A NEW FRAMEWORK FOR FMRI DATA ANALYSIS: MODELING, IMAGE RESTORATION, AND ACTIVATION DETECTION

Jianing Wei and Ilya Pollak

School of Electrical and Computer Engineering
Purdue University
West Lafayette, IN 47907

ABSTRACT

We propose a new model for event-related functional magnetic resonance imaging (fMRI), and develop a new set of tools for activation detection. A novel feature of our framework is the explicit modeling of the spatial correlation introduced by the scanner. We propose simple, efficient algorithms to estimate model parameters. We develop an activation detection algorithm which consists of two parts: image restoration and least-squares estimation of the parameters of the hemodynamic response function. During the image restoration stage, a total-variation-based approach is employed to restore each data slice, for each time index. The amplitude of the least-squares fit of the hemodynamic response function is then thresholded to yield an estimate of the activation map. We illustrate the promise of our method through several experiments with synthetic data as well as one example with real data.

Index Terms— Magnetic resonance imaging, parameter estimation, detection, image restoration, modeling

1. INTRODUCTION

Functional magnetic resonance imaging (fMRI), a noninvasive technique to study functional brain anatomy, has been an active area of research. During an fMRI experiment, a stimulus is presented to a subject, and a sequence of several volumetric images of the brain is taken. These images measure the concentration of oxygen in the blood as a function of the location in the brain. Since blood oxygenation levels are related to neuronal activity [11], accurate detection of the regions of increased activity from fMRI data would allow establishing which areas of the brain are responsible for processing the stimulus. However, since the fMRI measurement process results in severe noise and other degradations, accurate detection of activation regions is a challenging inverse problem. A paradigm that has been frequently employed to address this problem in recent years is to propose a forward model and develop a statistical estimation method to invert it [2, 4, 5, 6, 9, 10, 13].

In this paper, we focus on event-related fMRI which measures the response to a single short stimulus or a sequence of such stimuli. We develop a new activation detection algorithm. The input to our algorithm is fMRI data, and the output for each pixel is the decision which says whether or not the pixel is activated. Our algorithm has two main novel features:

- Our forward model explicitly incorporates spatial blurring introduced by the scanning process. We model this blurring as a convolution with a spatially isotropic Gaussian kernel.
- We use an image restoration algorithm based on constrained minimization of the total variation (TV) [7]. TV-based image restoration techniques [14] have been shown to reduce noise while preserving image edges [15]. These methods have been demonstrated to be robust to significant amounts of noise and blurring, and to be especially well suited to the situations where the image to be restored is piecewise constant. These properties make such techniques ideally matched to our problem of producing a binary activation map based on very noisy, blurred data.

We use the gamma-variate model [13] for the hemodynamic response function (HRF) [8]. HRF models a pixel's response when stimulated by an ideal impulse. To capture the temporal correlation of the data, we adopt the autoregressive noise model of [12]. We develop computationally efficient procedures to estimate the parameters of this model. We model the spatial correlation of fMRI data as the effect of a Gaussian blur. We develop a method to estimate the width of the Gaussian kernel. Our activation detection algorithm then consists of two main steps. First, we perform TV-based image restoration frame by frame. Then we compute the least-squares fit of the HRF and obtain the activation map by thresholding the amplitude parameters of the estimated HRF. Our proposed approach is tested on simulated data and compared with two other approaches. The performance is assessed using Receiver Operating Characteristic (ROC) curves. We also illustrate our method on real data.

2. MODEL FORMULATION

Real fMRI data is 4D: It has three spatial dimensions and one temporal dimension. However, for notational convenience as well as for the ease and speed of our implementations and experiments, we assume 3D data $y_{i,j,t}$ where i and j index two spatial dimensions, and t is a discrete time parameter.

We assume that, if stimulated with an ideal impulse, the pixel (i,j) produces the following response for time $t \geq 0$:

$$h_{i,j,t} = \begin{cases} x_{i,j} \left(\dfrac{t-\delta_{i,j}}{\tau_{i,j}} \right)^2 e^{-\frac{t-\delta_{i,j}}{\tau_{i,j}}} & t \geq \delta_{i,j}, \\ 0 & t < \delta_{i,j}, \end{cases} \tag{1}$$

where the parameters $x_{i,j}$, $\delta_{i,j}$, and $\tau_{i,j}$ have the following interpretations:

- $\delta_{i,j}$ is the delay between the presentation of the stimulus and the onset of the response;
- $2\tau_{i,j}$ is the time from the onset of the response to its peak;
- $4x_{i,j}/e^2$ is the peak amplitude of the response.

This work was supported in part by a National Science Foundation grant IIS-0329156.

This impulse response is also commonly referred to as the hemodynamic response function (HRF). The form of HRF used in Eq. (1) is taken from [3, 13]. Note that, if pixel (i, j) is not activated, then $x_{i,j} = 0$. Our approach to detect activation is therefore to estimate the amplitude parameter $x_{i,j}$ for each pixel's HRF, and threshold it. Activation is detected at pixel (i, j) if $x_{i,j}$ is above a threshold.

Each pixel's hemodynamic response is subject to physiological noise $\omega_{i,j,t}$, i.e., noise arising due to the subject's heartbeat, breathing, etc. We model this noise as an additive AR(1) process, following [12]. In other words, the physiological noise is modeled as the output of the following linear system driven by white noise $\epsilon_{i,j,t}$:

$$\omega_{i,j,t} = \rho\omega_{i,j,t-1} + \epsilon_{i,j,t}. \tag{2}$$

We assume that the white noise $\epsilon_{i,j,t}$ is zero-mean. The parameters of this model are therefore ρ (which we assume to be a real number between 0 and 1) and the standard deviation σ_ϵ of the noise $\epsilon_{i,j,t}$. We assume each of these parameters to be the same for all pixels and all times.

In addition, it is common to assume the presence of an additive baseline m, which is the signal level when no stimulus is presented [8]. However, we assume that an accurate estimate of the baseline may be obtained by averaging the response values measured when no stimulus is applied. This estimate can then be subtracted from the data. Therefore, in the remainder of the paper we assume zero baseline.

The sum $h_{i,j,t} + \omega_{i,j,t}$ of the the HRF and the physiological noise is subject to two further types of degradation introduced by the scanner: spatial blurring and noise. We model the former as the convolution with a spatially isotropic Gaussian kernel $g_{i,j}$ with zero mean and standard deviation σ. We model scanner noise as additive white noise $\eta_{i,j,t}$ with zero mean and standard deviation σ_η which is assumed to be a constant both spatially and temporally. Putting all the pieces of our model together results in the following:

$$y_{i,j,t} = \sum_k \sum_l g_{k,l}(h_{i-k,j-l,t} + \omega_{i-k,j-l,t}) + \eta_{i,j,t}. \tag{3}$$

3. ACTIVATION DETECTION

Our activation detection algorithm first estimates the HRF $h_{i,j,t}$ from the data $y_{i,j,t}$, and then estimates the amplitude parameters $x_{i,j}$ of the estimated HRF. The activation map is obtained through thresholding $x_{i,j}$. The remainder of this section describes the estimation algorithm. This algorithm relies on the estimates of the following model parameters: the standard deviation σ of the Gaussian blurring kernel, the physiological noise parameters ρ and σ_ϵ, and the scanner noise parameter σ_η. The procedures for estimating these parameters are developed in the next section.

The first part of our activation detection algorithm is accomplished through a total-variation (TV) based restoration algorithm [14] which is applied frame by frame—in other words, the data $\{y_{i,j,t}\}_{i,j}$ is restored for each fixed value of t. We therefore drop the t index for notational convenience. We denote the combined contribution of the physiological and scanner noise to pixel (i, j) in Eq. (3) by $v_{i,j}$:

$$v_{i,j} \equiv \sum_k \sum_l g_{k,l}\omega_{i-k,j-l} + \eta_{i,j}.$$

With this notation, our model of Eq. (3) can be rewritten as follows:

$$y_{i,j} = \sum_k \sum_l g_{k,l}h_{i-k,j-l} + v_{i,j}. \tag{4}$$

We formulate the following constrained total variation minimization problem, where we assume that our images are $N \times N$:

$$\begin{aligned} \min_{\{h_{i,j}\}} \quad & \sum_i \sum_j \sqrt{(h_{i+1,j} - h_{i,j})^2 + (h_{i,j+1} - h_{i,j})^2} \\ \text{subject to} \quad & \text{Eq. (4)} \\ \text{and} \quad & \frac{1}{N^2}\sum_i \sum_j v_{i,j}^2 \le \sigma_v^2. \end{aligned} \tag{5}$$

We follow [7] and use second-order cone programming (SOCP) to solve this optimization problem. Our SOCP is almost identical to that in [7], with the exception of the additional blurring term in the constraint of Eq. (4). This SOCP is solved using interior point method through MOSEK software [1]. The constraint parameter σ_v^2 is an estimate of the total noise variance, which can be obtained as follows:

$$\sigma_v^2 = \sum_l \sum_k g_{k,l}^2 \frac{\sigma_\epsilon^2}{1-\rho^2} + \sigma_\eta^2. \tag{6}$$

The next section shows how to estimate all the parameters used in the right-hand side of this equation.

The result of performing the above optimization for every time t, is a set of estimates of HRF, at every pixel location, and for every time t. We denote these estimates by $\hat{h}_{i,j,t}$. The second stage of our activation detection algorithm uses $\hat{h}_{i,j,t}$ in conjunction with the HRF model of Eq. (1) to produce least-squares estimates of the parameters $x_{i,j}$, $\tau_{i,j}$, and $\delta_{i,j}$. We then classify each pixel whose parameter $x_{i,j}$ is above a threshold as activated, and each pixel whose parameter $x_{i,j}$ is below the same threshold as non-activated.

4. PARAMETER ESTIMATION

We estimate the parameters from the fMRI data corresponding to the initial time interval of the experiment, when the subject is not presented with any stimulus. We let T be the duration of this interval. This results in the following simplified version of Eq. (3) for $t = 1, \ldots, T$:

$$y_{i,j,t} = \sum_k \sum_l g_{k,l}\omega_{i-k,j-l,t} + \eta_{i,j,t}. \tag{7}$$

We use these observations to estimate the model parameters: the standard deviation σ of the Gaussian blurring kernel, the physiological noise parameters ρ and σ_ϵ, and the scanner noise parameter σ_η. We rasterize the data and the three noise processes, for each time t, to define four vectors $\boldsymbol{y}_t$, $\boldsymbol{\omega}_t$, $\boldsymbol{\epsilon}_t$, and $\boldsymbol{\eta}_t$. For example, $\boldsymbol{y}_t = (y_{1,1,t} \ \ldots y_{1,N,t} \ y_{2,1,t} \ \ldots y_{2,N,t} \ \ldots y_{N,N,t})'$ for $t = 1, \ldots, T$, where the prime denotes the transpose of a vector. We let A be the Gaussian filtering matrix—i.e., the matrix representing the convolution with the spatial Gaussian blurring kernel. Using this notation, Eqs. (2) and (7) can be rewritten, respectively, as follows:

$$\begin{aligned} \boldsymbol{\omega}_t &= \rho\boldsymbol{\omega}_{t-1} + \boldsymbol{\epsilon}_t \\ \boldsymbol{y}_t &= A\boldsymbol{\omega}_t + \boldsymbol{\eta}_t \end{aligned}$$

Our algorithm for estimating the noise parameters ρ, σ_ϵ and σ_η is based on their relationships with the autocorrelation function of $\boldsymbol{y}_t$ which can be easily demonstrated:

$$\mathrm{E}[\boldsymbol{y}_t^T\boldsymbol{y}_t] = \frac{\sigma_\epsilon^2}{1-\rho^2}\mathrm{trace}\{AA'\} + N^2\sigma_\eta^2 \tag{8}$$

$$\mathrm{E}[\boldsymbol{y}_t^T\boldsymbol{y}_{t-1}] = \frac{\rho}{1-\rho^2}\sigma_\epsilon^2\mathrm{trace}\{AA'\} \tag{9}$$

$$\mathrm{E}[\boldsymbol{y}_t^T\boldsymbol{y}_{t-2}] = \frac{\rho^2}{1-\rho^2}\sigma_\epsilon^2\mathrm{trace}\{AA'\} \tag{10}$$

We denote $r \triangleq \mathrm{E}[\boldsymbol{y}_t^T \boldsymbol{y}_t]$, $r_1 \triangleq \mathrm{E}[\boldsymbol{y}_t^T \boldsymbol{y}_{t-1}]$, $r_2 \triangleq \mathrm{E}[\boldsymbol{y}_t^T \boldsymbol{y}_{t-2}]$. We use the following standard estimates of the autocorrelations r, r_1, and r_2:

$$\hat{r} = \frac{1}{T}\sum_{t=1}^{T}\sum_{i=1}^{N}\sum_{j=1}^{N} y_{i,j,t}^2 \quad (11)$$

$$\hat{r}_1 = \frac{1}{T-1}\sum_{t=2}^{T}\sum_{i=1}^{N}\sum_{j=1}^{N} y_{i,j,t}y_{i,j,t-1} \quad (12)$$

$$\hat{r}_2 = \frac{1}{T-2}\sum_{t=3}^{T}\sum_{i=1}^{N}\sum_{j=1}^{N} y_{i,j,t}y_{i,j,t-2} \quad (13)$$

If we let the i-th row of A be a_i, then $\mathrm{trace}\{AA'\} = N^2\|a_i\|_2^2$. Note that $\|a_i\|_2^2$ is the same for all rows and is only a function of σ. We denote this function by $\alpha(\sigma)$. Substituting the estimates of Eqs. (11-13) into Eqs. (8-10), replacing $\mathrm{trace}\{AA'\}$ with $N^2\alpha(\sigma)$, and solving for the noise parameters, we get the following estimates of the noise parameters:

$$\hat{\rho} = \hat{r}_2/\hat{r}_1 \quad (14)$$

$$\hat{\sigma}_\epsilon = \sqrt{\frac{\hat{r}_1(1-\hat{\rho}^2)}{\hat{\rho}\, N^2\alpha(\sigma)}} \quad (15)$$

$$\hat{\sigma}_\eta = \sqrt{\frac{\hat{r} - \hat{r}_1/\hat{\rho}}{N^2}} \quad (16)$$

We conclude this section by showing how to estimate the parameter σ which is needed both in order to compute an estimate of σ_ϵ via Eq. (15), and in order to estimate the blurring kernel in the constraint of Eq. (4) and the constraint parameter σ_v^2 of Eq. (6). The estimation of σ is based on its relationship with the covariance between $y_{i,j,t}$ and $y_{i,j+1,t}$. Since we are observing non-activated pixels, the means of $y_{i,j,t}$ and $y_{i,j+1,t}$ are 0. Thus the covariance between $y_{i,j,t}$ and $y_{i,j+1,t}$ is equal to $\mathrm{E}[y_{i,j,t}y_{i,j+1,t}]$. In addition, since the observations are time-stationary, $\mathrm{E}[y_{i,j,t}y_{i,j+1,t}]$ is independent of time t. We therefore drop the index t for notational convenience. It can be easily shown that:

$$\mathrm{E}[y_{i,j}y_{i,j+1}] = \frac{\sigma_\epsilon^2\langle a_i, a_{i+1}\rangle}{1-\rho^2}. \quad (17)$$

Note that $\langle a_i, a_{i+1}\rangle$ is a function of σ but not a function of i. We denote this function by $\beta(\sigma)$. We use the following standard estimate of $\mathrm{E}[y_{i,j}y_{i,j+1}]$:

$$\mathrm{E}[y_{i,j}y_{i,j+1}] = \frac{1}{N^2T}\sum_t\sum_i\sum_j y_{i,j,t}y_{i,j+1,t}.$$

Substituting the estimates of $\mathrm{E}[y_{i,j}y_{i,j+1}]$, σ_ϵ, and ρ into Eq. (17), we have the following result:

$$\frac{1}{N^2T}\sum_t\sum_i\sum_j y_{i,j,t}y_{i,j+1,t} = \frac{\hat{r}_1\beta(\sigma)}{\hat{\rho}N^2\alpha(\sigma)}. \quad (18)$$

By solving Eq. (18) for σ, we obtain the estimate of the smoothness of the image, i.e., the width of the Gaussian kernel.

5. EXPERIMENTAL RESULTS

We conduct an experiment to test the performance of the proposed activation detection approach on synthetic data and compare it with two other approaches.

Our synthetic dataset is $64 \times 64 \times 15$, representing 64×64 2D images collected at each of the 15 time points. The dataset is simulated using our model of Eqs. (1,2,3). The correct parameters of the HRF (see Eq. (1)) are $\tau_{i,j} = 2$, $\delta_{i,j} = 1.5$, and $x_{i,j} = 0.03$ for activated pixels which occupy a 4×4 rectangular region of the 64×64 image. Note that our algorithm does not know these values, and does not even know the fact that they are independent of i and j.

We define the PSNR as $20\log_{10}(p/\sigma_v)$ where $p = 0.0162$ is the peak amplitude of the HRF for every activated pixel, and σ_v is the standard deviation of the noise from Eq. (6). To conduct experiments with various levels of PSNR, we fix the ratio $\sigma_\epsilon^2/\sigma_\eta^2 = 7/16$ and increase both σ_ϵ and σ_η proportionally to decrease the PSNR. Specifically, we use four PSNR values: 20.5, 14.5, 11.8, 8.3. These values are unknown to our algorithm. The two remaining noise parameter values, also unknown to our algorithm, are $\sigma = 0.569$ and $\rho = 0.75$ in all experiments. For each of the four distinct parameter settings, we perform a Monte-Carlo run consisting of 100 simulations, and generate an empirical receiver operating characteristic (ROC). We also generate ROCs for two other activation detection approaches for comparison.

To demonstrate the added value of the restoration stage, we compare our algorithm to its simplified version where the restoration stage is omitted. In other words, the benchmark algorithm performs the least-squares fit of HRF based on raw data.

We also compare our method to the paired t-test method which has been used in literature as a benchmark [13]. This method assumes that that all activated pixels have the same HRF, and has the knowledge of the correct parameters of this HRF, including the amplitude. This correct HRF is used as the reference in the pixel-by-pixel paired t-test. Note that the knowledge of the correct HRF parameters is critical, and is available to this method but not to ours. The objective of this exercise is to compare our method with a much simpler method which, however, has much more information available to it.

The correct detection rate is plotted against the false alarm rate in Fig. 1 for four different PSNR values for our proposed approach and two other methods described above. Error bars indicate $\pm$ twice the standard deviation of the correct detection rate in 100 Monte-Carlo simulations. Note that our method is statistically indistinguishable from the t-test which uses the knowledge of the correct parameters of the HRF. In addition, our full method clearly outperforms its no-restoration version. For example, for PSNR = 11.8 and false alarm rate under 0.03, our full method achieves more than twice the correct detection rate of the method without the restoration stage. At the same level of PSNR, the full version of the algorithm needs 0.08 false alarm rate to achieve perfect correct detection rate, whereas the no-restoration version achieves perfect detection only at about 0.24 false alarms. Also note that for PSNR = 20.5, the ROC curve for our full method goes through the point with zero false alarms and 100% correct detections (see Fig. 1(a)). Finally, note that our full method degrades gracefully, even for quite low levels of PSNR.

In Fig. 2 we show an example of applying our algorithm to real data. In this experiment, the subject is presented with a visual checkerboard stimulus to the right hemifield. From the figure, we can see that the region of the detected activation is in the corresponding primary visual cortex.

6. CONCLUSION

We have proposed a new forward model for event-related fMRI which explicitly incorporates the effect of blurring introduced by the scanner. We have developed a new activation detection algorithm which

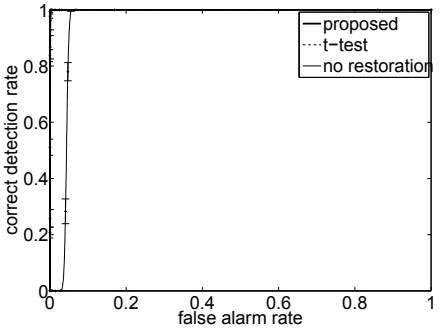

(a) ROC curves when PSNR is 20.5.

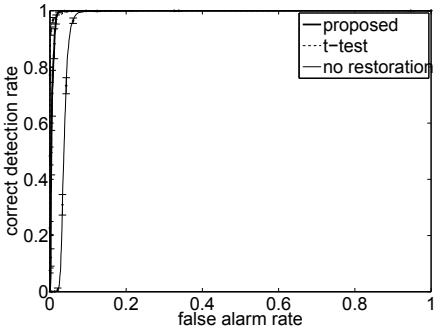

(b) ROC curves when PSNR is 14.5.

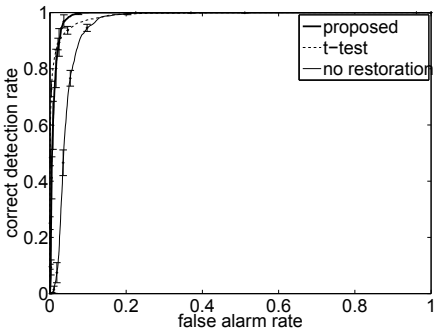

(c) ROC curves when PSNR is 11.8.

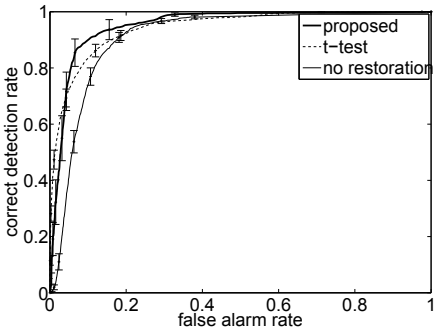

(d) ROC curves when PSNR is 8.3.

Fig. 1. Thick solid lines denote ROC curves for our proposed method. Dashed lines denote ROC curves for the t-test method. Thin solid lines denote ROC curves for the method of fitting without restoration. Error bars on each line indicate $\pm$ twice the standard deviation of the correct detection rates for the corresponding method.

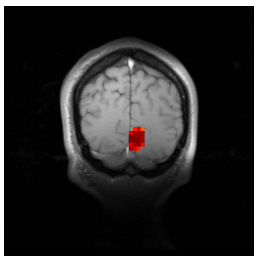

Fig. 2. Activation map of our proposed detection algorithm overlayed on top of anatomical image for a visual checkerboard stimulus presented to the right hemifield.

uses a total variation minimization algorithm to restore image data. We have also proposed novel ways of estimating model parameters. Our experimental examples showed that our method outperforms two benchmark methods on synthetic data. It also performs well in an example with real data. This illustrates the promise of our model and algorithm.

7. REFERENCES

[1] E.D. Andersen, C. Roos, and T. Terlaky. On implementing a primal-dual interior-point method for conic quadratic optimization. *Math. Program.*, 95(2):249–277, 2003.

[2] M.A. Burock and A.M. Dale. Estimation and detection of event-related fMRI signals with temporally correlated noise: a statistically efficient and unbiased approach. *Human Brain Mapping*, 11:249–260, 2000.

[3] A.M. Dale and R.L. Buckner. Selective averaging of rapidly presented individual trials using fMRI. *Human Brain Mapping*, 5:329–340, 1997.

[4] X. Descombes, F. Kruggel, and D.Y.V. Cramon. Spatial-temporal fMRI analysis using Markov Random fields. *IEEE Transactions on Medical Imaging*, 17(6):1028–1039, 1998.

[5] K.J. Friston, A.P. Holmes, K.J. Worsley, J.P. Poline, C.D. Frith, and R.S.J. Frackowiak. Statistical parametric maps in functional imaging: a general linear approach. *Human Brain Mapping*, 2:189–210, 1995.

[6] K.J. Friston, P. Fletcher, O. Josephs, A. Holmes, M.D. Rugg, and R. Turner. Event-related fMRI: characterizing differential responses. *NeuroImage*, 7:30–40, 1998.

[7] D. Goldfarb and W. Yin. Second-order cone programming methods for total variation-based image restoration. *SIAM J. Sci. Comput.*, 27(2):622–645, 2005.

[8] P. Jezzard, P.M. Matthews, and S.M. Smith. Functional MRI: an introduction to methods, New York, Oxford University Press, 2001.

[9] O. Josephs, R. Turner, and K.J. Friston. Event-related fMRI. *Human Brain Mapping*, 5:243–248, 1997.

[10] J. Kim, J.W. Fisher III, A. Tsai, C. Wible, A. Willsky, and W.M. Wells III. Incorporating spatial priors into an information theoretic approach for fMRI data analysis. *MICCAI*, 62–71, 2000.

[11] S. Ogawa, T.M. Lee, A.R. Kay, and D.W. Tank. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proc. Natl. Acad. Sci. USA*, 87:9868–9872, 1990.

[12] P.L. Purdon, V. Solo, R.M. Weisskoff, and E.N. Brown. Locally regularized spatiotemporal modeling and model comparison for function MRI. *NeuroImage*, 14:912–923, 2001.

[13] A.A. Rao and T.M. Talavage. Clustering of fMRI data for activation detection using HDR models. *Proceedings of the 26th Annual International Conference of the IEEE EMBS*, 1876–1879, 2004.

[14] L.I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.

[15] D. Strong and T. Chan. Edge-preserving and scale-dependent properties of total variation regularization. *Inverse Problems*, 19:S165–S187, 2003.