are restricted to positive values. To this end, we first make a change of variables

$$\tilde{y}_1(t) = y_1(t) + \alpha y_2(t),$$
$$\tilde{y}_2(t) = y_1(t) - \alpha y_2(t), \quad (9)$$

where $0 < \alpha < 1$ is a user defined constant. Using $\tilde{y}_1(t)$ and $\tilde{y}_2(t)$ in a constrained analysis with non-negative weights w_{y_i} ensure that the average response $y_1(t)$ always is given a larger weight than the residual variations in $y_2(t)$ [1]. Hence, plausible response shapes are guaranteed. A guideline for choosing α is offered by examining to what extent the basis functions are needed in order to reconstruct the simulated responses we started with. An example of the above procedure is given in the Results section and the generalization to more basis functions is described in the Appendix.

Spatial Basis Functions

Just as it is possible to express the hemodynamic response model in terms of temporal basis functions, we can build spatial filters by combining spatial filter basis functions. Traditional fMRI analysis employs a fixed Gaussian filter for smoothing the images, which is homologous to using a single basis function for modelling the temporal hemodynamic response. While we have prior knowledge of the expected shape of the hemodynamic response, we usually have no prior information about the shapes of the activated brain regions. It can therefore be expected to be even more beneficial to use a set of filter basis functions of which filters matched to the actual shapes of the activated brain regions can be constructed. Suitable spatial basis functions are so-called *steerable filters* (Knutsson et al., 1983; Freeman and Adelson, 1991; Granlund and Knutsson, 1995), which are well known in the image processing community. The procedure for designing a set of 2D filter functions is illustrated in Fig. 1. Figure 1a shows a Gaussian smoothing filter $f(\mathbf{z})$ with a certain user defined size. The reader can replace this Gaussian filter with his/her favourite low pass filter for smoothing the fMRI data. In Fig. 1b the filter kernel space is partitioned into four parts, one isotropic central part $g_{iso}(\mathbf{z})$ and three oriented parts $g_i(\mathbf{z}), i = 1 \dots 3$. These functions are used for weighting the original filter and they sum to one in every point. The oriented weight functions $g_i(\mathbf{z})$ are designed in such a way that they can be linearly combined to *any* orientation. This interpolation property allows us to steer the final filters and to create anisotropic filter kernels. The directions $\hat{\mathbf{n}}_i$ of the oriented weight functions are for 2D filters given by:

$$\hat{\mathbf{n}}_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \hat{\mathbf{n}}_2 = \begin{pmatrix} 1/2 \\ \sqrt{3}/2 \end{pmatrix} \text{ and } \hat{\mathbf{n}}_3 = \begin{pmatrix} -1/2 \\ \sqrt{3}/2 \end{pmatrix}. \quad (10)$$

The oriented weight functions are then obtained as

$$g_i(\mathbf{z}) = \frac{4}{3}\left(1 - g_{iso}(\mathbf{z})\right)\left(\left(\frac{\mathbf{z}^T\hat{\mathbf{n}}_i}{\|\mathbf{z}\|}\right)^2 - \frac{1}{4}\right), i = 1 \dots 3. \quad (11)$$

The center isotropic weight $g_{iso}(\mathbf{z})$ is user defined. A natural choice is a Gaussian shaped kernel with a width equal to half the original $f(\mathbf{z})$ filter size. Weighting the original Gaussian low pass filter with each of these four functions yields the spatial filter basis functions,

$$\begin{aligned} f_{iso}(\mathbf{z}) &= g_{iso}(\mathbf{z})f(\mathbf{z}), \\ f_i(\mathbf{z}) &= g_i(\mathbf{z})f(\mathbf{z}),\ i = 1 \dots 3. \end{aligned} \quad (12)$$

[1] $y(t) = (w_{y_1} + w_{y_2})\, y_1(t) + \alpha\,(w_{y_1} - w_{y_2})\, y_2(t),\ w_{y_1} > 0, w_{y_2} > 0$

These functions are shown in the box in Fig. 1c. By simply adding the basis filters together, the original low pass filter is reconstructed. However, by weighting them differently other filter shapes are possible, for example anisotropic low pass filters in *any* orientation as well as isotropic filters with different sizes can be constructed. A finer partitioning of the filter space can be obtained at the expense of more filter functions. It can however be shown that three oriented filters is the minimum number that gives the steerable property. For 3D filtering the minimum number of oriented weight functions is six (Granlund and Knutsson, 1995), distributed as:

$$\begin{aligned} &\hat{\mathbf{n}}_1 = \begin{pmatrix} a \\ 0 \\ b \end{pmatrix}, \hat{\mathbf{n}}_2 = \begin{pmatrix} -a \\ 0 \\ b \end{pmatrix}, \hat{\mathbf{n}}_3 = \begin{pmatrix} b \\ a \\ 0 \end{pmatrix}, \\ &\hat{\mathbf{n}}_4 = \begin{pmatrix} b \\ -a \\ 0 \end{pmatrix}, \hat{\mathbf{n}}_5 = \begin{pmatrix} 0 \\ b \\ a \end{pmatrix}, \hat{\mathbf{n}}_6 = \begin{pmatrix} 0 \\ b \\ -a \end{pmatrix}, \\ &a = \frac{2}{\sqrt{10 + 2\sqrt{5}}}, \quad b = \frac{1 + \sqrt{5}}{\sqrt{10 + 2\sqrt{5}}}. \end{aligned} \quad (13)$$

The procedure for generating the filter basis functions is identical to the 2D case except for some constants in the weight functions:

$$g_i(\mathbf{z}) = \left(1 - g_{iso}(\mathbf{z})\right)\left(\left(\frac{\mathbf{z}^T\hat{\mathbf{n}}_i}{\|\mathbf{z}\|}\right)^2 - \frac{1}{6}\right), i = 1 \dots 6. \quad (14)$$

Finally, it is natural to allow only positive weights when combining the basis filters in Fig. 1c to obtain a pure smoothing effect. Hence, the constrained CCA fits very nicely in this context.

Putting the Pieces Together

To summarize the work so far, we have a set of temporal basis functions $\mathbf{y}(t) = [y_1(t),\, y_2(t)]^T$, a set of spatial filter basis functions and constrained CCA that takes two multidimensional data sets as input. Here these inputs will be in form of time series. Now assume that we would like to examine a certain voxel for activity. Then we first apply the filter basis functions to the voxel neighborhood and obtain a number of time series $\mathbf{x}(t) = [x_1(t), \dots x_k(t)]^T$, one for each basis filter (i.e. four in 2D and seven in 3D). A first approach is to plug the $\mathbf{x}(t)$ and $\mathbf{y}(t)$ time series directly into the constrained CCA. From the Theory section we know that the CCA finds the best linear combinations of these time series, $x(t) = \mathbf{w}_x^T\mathbf{x}(t)$ and $y(t) = \mathbf{w}_y^T\mathbf{y}(t)$, in the sense that the resulting time series $x(t)$ and $y(t)$ correlate best. Hence, the best matching hemodynamic response and the spatial filter that best extracts this response are found simultaneously. However, for better utilization of the constraints, for computational purposes and for improved detection properties, we instead choose the following strategy (see Fig. 2).

1. Perform constrained CCA with the time series obtained from the oriented basis filters and the temporal basis functions. The change-of-variables trick in Eq. 9 is used for the temporal basis functions in $\mathbf{y}(t)$.

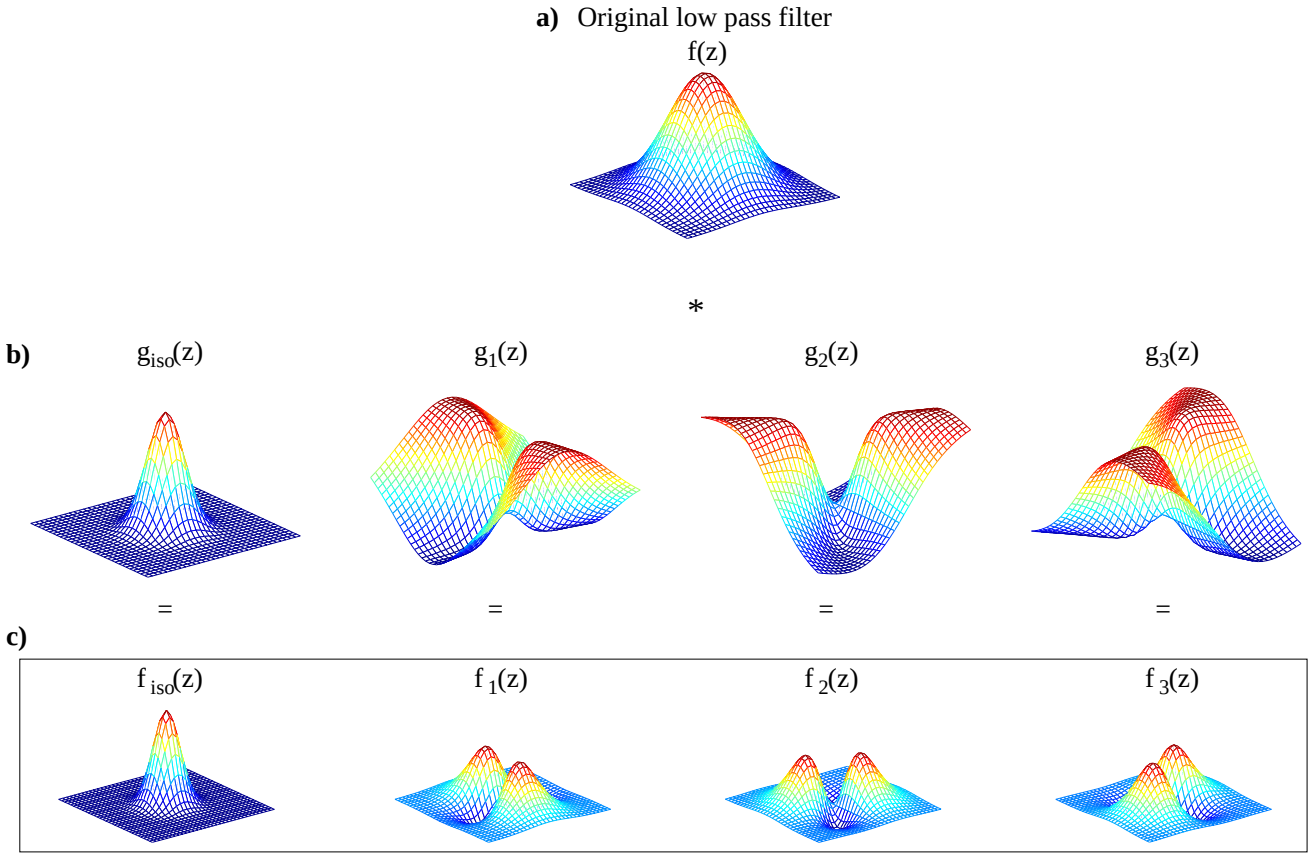

Figure 1: The construction of a set of 2D spatial basis filters. In a), a smoothing filter traditionally used for filtering the fMRI images is shown. In b), four weight functions divide the filter space into four parts. These functions sum to one in every point. In c) the final spatial basis functions are obtained by weighting the original filter with the weight functions.

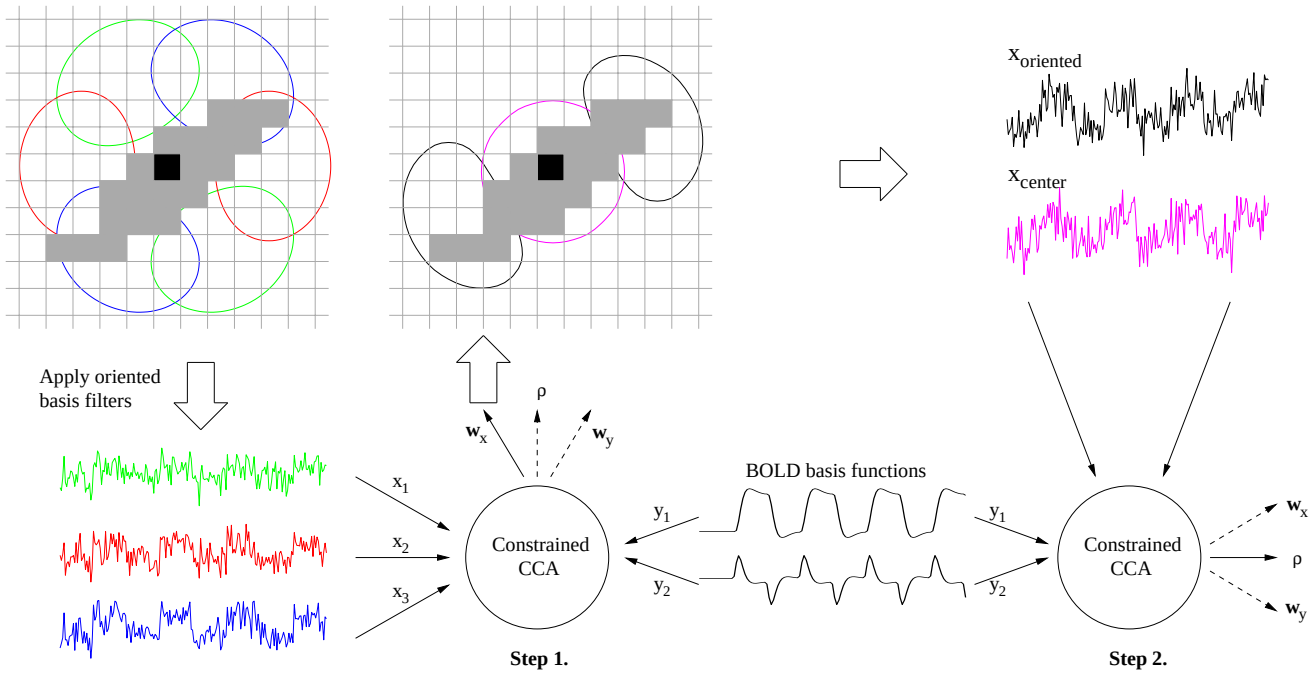

Figure 2: A schematic overview of the analysis steps. At the top left an active area in the brain is shown and the black voxel is currently under analysis. First the oriented spatial basis functions, indicated by the red, green and blue contours, are applied. The resulting time series are used as input to the constrained CCA in Step 1. The regression weight vector $\mathbf{w}_x$ can be interpreted in terms of a new filter oriented for best extraction of the hemodynamic response (black contour). The resulting optimal oriented timecourse is used together with the center filtered timecourse (purple) as input to the constrained CCA in Step 2. The resulting correlation coefficient is assigned to the center voxel.

2. Perform constrained CCA with the time series obtained from the central isotropic basis filter augmented with the "oriented" time series resulting from the above step and the temporal basis functions. The change-of-variables trick is used for *both* the $\mathbf{x}(t)$ and $\mathbf{y}(t)$ variables (see below). Assign the resulting correlation coefficient to the voxel.

In step 1, the constrained CCA finds the linear combination of the oriented basis filters that produces the time series that correlate best with the hemodynamic response model. Due to the interpolating property of the steerable filters, this can be interpreted as finding the best orientation to apply a spatial filter. Since the center basis filter $f_{iso}(\mathbf{z})$ does not provide any orientational information, it is excluded in this step to gain computational speed (remember that the computational complexity of constrained CCA grows exponentially with the number of input variables). In step 2, the center filtered time series and the resulting "oriented" time series from step 1 are combined and a final, possibly anisotropic, low pass filter is implicitly constructed. This spatial filter is adapted to extract a hemodynamic response in the neighborhood currently under analysis. To emphasize that the center neighborhood is more important than the surroundings, a change of variables is used also for the $\mathbf{x}$-variables in the second step,

$$\begin{aligned} \tilde{x}_1(t) &= x_{center}(t), \\ \tilde{x}_2(t) &= x_{center}(t) + x_{oriented}(t). \end{aligned} \tag{15}$$

When combining $\tilde{x}_1(t)$ and $\tilde{x}_2(t)$ in Eq. 15 with *positive* regression coefficients, the more important center filtered time series $x_{center}(t)$ cannot be ignored [2]. Another way of viewing this is that reasonable adaptive filter shapes is guaranteed in the same way as the change of variables guarantee that plausible hemodynamic response shapes are created of the temporal basis functions.

This two-step procedure exploits the possibilities offered by the constrained CCA in a better way compared to the more direct one-step procedure first suggested. Also, since there are fewer variables in each step, computational speed is gained.

IV. RESULTS

In this section we elucidate different aspects of the methods proposed above and provide results that demonstrate the improved detection performance. Matlab code is available on request.

Example of Temporal Basis Functions Design

A commonly used model for the hemodynamic response is a convolution between the binary reference function indicating stimulus presentation times and an impulse response consisting of a difference of two Gamma functions. In total, there are five parameters controlling the shape of this impulse response. Generally these parameters are fixed, for example to the values empirically estimated by G.H. Glover (Glover, 1999). However, many different choices of parameter values result in plausible hemodynamic response shapes, Fig. 3 shows some examples. As the parameters control the response shape in a nonlinear fashion, it is not

[2] $x(t) = (w_{x_1} + w_{x_2})\, x_{center}(t) + w_{x_2}\, x_{oriented}(t),\ w_{x_1} > 0, w_{x_2} > 0$

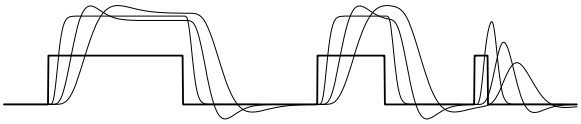

Figure 3: Possible shapes of the hemodynamic response when a difference between two Gamma functions is used as impulse response in a convolutive hemodynamic response model. The square wave function represents the putative reference function.

straightforward to relax this model and use it with freely varying parameters (Friman et al., 2002). Instead, we apply the procedure described in the Methods section and begin by generating 500 different responses with random parameter values within reasonable ranges. The average response and the first principal

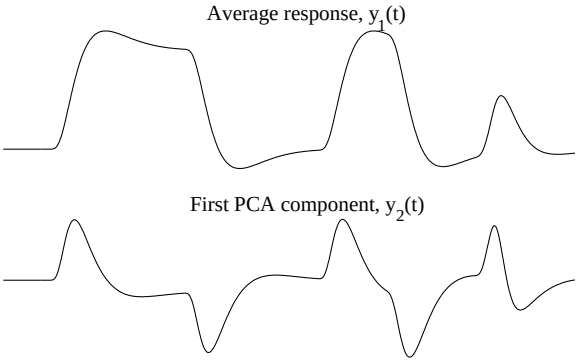

Figure 4: The resulting temporal basis functions when PCA has been applied to simulated responses similar to the examples in Fig. 3

"eigen"-response are shown in Fig. 4. The first principal component accounts for about 80% of the variation around the average response. Consequently, it is possible to reconstruct the simulated responses quite well using a linear model with $y_1(t)$ and $y_2(t)$ in Fig. 4 as basis functions. To further refine the model we make the change of variables suggested in Eq. 9, see Fig. 5. The α-

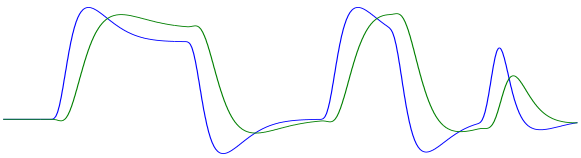

Figure 5: The final temporal basis functions obtained after the variable transformation of the timecourses in Fig. 4. Green: $\tilde{y}_1(t) = y_1(t) - 0.3y_2(t)$, blue: $\tilde{y}_2(t) = y_1(t) + 0.3y_2(t)$.

parameter was determined to 0.3 by inspecting how much of each timecourse in Fig. 4 that was needed to reconstruct the original simulated timecourses. In effect, the regression weight for the average response $y_1(t)$ is forced to be about three times larger than the weight for the first "eigen"-timecourse $y_2(t)$. The final model is obtained as $y(t) = w_{y_1}\tilde{y}_1(t) + w_{y_2}\tilde{y}_2(t)$, where w_{y_1} and w_{y_2} are non-negative and are found by the constrained CCA. The shapes in Fig. 5 then represent the extremes of the range of hemodynamic response shapes that can be produced. Hence, it is not possible to construct unrealistic response shapes, implying good specificity since fits to noise are very unlikely.

Fixed vs Steerable Spatial Filtering

Here we demonstrate the advantage of using the steerable basis filters compared to the standard fixed Gaussian smoothing. For this purpose, artificial activity was embedded in a real fMRI data set acquired when no specific task was performed. As we do not want to be dependent on uncertainties about the shape of the hemodynamic response it was simply taken to be a boxcar (the exact shape is for this comparison irrelevant as long as it is known). Different patterns of activity with different strengths were embedded in the data, see top image in Fig. 6. The constrained CCA method was then applied using two different sizes (FWHM 4 mm and 6 mm) of the original Gaussian spatial filter. The known boxcar shape was used as single temporal basis function and the $\mathbf{y}(t)$ variable was therefore one-dimensional in this particular example. For comparison, equivalent results when applying the Gaussian smoothing filters directly to the images were also calculated. In this latter approach, the CCA collapses into an ordinary correlation analysis. Thresholds were subsequently selected so that no areas outside the active voxels and their direct neighbors were declared active. The results are reported in Fig. 6. For both examined filter sizes the steerable filtering approach clearly is superior to the fixed Gaussian filtering.

The ability of the steerable filters to adapt to the local activity patterns is illustrated in Fig. 7. Note the adaptivity in both *orientation* and *size* which renders the result relatively insensitive to the choice of original filter size.

Constrained vs Unconstrained Analysis

In order to illustrate the improved detection performance obtained with non-negativity constraints imposed on the regression weights, we use data from a mental calculation fMRI experiment (B_0 1.5 T, FOV 20 cm, TE 60 ms, TR 2 s, slice thickness 6 mm, 128×128 matrix). Fixed and adaptive spatial filtering were combined with unconstrained and constrained regression weights. A Gaussian spatial filter (FWHM 8 mm) was used for the fixed filtering and to create the spatial basis functions for the adaptive filtering. The proposed PCA method was used to create the temporal basis functions. An ordinary CCA was employed for the unconstrained analyses, which combined with fixed filtering becomes equivalent to the GLM approach. The resulting correlation maps for a slice are shown in Fig. 8. The benefit of the constraints imposed on the temporal basis functions can be seen at the area indicated with the arrow in Fig. 8. This area vanishes entirely when the temporal constraints are introduced and only plausible temporal shapes are allowed. The effect of the adaptive filtering is less pronounced when applied to real data compared with the synthetic data used in the previous section. Still, a constrained analysis with adaptive filtering yields more distinct correlation maps with better contrast and less blurring of active areas. This should be attributed to the ability to create different filter sizes rather than anisotropic filters. Without constraints the adaptive filters find spurious correlations due to the large freedom in constructing suitable filters. To summarize, the effect of the non-negativity constraints is suppressed correlation levels in non-active

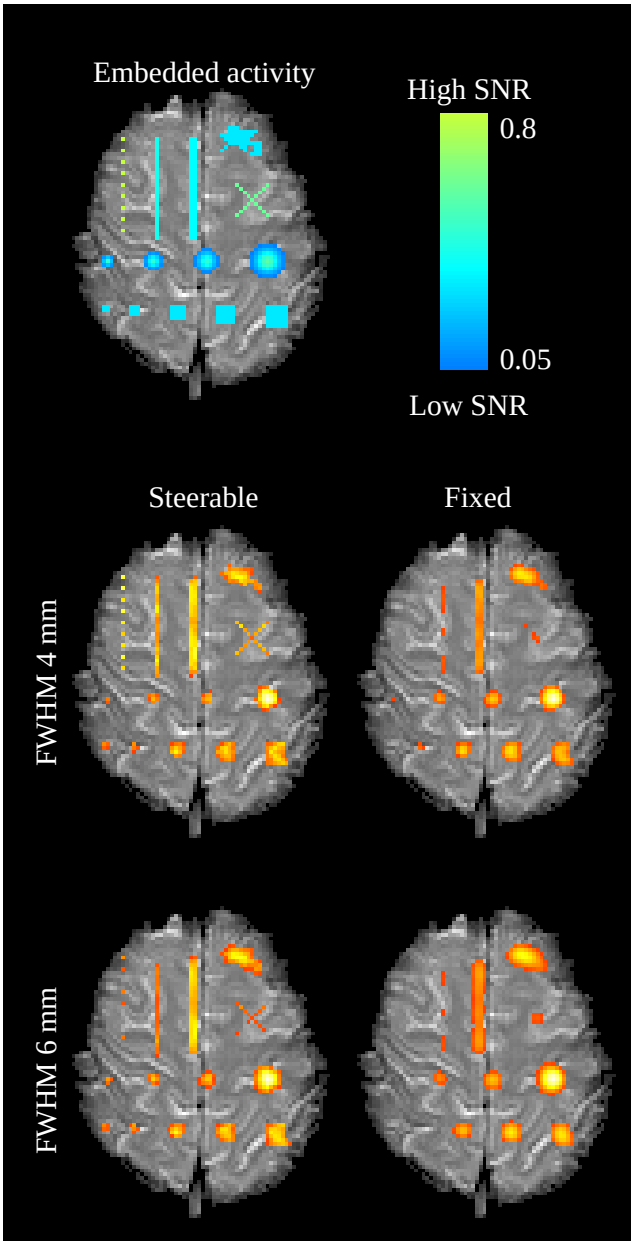

Figure 6: A comparison between spatial filtering using a fixed Gaussian kernel and the adaptive filtering obtained when using the steerable filters as spatial basis functions. The top panel illustrates the artificially embedded activity while the bottom panels show the result for two different filter sizes.

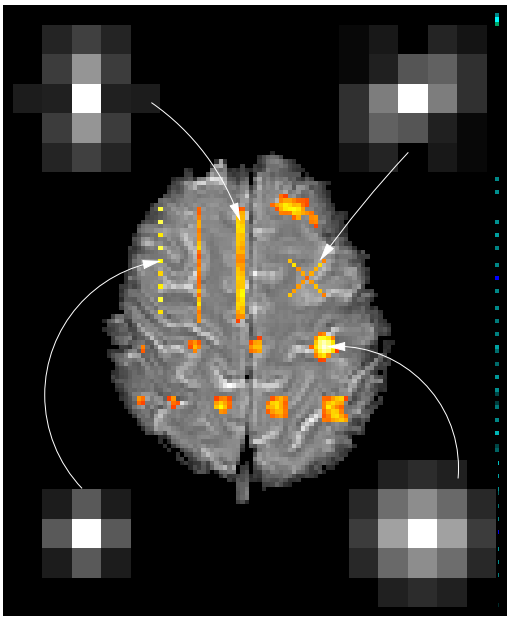

Figure 7: Adaptive filter shapes that are produced by combining the spatial basis functions at different locations in the artificial data set.

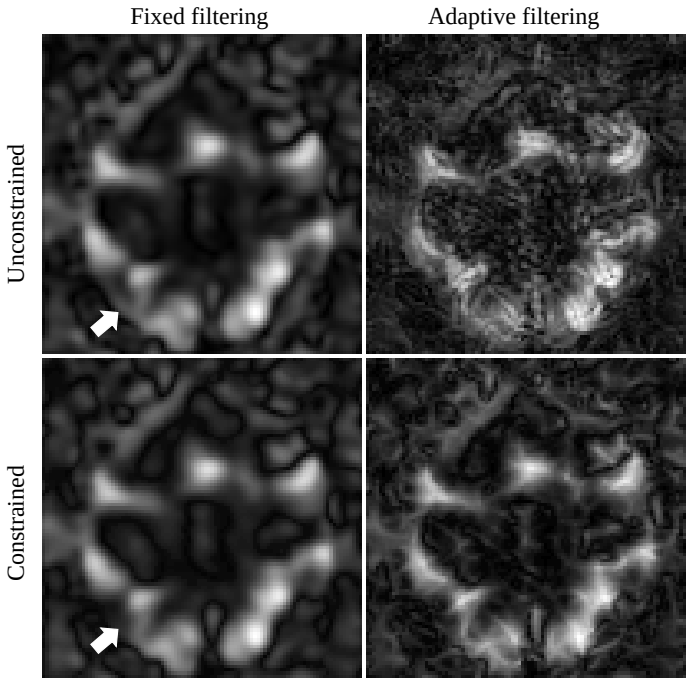

Figure 8: Correlation maps obtained with and without the non-negativity constraints imposed on the regression weights and with fixed and adaptive spatial filtering.

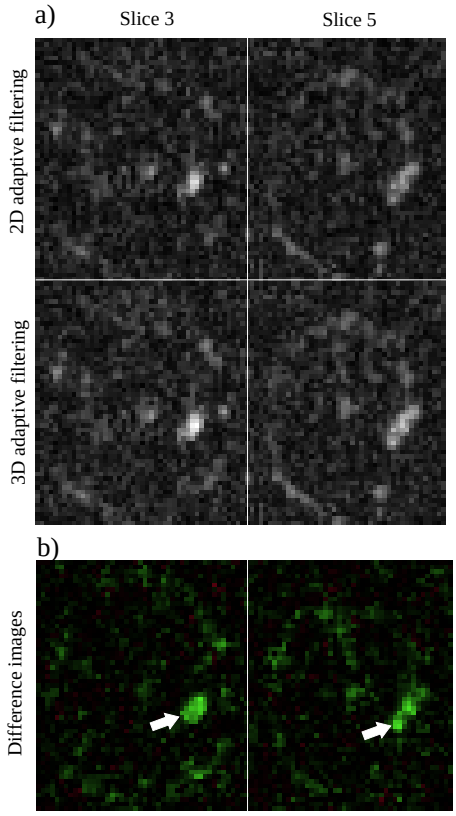

Figure 9: a) shows the correlation maps for two slices when 2D and 3D adaptive spatial filtering have been applied. In b) the difference between the correlation maps are shown. Green indicates that 3D filtering results in higher correlation than 2D filtering while red indicates the opposite.

areas while the correlation levels in active voxels are maintained high, i.e. an improved specificity.

2D vs 3D Spatial Filtering

In this subsection, 2D and 3D adaptive spatial filtering are compared. Generally, higher sensitivity is expected with a 3D adaptive filtering approach. It must however be ensured that the gain in sensitivity supersedes the loss in specificity induced by the larger number of degrees of freedom that accompanies 3D filtering. For the comparison, data from a right hand finger tapping experiment with isotropic voxel size is used (B_0 1.5 T, FOV 24 cm, TE 60 ms, TR 2 s, slice thickness 3 mm, $80\times80\times12$ matrix, $3\times3\times3$ mm voxels). The 2D and 3D spatial basis filter sets were constructed from an original Gaussian filter kernel (FWHM 5 mm) as described in the Methods section. The result for two slices are shown in Fig. 9. While the correlation maps for the 2D and 3D filtering in Fig. 9a visually appear quite similar, the difference images in Fig. 9b reveal that 3D filtering seems to improve the possibilities to detect the active areas. The better performance of the 3D filtering is obtained at the cost of extra computations. The computational time is of the order of seconds for a 2D analysis and of the order of minutes for a 3D analysis. Both alternatives are however feasible in practice.

V. DISCUSSION

In this paper a number of improvements of fMRI data analysis have been presented. First it may be appropriate to relate the proposed methods to the widely used GLM approach. The GLM is equivalent to an *unconstrained* CCA with just one spatial basis filter, usually a Gaussian smoothing filter is chosen. The design matrix used in the GLM is in this work denoted $\mathbf{y}(t)$ (drifts etc. excluded). By generalizing the GLM into CCA, the option of using spatial basis functions is enabled. This is a natural spatial counterpart of the traditional use of temporal basis functions for modelling the hemodynamic response. It is well known that an active brain area is best detected with a spatial filter adapted to the shape of that area. The filters constructed by the steerable filter set can adapt to both the orientation and size of the active areas. The advantages obtained with the adaptive filtering were evident when applied to the synthetic activation shapes. In real data where active areas are less extremely shaped the ability to create different filter sizes seems more important than the ability to create anisotropic filters. A more compact set of spatial basis functions which steers only in scale may therefore provide even better delineation of active brain regions.

Another improvement is the constrained analysis approach. Constraining the regression weights can be seen as a way of introducing prior information into the analysis. We know beforehand what a reasonable hemodynamic response looks like and that the spatial filter should have a smoothing effect. The constraints prevent fits to noise and the obvious result is an improvement in detection performance. The price is a slight increase in computation time.

Furthermore, a general approach for generating compact subspace models of the hemodynamic response has been introduced. The shape of the hemodynamic response can today be considered well known and using an excess of temporal basis functions is more likely to deteriorate than improve the analysis result. For this reason we use only two temporal basis functions but it is possible to generalize to more basis functions, see the Appendix. Recently, a similar approach for finding suitable temporal basis functions was presented by Hossein-Zadeh and Ardekani (Hossein-Zadeh and Ardekani, 2002). The minor difference is that Hossein-Zadeh et al. make a PCA directly on the impulse response in a convolutive model. The method proposed here is slightly more general and is applicable to more sophisticated hemodynamic response models, for example Buxton's Balloon model (Buxton et al., 1998).

Finally, an issue not addressed in this article is how a correlation threshold for the correlation maps is found. The null distribution for the constrained canonical correlation coefficient is unknown and given the autocorrelated noise structure in fMRI data, an analytical solution to this problem is not likely possible. Thus, one has to resort to non-parametric statistical methods for finding an appropriate threshold. Such procedures have been presented elsewhere (Das and Sen, 1996; Locascio et al., 1997; Bullmore et al., 2001; Nichols and Holmes, 2001). In this context it should also be stressed that pre-processing in form of temporal filtering, in order to create an uniform temporal autocorrelation structure across voxels, can be applied independently of the analysis improvements suggested in this work.

VI. CONCLUSIONS

An analysis method based on constrained Canonical Correlation Analysis has been presented. The approach is a natural extension of the General Linear Model analysis technique and it therefore possesses nice properties such as intuitiveness and transparency. The concept of spatial filter basis functions has also been introduced, as well as a method for constructing suitable temporal basis functions. The Results section show how these different and novel parts contribute to significant improvements of the detection performance, with a computation time well within limits for practical use.

ACKNOWLEDGEMENTS

The financial support from the Swedish Research Council and Linköping University Hospital funds is gratefully acknowledged.

APPENDIX

We have used two basis functions to model the hemodynamic response because it strikes a good balance between sensitivity, specificity and computational complexity. It is possible to include for example also the second principal component (in addition to the mean response and first principal component) in the temporal model. The following change of variables needs to be done when a constrained analysis method is used,

$$\begin{aligned}
\tilde{y}_1(t) &= y_1(t) + \alpha y_2(t) + \beta y_3(t), \\
\tilde{y}_2(t) &= y_1(t) + \alpha y_2(t) - \beta y_3(t), \\
\tilde{y}_3(t) &= y_1(t) - \alpha y_2(t) + \beta y_3(t), \\
(\tilde{y}_4(t) &= y_1(t) - \alpha y_2(t) - \beta y_3(t)),
\end{aligned}$$

where $0 < \beta \leq \alpha < 1$. As before, α and β indicate the strength of the two first principal components in relation two the mean response $y_1(t)$. Only the first three variables $\tilde{y}_1(t)$, $\tilde{y}_2(t)$ and $\tilde{y}_3(t)$ should be used. The $\tilde{y}_4(t)$ equation can be written as a linear combination of the first three equations and it is included just to illustrate how the change of variables trick generalizes to higher dimensions. Following this pattern it is straightforward to add more basis functions.

References

Aguirre, G., Zarahn, E., and D'Esposito, M. (1998). The variability of human, BOLD hemodynamic responses. *NeuroImage*, 8(4):360–369.

Anderson, T. W. (1984). *An Introduction to Multivariate Statistical Analysis*. John Wiley & Sons, second edition.

Bullmore, E., Long, C., Suckling, J., Fadili, J., Calvert, G., Zelaya, F., Carpenter, T., and Brammer, M. (2001). Colored noise and computational inference in neurophysiological (fMRI) time series analysis: resampling methods in time and wavelet domains,. *Human Brain Mapping*, 12(2):61–78.

Buxton, R., Wong, E., and Frank, L. (1998). Dynamics of blood flow and oxygenation changes during brain activation: the Balloon model. *Magnetic Resonance in Medicine*, 39(6):855–864.

Das, S. and Sen, P. (1994). Restricted canonical correlations. *Linear Algebra and its Applications*, 210:29–47.

Das, S. and Sen, P. (1996). Asymptotic distribution of restricted canonical correlations and relevant resampling methods. *Journal of Multivariate Analysis*, 56(1):1–19.

Freeman, W. T. and Adelson, E. H. (1991). The design and use of steerable filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(9):891–906.

Friman, O., Borga, M., Lundberg, P., and Knutsson, H. (2002). Detecting neural activity in fMRI using maximum correlation modeling. *NeuroImage*, 15(2):386–395.

Friston, K., Mechelli, A., Turner, R., and Price, C. (2000). Nonlinear responses in fMRI: The Balloon model, Volterra kernels, and other hemodynamics. *NeuroImage*, 12(4):466–477.

Glover, G. (1999). Deconvolution of impulse response in event-related BOLD fMRI. *NeuroImage*, 9(4):416–429.

Granlund, G. H. and Knutsson, H. (1995). *Signal Processing for Computer Vision*. Kluwer Academic Publishers. ISBN 0-7923-9530-1.

Hossein-Zadeh, G. and Ardekani, B. (2002). A signal subspace for modeling the hemodynamic response function in fMRI. In *Proc. Intl. Soc. Mag. Reson. Med. 10 (ISMRM'02)*.

Hotelling, H. (1936). Relations between two sets of variates. *Biometrika*, 28:321–377.

Knutsson, H., Wilson, R., and Granlund, G. H. (1983). Anisotropic non-stationary image estimation and its applications — Part I: Restoration of noisy images. *IEEE Transactions on Communications*, 31(3):388–397. Report LiTH–ISY–I–0462, Linköping University, Sweden, 1981.

Locascio, J., Jennings, P., Moore, C., and Corkin, S. (1997). Time series analysis in the time domain and resampling methods for studies of functional magnetic resonance brain imaging. *Human Brain Mapping*, 5(3):168–193.

Nichols, T. and Holmes, A. (2001). Nonparametric permutation tests for functional neuroimaging: A primer with examples. *Human Brain Mapping*, 15(1):1–25.