Visualizing Brain Activity from fMRI Data

Firdaus Janoos, Okan Irfanoglu, Raghu Machiraju, and Istvan Akos Morocz

Abstract— Classically, visualization the time-varying data acquired during fMRI experiments is performed using static activation maps obtained by testing voxels for the presence of significant activity, using parametric or non-parametric methods. The models used in these analysis methods have a number of tunable settings, which profoundly impact the detection of active brain areas. Also, studying the temporal dependencies and cascading effects of brain activation from these static maps is challenging. In this paper, we propose a methodology to visualize brain function in a time-varying fashion. This will allow neurologists to not only validate the static activation maps, but also to develop a better understanding of temporal characteristics of the activation patterns. The method is based on extracting the underlying BOLD signal from the data, de-noising and de-trending using a wavelet representation, time-windowed hierarchical clustering of voxels in the wavelet domain, followed by a spatio-temporal visualization of brain activity. We also demonstrate the utility of the method by presenting two use-cases on a real data-set.

Index Terms—Spatio-temporal visualization, fMRI, BOLD

1 Introduction

Functional magnetic resonance imaging (fMRI) is a very powerful tool for investigating brain activity to identify functionally specialized responses and characterize its functional anatomy [20]. Functional neuro-imaging scanners acquire spatio-temporal scalar data, with a time-series associated at every voxel. These time-series data are seldom viewed directly, but instead it is subject to elaborate processing [5, 14]. The results of the analysis are static activation maps and independent clusters, although there are beginning efforts today to characterize the connectivity among brain regions. But just like in the case of electrophysiological brain studies using EEG or MEG, where temporal resolution is in the millisecond range (in stark contrast to fMRI), there is growing evidence in the metabolic imaging community that analyzing a given neural activation pattern in the temporal domain results in a much more precise and specific mapping of the metabolic signature to a given neural stimulus. As imaging pulse sequences used for fMRI become more efficient in their temporal and spatial resolution, tools that can efficiently depict cascadic and serial brain recruitment in task performance, become ever more important. Such activity road-maps shown in cine mode, as used for years in EEG and MEG, but now enhanced by the tomographic fidelity of MRI, will crucially enhance our understanding of physiological and pathological conditions of brain responses which are not apparent from static maps.

Currently, there are two salient approaches to the generation of these static activation maps:

a Confirmatory methods which parametrically fit a prior model to the data using general linear models (GLM) and analysis-of-variance techniques. The resulting activation patterns are obtained in the form of statistical parametric or probability parametric maps [11, 15, 32, etc.].

b Exploratory methods like PCA [4, 37, etc.], ICA [25, etc.], CCA [12, etc.] or clustering techniques [13, 17, 26, 28, etc.]. These methods create spatial maps of voxels labeled based on the similarity/disparity of their time courses.

All these methods have a large number of parameters that need to be tuned, and that profoundly affects the detection of activated brain regions. Therefore, there is also a need for a tool that allows the neurologist to be able to visually assess the fidelity and veracity of any analysis result in order to find the best settings for the data-set under consideration.

However, in order to provide a meaningful visualization of the fMRI time signal, multiple problems have to be solved. The first challenge is extracting the BOLD signal (described below) at every voxel from the acquired data. The time-series data consists of four components:

- *Firdaus Janoos is with The Ohio State University, OH, USA, E-mail:*firdaus@ieee.org
- *Okan Irfanoglu is with The Ohio State University, OH, USA*
- *Raghu Machiraju is with The Ohio State University, OH, USA*
- *Istvan Akos Morocz is with Brigham and Women's Hospital, MA, USA*

Technical Report

1. A structural component, which measures the anatomical features of the brain (much like a conventional MR image). The majority of the signal energy is contained in this component, and it is constant across time.

2. The *blood oxygenation level dependent* (BOLD) signal, which measures the concentration of oxygenated blood in a voxel. It provides an indirect indication of brain activity, and depends upon the correlation between local neuronal activation and regional changes in blood flow and blood oxygenation levels [6]. While the exact mechanics linking neuronal activity to the BOLD signal is an active area of research [16], it appears that after an activation there is an initial drop in the levels of oxygenated hemoglobin in response to increased metabolic activity. This followed by a surge in the level of oxygenated blood, a peak and then a return to the baseline with a slight undershoot [18], lasting from 4 to 7 seconds. While the exact shape of this hemodynamic response function (HRF) is highly variable across brain regions, stimuli, and subjects [1], the shape of a theoretical (canonical) HRF is shown in Fig. 1. The BOLD signal constitutes only a few percent of the acquired fMRI signal energy.

3. Instrumental and physiological noise which tends to exhibit a $1/f$ spectrum associated with fractional Brownian motion (fBm) [7].

4. Baseline drifts, which are long term low frequency drifts in baseline signal intensity, the cause of which is not totally understood [34], and which is a major hurdle in any analysis of the data.

An example of a raw time course through an activated voxel is shown in Fig. 4(a) . This signal contains, along with the desired BOLD signal, the structural component and is corrupted by noise and baseline drift. Fig. 2 shows the one axial slice of the raw fMRI data at five consecutive time-points. From these figures it can be seen that visualizing the raw signal by itself is meaningless, and some method to extract the relevant information is needed. In this paper, we use a wavelet based estimation method to remove the noise, correct the baseline drift and extract the BOLD signal from the data, following a model proposed by Meyer [27].

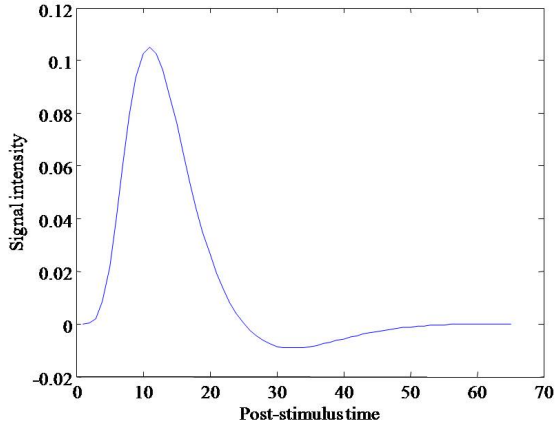

Fig. 1. The shape of the canonical hemodynamical response function (HRF) [1]. However, in reality, the exact shape of HRF is highly variable.

Another major problem is that of visual clutter. A typical fMRI volume consists of $64 \times 64 \times 32$ voxels, and making sense of the independent time-course through every voxel is an impossible task. Existing fMRI analysis tools allow graphing the estimated BOLD signal through an individual voxel (or the average BOLD signal through a user-specified VOI). However, by eliminating the spatial aspect of the signal, the utility of these visualizations for analysis of the regional cerebral blood flow (rCBF) is limited. Additionally, the activation time courses through similarly activated voxels are not exactly the same, and a temporally varying visualization of the per-voxel BOLD signals results in speckling and sparkling visual artifacts. To deal with these problem, we group voxels into coherent clusters based on the similarity of their time courses, and modulate their intensity with the representative time course of the cluster. Clustering is performed using a time-windowed hierarchical clustering algorithm in the wavelet descriptor space. An another advantage of this step, is that by clustering and averaging across similarly activated voxels, we are able to

improve the signal SNR. The natural number of clusters in data is identified using the Mixture Regression Criteria (MRC) [29].

The extracted BOLD signal is rendered as a time-varying cine overlaid on a structural MR image of the same subject, recreating the effect of oxygenated blood flow. The experimental conditions are also displayed in synchrony with BOLD signal, allowing the user to visually correlate experimental events with the hemodynamic response. This visualization helps the user to both validate experimental hypotheses and explore for ancillary effects. The BOLD signal, as extracted by our technique, is shown in Fig. 3, for the same axial slice at the same time points, as in Fig. 2.

The main contribution of this paper is a method and a tool for visualizing the measured BOLD signal from fMRI studies to help verify the activation maps generated by standard toolkits. Moreover, the system was designed to help develop an understanding of the temporal and spatial patterns of brain activity. To realize these goals, it was important that the techniques employed not add any additional uncertainty, in terms of restrictive assumptions on the shape of the hemodynamic response function, the linearity of the response, nature and number of the clusters, or through algorithms with a large number of user-tunable parameters. The wavelet based estimation, the hierarchical clustering and the model size selection methods used in this paper conform to extremely weak assumptions about the nature of the data and have very few parameters that need to be set, thereby adhering to these principles.

The organization of the paper is as follows: In Section 2, we go over current literature in the domain of spatio-temporal visualization, with special attention to medical images. The wavelet based estimation of the BOLD signal is presented in Section 3. Then in section 4, we describe our hierarchical clustering algorithm, with cluster selection using the MRC. Section 5 gives a description of the visualizations used and the tool. We then demonstrate the utility of the tool by going through a use case on a real data-set in Section 6. Finally in Section 7, we conclude with some remarks on the tool and directions of future investigation.

2 Related Work

There is a large body of work for visualizing data-sets with temporal dependencies. See [2] for a plenary survey on visualizing time-series data. Most of the prior art in visualization of spatio-temporal visualization has been concerned with the problem of dealing with massive data sizes and rendering them accurately and efficiently [24]. It has also been studied in the context of flow-field visualization [23, 36, 39], with the focus mainly on extracting the relevant features from the data and visualizing them in an interactive fashion. A discussion of these methods is out of the scope of the current work here.

In the case of medical data-sets, existing methods [3, 10, 31] concentrate on a comprehensive evaluation of the time-behavior of the data. They provide static visualizations by building transfer functions that represent the temporal characteristics of the data, thereby effecting a visual clustering/segmentation of the data in a temporal feature space. Even though we perform a feature extraction and clustering of the data, our aim is quite the opposite in that we want to render the time varying data itself, in order to allow the investigator to draw inferences about the underlying activity patterns. In the case of fMRI, there are many tools currently available for analysis and visualizing, including AFNI [8], BrainMap [22], Brede Toolbox [30], SPM [15], FSL [35]. However, none of them provide time-varying visualizations of the BOLD signal in the brain. Tory *et al.* [38] discuss several methods for visualizing time-varying medical data. They explore visualization methods, which include iso-surfaces, direct volume rendering and vector visualization using glyphs, to produce an animation of consecutive time steps that illustrate how the intensity and gradient change

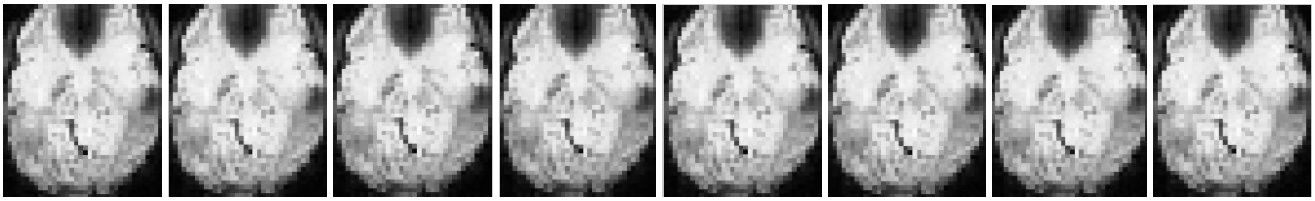

(a) $t = 2s$ (b) $t = 4s$ (c) $t = 6s$ (d) $t = 8s$ (e) $t = 10s$ (f) $t = 12s$ (g) $t = 14s$ (h) $t = 16s$

Fig. 2. An axial slice from the raw fMRI data at eight consecutive time points for the same brain image. From the images, it is extremely hard to deduce the activation patterns and the hemodynamic response function.

with time. We are presently investigating the application of such methods to our visualizations.

3 BOLD Signal Estimation

To relate the BOLD signal $s_{\mathbf{x}}(t)$ with the time course $y_{\mathbf{x}}(t)$ measured at T time points for each voxel $\mathbf{x}$, a partial linear model [27] is used:

$$y_{\mathbf{x}}(t) = \theta_{\mathbf{x}}(t) + s_{\mathbf{x}}(t) + \upsilon_{\mathbf{x}}(t), \qquad t = 1 \ldots N-1. \tag{1}$$

Here, $\theta_{\mathbf{x}}(t)$ is the baseline drift and $\upsilon_{\mathbf{x}}(t)$ is correlated noise with the $1/f$ spectrum of a long memory fBm noise [7].

The baseline drift is restricted to the wavelet coefficients at large scales, since it happens due to phenomena that vary gradually over relatively large periods of time as compared to the HRF [34]. By selecting a wavelet basis with p vanishing moments and assuming a polynomial approximation of the drift with order less than p, $\theta_{\mathbf{x}}(t)$ will have a sparse representation in this basis. Thus the drift can be modeled in the wavelet domain by:

$$\theta_{\mathbf{x}}(t) = c^{\theta}_{J,0}\Phi(2^{-J}t) + \sum_{j=J_0}^{J}\sum_{k=0}^{\frac{N}{2^j}-1} w^{\theta}_{j,k}\psi(2^{-j}t-k)\,. \tag{2}$$

Here, $\Phi(t)$ is the scaling function, $\psi(t)$ is the wavelet, while $c^{\theta}_{j,k}$ and $w^{\theta}_{j,k}$ are the scaling and wavelet coefficients of $\theta_{\mathbf{x}}(t)$ respectively, at scale j and translation k. J corresponds to the coarsest scale and typically $T = 2^J$. This model of baseline drift assumes that all the wavelet coefficients $w^{\theta}_{j,k}$ at scale $1 \le j \le J_0$ are zero. The Cohen-Daubechies-Feauveau $9/7$ bi-orthogonal spline wavelets with $p = 4$ vanishing moments are used because they are almost orthogonal, have finite support and linear phase and do not require special treatment at boundaries [27].

The wavelet transform provides an approximation of the Karhunen-Loève transform (KLT) for the fBm noise [40]. The KLT, which is an extension of the PCA to Hilbert spaces, de-correlates a random process by projecting it on to an orthogonal basis that are the eigenfunctions of the auto-covariance kernel. Specifically, if $K_y(t_1,t_2) = E[y(t_1)y^*(t_2)] - E[y(t_1)]E[y(t_2)]$ is the auto-covariance kernel of a stochastic process $y(t)$, then KLT is the decomposition of the process onto an orthogonal basis $e_i(t), i = 1\ldots\infty$, such that $y(t) = \sum_i Z_i e_i(t)$. Here $e_i(t)$ are the eigenfunctions of $K_y(t_1,t_2)$, and Z_i are uncorrelated random variables. Therefore, the wavelet transform $\mathscr{W}\{\upsilon_{\mathbf{x}}\}$ of the fBm noise $\upsilon_{\mathbf{x}}(t)$ is composed of coefficients which are *almost* de-correlated[1] [40]. The noise coefficients $[c^{\upsilon}_{J,0}, w^{\upsilon}_{J,0}, \ldots w^{\upsilon}_{1,T/2-1}]^{\mathbb{T}}$ are well approximated by normally distributed independent random variables, with co-variance matrix $\Sigma_{\upsilon} = \mathrm{diag}(\sigma_J^2, \ldots \sigma_1^2)$, which can be estimated at each scale j by a robust estimator [21].

Now applying the wavelet transform $\mathscr{W}$ to both sides of eqn. 1, we get:

$$\mathscr{W}\{y_{\mathbf{x}}\} = \mathscr{W}\{\theta_{\mathbf{x}}\} + \mathscr{W}\{s\} + \mathscr{W}\{\upsilon_{\mathbf{x}}\}, \tag{3}$$

Let $t_0 = 2^{-J_0+1}T$ be the non-zero wavelet coefficients of the baseline drift, i.e. $\mathscr{W}\{\theta_{\mathbf{x}}\} = [c^{\theta}_{J,0}, d^{\theta}_{J,0}, \ldots, d^{\theta}_{J_0,t_0-1}, 0\ldots 0]^{\mathbb{T}}$. Let $\mathbb{P}_{t_0}(\mathbf{z}) = [z_0, \ldots, z_{t_0-1}, 0, \ldots, 0]^{\mathbb{T}}$ be the projector of a vector $\mathbf{z}$ onto the first t_0 coordinates, and $\mathbb{Q}_{t_0}\mathbf{z} = [0, \ldots, 0, z_{t_0}, \ldots, z_{T-1}]^{\mathbb{T}}$ be the projector onto the last $T - t_0$ coordinates. Then, we can estimate $\hat{s} = \mathscr{W}^{-1}\left\{\mathbb{Q}_{t_0}\Sigma_{\upsilon}^{-1}\mathscr{W}\{y\}\right\}$ and $\hat{\theta} = \mathscr{W}^{-1}\left\{\mathbb{P}_{t_0}\Sigma_{\upsilon}^{-1}(\mathscr{W}\{y\} - \mathscr{W}\{\hat{s}\})\right\}$. Thus, by performing the processing in

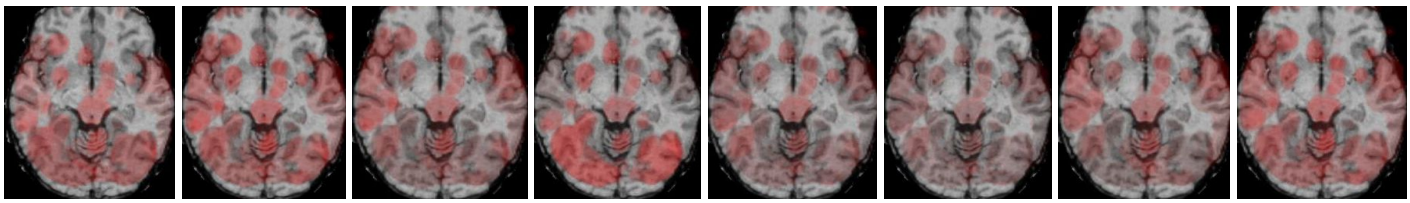

(a) $t = 2s$ (b) $t = 4s$ (c) $t = 6s$ (d) $t = 8s$ (e) $t = 10s$ (f) $t = 12s$ (g) $t = 14s$ (h) $t = 16s$

Fig. 3. The same axial slice of the structural MR image overlaid with the extracted BOLD signal at the same eight consecutive time points.

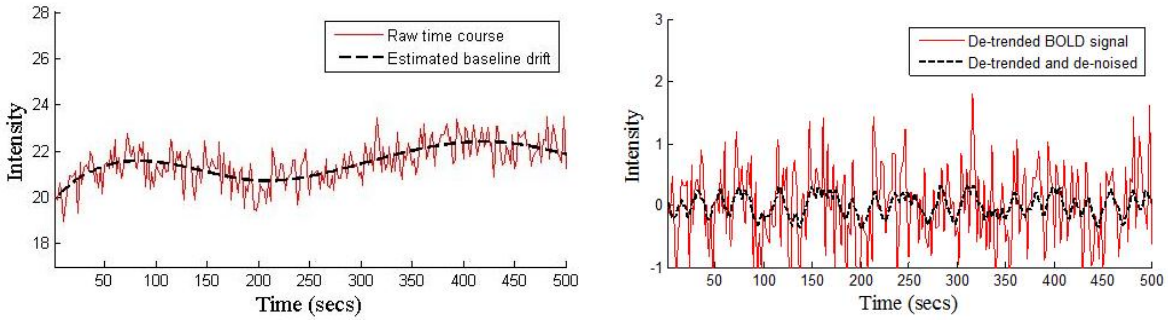

(a) Raw time course

(b) De-trended and de-noised time course

Fig. 4. Fig.(a) shows the raw time course (red, full), with an intensity offset due to the structural component in the fMRI signal, plus the baseline drift (black, dashed) and noise. Fig.(b) shows the de-trended BOLD signal with noise (red, full) and the de-trended and de-noised BOLD signal (black, dashed)

the wavelet domain, we can effectively remove the contamination of the baseline drift, as well as the long memory noise. The results of this step are shown in Fig. 4 for a single voxel.

The optimal value of J_0 is selected using a version of the Akaike Information Criteria corrected for small sample sizes [19], given as:

$$\mathrm{AIC}_c(t_0) = \log \hat{\sigma}^2 + \frac{T + t_0}{T - t_0 - 2}, \tag{4}$$

where, $\hat{\sigma}^2$ is the maximum likelihood estimate of the reconstruction error, given by $\hat{\sigma}^2 = 1/T||y - \hat{s}||$. The AIC_c criterion finds a balance between the error of the data due to a model (reconstruction error) and the complexity of the model (the number of non-zero wavelet coefficients). The desired $J_0 = 1 - \log_2 t_0/T$ occurs at the minimum of this function.

We have observed that there is a fairly large latitude both in the selection of p and J_0, which can be explained by the fact that the baseline drift occurs at scales significantly larger than the BOLD response.

4 CLUSTERING

Clustering of fMRI voxels has classically been used to find volumes of interest (VOIs) exhibiting similar time courses, and then to test these VOIs for activation. Our goal, on the other hand, is to replace the large number of independent BOLD signals with a few representative signals, primarily to reduce the visual clutter. A secondary purpose is to loosely group similarly activated voxels, thereby giving the investigator a general sense of regions that exhibit a similar hemodynamic response.

We first define a generalized distance metric between the estimated HRFs of two voxels $\mathbf{x_i}, \mathbf{x_j}$ as:

$$d^2(s_{\mathbf{x_i}}, s_{\mathbf{x_j}}) = \left[\mathscr{W}\{s_{\mathbf{x_i}}\} - \mathscr{W}\{s_{\mathbf{x_j}}\}\right]^{\mathbb{T}} \Lambda \left[\mathscr{W}\{s_{\mathbf{x_i}}\} - \mathscr{W}\{s_{\mathbf{x_j}}\}\right], \tag{5}$$

where $\Lambda = \mathrm{diag}(2^J, 2^{J-1}, ..., 1)$ provides a weighting towards larger scale wavelet coefficients, whereby accentuating differences between voxels at large signal scales, while suppressing differences at smaller scales. The representative signal s_c of a cluster $c = \{\mathbf{x_i}\}$ is defined as the one that minimizes the within-cluster variance, given as:

$$s_c = \arg\min_s \sum_i d\left(s_{\mathbf{x_i}}, s\right)^2, \tag{6}$$

$$\text{which gives} \quad s_c = \mathscr{W}^{-1}\left\{\sum_i \mathscr{W}\{s_{\mathbf{x_i}}\}\right\} \tag{7}$$

The clustering is then performed using an agglomerative hierarchical clustering algorithm as follows:

[1] Strictly, the correlation of the wavelet coefficients $w^{\upsilon}_{j,k}$ and $w^{\upsilon}_{j,k'}$ decays like $\mathscr{O}(|2^{-j}k - 2^{-j}k'|2^{-p})$.

i For each voxel $\mathbf{x}_i$, create one cluster c_i of weight $w_i = 1$. Initialize the cluster dissimilarity matrix $D_{i,j} = d^2(s_{\mathbf{x_i}}, s_{\mathbf{x_j}})$.

ii Find the two clusters c_i and c_j, such that $D_{i,j}$ is minimum. Merge the two clusters into a new cluster $c_n = c_i \bigcup c_j$ with weight $w_n = w_i + w_j$.

iii For all clusters c_k with $k \neq i, j$, update the dissimilarity matrix by the formula

$$D_{n,k} = \frac{(w_i + w_k)D_{i,k} + (w_j + w_k)D_{j,k} + (w_k)D_{i,j}}{w_i + w_j + w_k} \tag{8}$$

iv Remove clusters c_i and c_j from the set of clusters and repeat until there is only one cluster left.

Since the tool is to designed to visualize the measured BOLD signal with good fidelity, it is imperative that the algorithm be conservative and not over-cluster regions. To automatically select the correct level of clustering, the Mixture Regression Criteria (MRC) [29] is used. The MRC was designed to deal with the problems of over-clustering associated with other model selection methods like the AIC. The MRC formulation used here (eqn. 9) consists of two terms: the first measures the lack of fit (i.e. within-cluster variance), while the second penalizes a model with too many small clusters.

$$\mathrm{MRC}(K) = \sum_{k=1}^{K} n_k \log(\sigma_k^2) - 2\sum_{i=1}^{k} n_k \log \frac{n_k}{N}. \tag{9}$$

Here K is the number of clusters, n_k is the number of voxels in cluster c_k, and σ_k^2 is the variance of cluster c_k defined as $1/(n_k - 1)\sum_i d^2(s_{\mathbf{x}_i}, s_{c_k})$. This formulation of $\mathrm{MRC}(K)$ has been shown to have a distinct minimum at a finite K, and we choose this as the appropriate value for the number of clusters.

To capture the temporally transient nature of the brain response to experimental events (due to habituation, shift-of-attention, etc.), rather than cluster the BOLD signal across the entire duration of the experiment ($\approx$ 1200s in our data-sets), we cluster the BOLD signal across shorter durations, with a sliding time window. As the clustering is performed in the wavelet domain, the support of the largest scale wavelet significant to the HRF serves as good indicator for the window size. To elaborate, in the wavelet decomposition $\mathscr{W}\{\hat{s}_{\mathbf{x}}\}$ we determine the largest scale j at which the wavelet coefficient $d_{j,k}^{\hat{s}_{\mathbf{x}}}$ is significantly different from zero, for all voxels $\mathbf{x}$. This is done using a Student's two-sided t-test with significance level $\alpha = 0.05$ and $N-1$ degrees of freedom. The scale of this wavelet determines the temporal window size in which transient clustering is performed.

The hierarchical clustering algorithm has a complexity of $\mathscr{O}(N^2)$, where $N \approx \mathscr{O}(10^5)$. In order to reduce the computational burden, we cull all those voxels where the energy of the BOLD signal in the current time window is less than a certain threshold. This also serves to clean up the visualization by restricting it to voxels of significant activation. This threshold is selected manually and is the only parameter of the system. This step typically reduces the number of voxels by half. In spite of this, the clustering has to be performed off-line.

5 VISUALIZATION

The BOLD signal at every voxel is estimated (Section 3) and the signals are loosely clustered, in a time-windowed fashion, into a sparse set of clusters (Section 4). The independent BOLD signals are replaced by their cluster centroid, in order to reduce the visual speckling effect. This step has the added benefit of improving the reconstructed signal SNR and smoothness, by averaging across voxels in a cluster, in the wavelet domain.

The visualization (Fig. 5(a)) consists of a structural MR image in which the BOLD effect $s_{\mathbf{x}}(t)$ at voxel $\mathbf{x}$ at time point t modulates the intensity of the red color channel at that voxel. The visualization is displayed in the three orthogonal views (Axial, Coronal and Sagittal). and in a 3D display. Here, two modes are provided: (a) the volume sliced along the three orthogonal planes, or (b) a surface rendering of the cortical surface and internal neuro-anatomic structuresCortical surface extraction and skull stripping from the anatomical MR images was done using FreeSurfer [9,33] A time slider allows the user to change the time point t in the visualization. Alternatively, the time-series can be displayed in a cine mode, for the entire experimental duration or only for particular experimental events.

A second UI window shows the experimental conditions (external stimuli, subject responses, and experimental parameters [20]) as a set of graphs time-locked with BOLD effect display. The conditions are convolved with the

canonical hemodynamic response function (see Fig. 1), in order to help the user seek and understand correlations between the measured and expected responses (Fig. 5(b)). The experimental conditions for our data-set are described in Section 6.1. As the user steps through the BOLD signal visualization in time, this display advances in synchrony, giving the user a reference frame with respect to the experimental conditions.

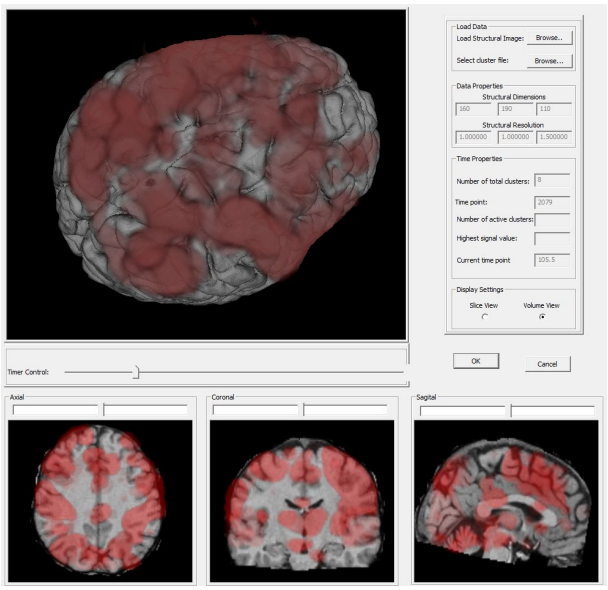

(a) BOLD signal visualization

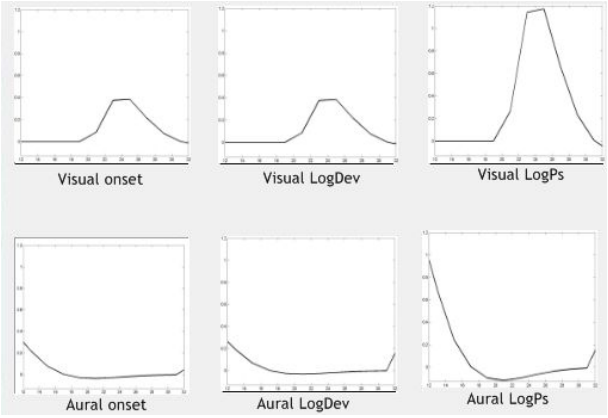

(b) Experimental conditions

Fig. 5. Fig.(a) shows the BOLD effect visualization. The top panel is 3D visualization along three orthogonal cutting planes. The bottom panel shows the image in three orthogonal views (Axial, Coronal and Sagittal). Fig.(b) shows the experimental conditions, graphed in synchrony with the BOLD effect visualization. The conditions are convolved with a canonical HRF (Fig. 1), in order to help the user understand correlations between the measured and expected responses.

6 Results

In this section we present a couple of examples of how this tool can help an investigator refine the analysis, validate the generated activation maps generated, and better understand the hemodynamic response to specific hypotheses. In this study, we used the SPM software[2], version 5, which is a standard toolkit for fMRI analysis. It uses a General Linear Model (GLM) description of the experimental paradigm to measure the response at each voxel to the experimental conditions. This is followed by statistical hypothesis testing (t and F–tests) to determine significant activation in voxels, which are displayed as statistical parametric maps. To analyze a data-set, a number of parameters need to be configured. For example, during a typical analysis session the following need to be set:

[2] `http://www.fil.ion.ucl.ac.uk/spm/`

i the experimental design, which can be specified in many different ways resulting in different kinds of analysis,

ii the model for the HRF (canonical HRF, a Fourier description of the HRF, a gamma function description or a finite impulse response description),

iii the intensity re-normalization parameters to deal with baseline drift and variations in image intensity over time,

iv the hypotheses to be tested, specified as contrast vectors for t or F-tests,

v the type of conjunction to be used, in case multiple hypotheses need to be tested simultaneous, all of which have significantly different meanings,

vi the α-value for statistical significance,

vii the method to correct for the Multiple Comparisons (MC) problem.

For more details, the reader is referred to [11]. Over here we will present two use-cases, one of how the tool can be used to select an appropriate α-value, and the other on how the tool can be used to explore hypotheses and temporal dependencies of brain activation.

6.1 Material

Our data-set was collected from a study where participants underwent functional and structural MRI while assessing the incorrectness of multiplication results. An irregular, self-paced experimental paradigm was used, consisting of repeated multiplication tasks. See Fig.6 for an illustration of a single trial. In each trial, subjects were exposed visually or aurally to simple multiplication problems with single digit operands, e.g. 4×5 , and had to decide if the solution subsequently offered was either: (a) *close* (within $\pm 25\%$ of the correct answer), (b) *too small* or (c) *too big* with respect to the correct result. Subjects answered by pressing a button with the index finger of the dominant hand for *too small*, the middle finger for *close*, and the ring finger for *too big*. Each trial of the experiment consisted of the following phases:

i the visual or aural presentation of the numbers for 2.5s,

ii an equal sign (=) for 0.3s followed by the incorrect result (0.8s),

iii a judgement period of variable duration up to 4s followed by the response event,

iv rest condition with fixation point of 1s until the beginning of the next trial.

Stimulus onset asynchrony (SOA) ranged from around 4s to a maximum of 8.6s. The aim of the experiment was to

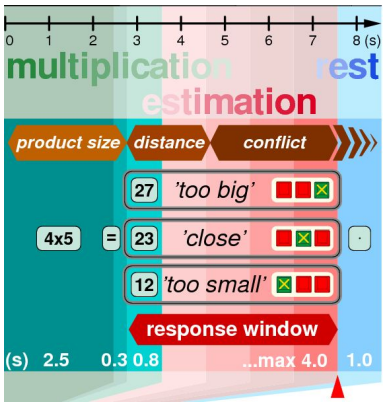

Fig. 6. Each trial consists of four distinct phases: i) the two numbers are either displayed or sounded (duration 2.5s), ii) a gap (0.3s), followed by an incorrect result (0.8s), iii) a variable duration interval of up to 4s for the subject to respond, iv) a rest duration of 1s before the start of the next trial.

identify the effect of the following factors on brain function during arithmetic processing:

i the size of the multiplication product (*LogPs* effect),

ii the distance of the displayed result from the correct solution (*LogDev* effect).

Acquisition was done on a General Electric 3-Tesla MRI scanner with a quadrature head coil using a 3D PRESTO BOLD pulse sequence with phase navigator correction and the following specifications: echo time 40ms, repetition time 26.4ms, echo train length 17, flip angle 17°, volume scan time 2.64s, and voxel size $3.75 \times 3.7 \times 4$mm. A typical session lasted from 15 to 18 minutes, with 130 trials. At the end of the study, an anatomical MR scan (sagittal 3D-SPGR) was acquired with a slice-thickness of 1.2mm and in-plane resolution of 0.94mm.

Raw data were reconstructed offline and routine pre-processing[3] was done in SPM5. The finite impulse response method of HRF modeling was used for three reasons: i) to enhance the temporal independence of the analysis, ii) to increase the sensitivity of cascading activations among brain regions, and iii) to optimally account for the irregular paradigm design and briefness of the rest-state condition.

6.2 Use Cases

For the first example, consider the question $\mathbb{Q}_1$:"Find those voxels activated only during visual presentation, and not during aural presentation". This hypothesis is tested as a one-sided t test with three different settings: (a) $\alpha = 0.01$ with no MC correction, (b) $\alpha = 0.001$ with no MC correction, (c) $\alpha = 0.01$ with the False Detection Rate correction method.

The resulting statistical parametric maps are shown in Fig. 7, as Maximum Intensity Projections (MIP) on three orthogonal views of a glass brain. Grey indicates regions of activation. The threshold for the map in Fig. 7(a) is too low and has resulted in too may false positives. It is very easy to confirm this by stepping through visualization of the BOLD signal in the entire brain, for aural and visual presentations separately and examining the regions that co-activate and those that do not. Even though the maps in Figs. 7(b) and (c) are better, there are still differences between the two, and some method to select between the two is required. For e.g. consider the area highlighted by the red arrow in Fig. 7(b). It is marked active in Fig. 7(b) but not in Fig. 7(c). However, by stepping through the BOLD visualization across multiple visual trials and aural trials separately, we see it being activated in both cases, albeit with a delay after number presentation. This observation is corroborated by the fact that the region (left dorsal pre-frontal cortex) is not known to be associated with visual tasks, but rather with number processing and working memory.

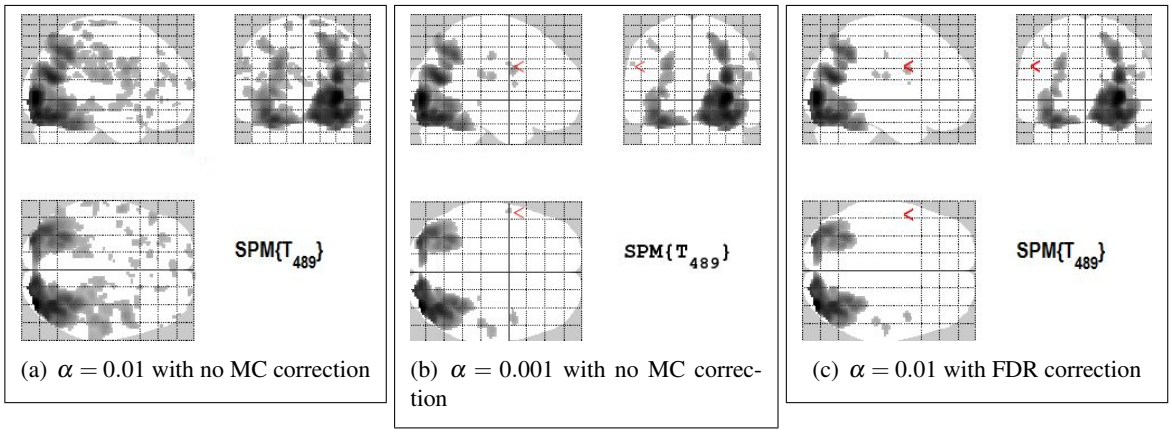

(a) $\alpha = 0.01$ with no MC correction

(b) $\alpha = 0.001$ with no MC correction

(c) $\alpha = 0.01$ with FDR correction

Fig. 7. MIPs of the activity maps for $\mathbb{Q}_1$ for three different settings of the significance level and MC correction. Darker shades of grey indicates higher levels of activation.

A second use case for the tool involves trying to visually validate the neuro-functional hypothesis that during the aural multiplication task, there is still a certain amount of activation in the visual cortices, but this is mainly a feedback effect. To validate this hypothesis, the regions of interest in the visual cortex are located in the tool, and those epochs corresponding to aural activations are played through. After a few runs through the data-set, one can see activations in

[3] spatial realignment for motion correction, spatial normalization to MNI brain space, co-registration of the functional and structural scans, and spatial smoothing with a 12mm Gaussian filter

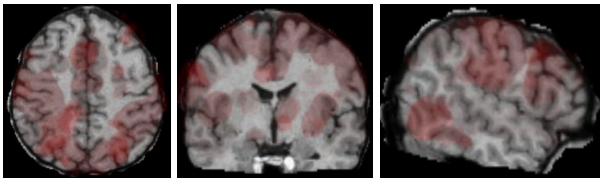

(a) Axial view (b) Coronal view (c) Sagittal view

Fig. 8. The BOLD signal visualizations at one time point of the axial (Fig. (a)), coronal (Fig. (b)) and sagittal planes (Fig. (c)) at the location indicated by the red arrow in Fig. 7(c).

the visual cortex that are temporally correlated with the aural presentation of the numbers. While this exercise does not rigorously confirm that the hypothesis is true, it does provide an indication of the effect and of its probable locations, which can be followed up by various statistical tests. Using this tool, it is also seen that while auditory and visual tasks differ significantly in multiple regions of the brain at the start of a trial, towards the middle and the end the functional response is similar, as expected.

7 Conclusion

In this paper, we have presented a method to visualize the time-varying BOLD signal acquired during neuro-imaging experiments, to address a need for studying brain function. The method consists of two main steps:

1. extracting the BOLD signal at every voxel from the raw fMRI data and correcting for noise and baseline drifts in signal intensity, using a wavelet based technique,

2. time-windowed hierarchical clustering which improves the visualization by reducing the number of independent components in the data and by increasing the SNR of the signal through selective averaging.

We presented two use-cases of the method highlighting its utility in validating the results generated by standard fMRI tools and in providing a mechanism for the exploration of the data and for the corroboration of hypotheses about brain function. We also demonstrated how the method can be used to understand the cascadic effects and temporal dependencies of brain activation. Such hypotheses are not easily tested and visualized with static activation maps.

We are currently investigating methods of locating cortical activations (as compared to whole brain) and displaying them on a surface rendering of the cerebral cortex, with and without flattening. We are also looking at techniques to integrate static activation maps and the time-varying BOLD effect into a single visualization scheme.

References

[1] G. K. Aguirre, E. Zarahn, and M. D'Esposito. The variability of human, BOLD hemodynamic responses. *NeuroImage*, 8(4), Nov 1998.

[2] W. Aigner, S. Miksch, W. Muller, H. Schumann, and C. Tominski. Visual methods for analyzing time-oriented data. *IEEE Trans. Vis. Comp. Graph.*, 14(1):47–60, 2008.

[3] H. Akiba, N. Fout, and K.-L. Ma. Simultaneous classification of time-varying volume data based on the time histogram. In *EuroVis*, pages 473 – 592, 2006.

[4] A. H. Andersen, D. Gash, and M. Avison. Principal component analysis of the dynamic response measured by fMRI: a generalized linear systems framework. *Mag. Res. Imag.*, Jul 1999.

[5] P. A. Bandettini, A. Jesmanowicz, E. C. Wong, and J. S. Hyde. Processing strategies for time-course data sets in fMRI of the human brain. *Magn Reson Med.*, 30(2):167–173, 1993.

[6] P. A. Bandettini, E. C. Wong, R. S. Hinks, R. S. Tikofsky, and J. S. Hyde. Time course EPI of human brain function during task activation. *Magn. Reson. Med.*, 25:390–397, Jun 1992.

[7] E. Bullmore, C. Long, J. Suckling, J.Fadili, G. Calvert, F. Zelaya, T.A.Carpenter, and M. Brammer. Colored noise and computational inference in neurophysiological (fMRI) time series analysis: Resampling methods in time and wavelet domains. *Hum. Brain Map.*, 12(2), Feb 2001.

[8] R. Cox. Software for analysis and visualization of functional magnetic resonance neuroimages. *Comp. Biomed. Res.*, 29:162–17, 1996.

[9] A. Dale, B. Fischl, and M. Sereno. Cortical surface-based analysis i: Segmentation and surface reconstruction. *NeuroImage*, 9:179–194, 1999.

[10] Z. Fang, T. Möller, G. Hamarneh, and A. Celler. Visualization and exploration of time-varying medical image data sets. In *GI '07: Proceedings of Graphics Interface 2007*, pages 281–288, New York, NY, USA, 2007. ACM.

[11] R. Frackowiak, K. Friston, C. Frith, R. Dolan, C. Price, S. Zeki, J. Ashburner, and W. Penny. *Human Brain Function*. Academic Press, 2nd edition, 2003.

[12] O. Friman, M. Borga, P. Lundberg, and H. Knutsson. Exploratory fMRI analysis by autocorrelation maximization. *Neuroimage*, 16:454–464, 2002.

[13] O. Friman, M. Borga, P. Lundberg, and H. Knutsson. Adaptive analysis of fMRI data. *NeuroImage*, 19(3), July 2003.

[14] K. Friston. *Brain Mapping: The Methods*, chapter Statistical Parametric Mapping and Other Analyses of Functional Imaging Data, pages 363–396. Acad. Press, 1996.

[15] K. Friston, A. Holmes, K. Worsley, J. Poline, C. Frith, and R. Frackowiak. Statistical parametric maps in functional imaging: A general linear approach. *Hum. Brain Map.*, 2(4):189–210, 1995.

[16] D. Glaser, K. Friston, A. Mechelli, and R. T. C. Price. *Human Brain Function*, chapter Hemodynamic modeling. Acad. Press, 2 edition, Dec 2003.

[17] C. Goutte, P. Toft, E. Rostrup, F. Nielsen, and L. Hansen. On clustering fMRI time series. *NeuroImage*, 9(3), March 1999.
[18] X. Hu, E. Yacoub, T. H. Le, E. R. Cohen, and K. Ugurbil. *Functional MRI*, chapter Functional MRI decrease at the onset of stimulation, pages 243–252. Springer-Verlag, 2 edition, Dec 1999.
[19] C. M. Hurvich and C.-L. Tsai. Regression and time series model selection in small samples. *Biometrika*, 76(2):297–307, 1989.
[20] P. Jezzard, editor. *Functional Mri: An Introduction to Methods*. Oxford Univ. Press, 2003.
[21] I. M. Johnstone and B. W. Silverman. Wavelet threshold estimators for data with correlated noise. *J. Roy. Stat. Soc.*, 59(2):319–351, 1997.
[22] J. Lancaster, E. Chan, S. Mikiten, S. Nguyen, and P. Fox. BrainMap^{tm} search and view. *NeuroImage*, 5:634, 1997.
[23] G. Li, U. Bordoloi, and H. W. Shen. Chameleon: An interactive texture-based framework for visualizing three-dimensional vector fields. In *IEEE Vis*, pages 241–248, 2003.
[24] K. L. Ma and E. B. Lum. *Visualization Handbook*. 2003.
[25] M. J. McKeown, S. Makeig, G. G. Brown, T.-P. Jung, S. S. Kindermann, A. J. Bell, and T. J. Sejnowski. Analysis of fMRI data by blind separation into independent spatial components. *Hum. Brain Map.*, 6:160–188, 1998.
[26] F. Meyer and J. Chinrungrueng. Local clustering of fMRI time series in the frequency domain. *Med. Image.Ana*, Feb 2005.
[27] F. G. Meyer. Wavelet-based estimation of a semiparametric generalized linear model of fMRI time-series. *IEEE Trans. Med. Imag.*, 22(3):315–322, Mar 2003.
[28] A. Meyer-Baese, A. Wismueller, and O. Lange. Comparison of two exploratory data analysis methods for fMRI: unsupervised clustering versus independent component analysis. *IEEE Trans. Inf. Tech. Biomed.*, 8(3):387–398, Sep 2004.
[29] P. A. NAIK, P. SHI, and C.-L. TSAI. Extending the Akaike Information Criterion to mixture regression models. *J. Am. Stat. Assc.*, 102(477), Mar 2007.
[30] F. Nielsen and L. K. Hansen. Experiences with Matlab and VRML in functional neuroimaging visualizations. Apr 2000.
[31] S. Oeltze, H. Doleisch, H. Hauser, P. Muigg, and B. Preim. Interactive visual analysis of perfusion data. *IEEE Trans. Vis. Comp. Graph.*, 2007.
[32] U. Ruttimann, M. Unser, R. Rawlings, D. Rio, N. Ramsey, V. Mattay, D. Hommer, J. Frank, and D. Weinberger. Statistical analysis of fMRI data in the wavelet domain. *IEEE Trans. Med. Imag.*, 17(2):142–154, 1998.
[33] F. Segonne, A. Dale, E. Busa, M. Glessner, U. Salvolini, H. Hahn, and B. Fischl. A hybrid approach to the skull-stripping problem in MRI. *NeuroImage (in press)*, 2007.
[34] A. Smith, B. Lewis, U. Ruttimann, F. Ye, T. Sinnwell, Y. Yang, J. Duyn, and J. Frank. Investigation of low frequency drift in fMRI signal. *NeuroImage*, 9(5), May 2003.
[35] S. Smith, M. Jenkinson, M. Woolrich, C. Beckmann, T. Behrens, H. Johansen-Berg, P. Bannister, M. D. Luca, I. Drobnjak, D. Flitney, R. Niazy, J. Saunders, J. Vickers, Y. Zhang, N. D. Stefano, J. Brady, , and P. Matthews. Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage*, 23(S1):208–219, 2004.
[36] A. Telea and J. J. van Wijk. 3d IBFV: Hardware-accelerated 3d flow visualization. In *IEEE Vis*, pages 233–240, 2003.
[37] B. Thirion and O. Faugeras. Dynamical components analysis of fMRI data through kernel PCA. *Neuroimage*, Sep 2003.
[38] M. Tory, N. Rober, T. Möller, A. Celler, and M. Atkins. 4d space-time techniques: a medical imaging case study. In *IEEE Vis*, pages 473 – 592, 2001.
[39] D. Weiskopf, F. Schramm, G. Erlebacher, and T. Ertl. Particle and texture based spatiotemporal visualization of time-dependent vector fields. In *IEEE Vis*, pages 639–646, 2005.
[40] G. W. Wornell. A Karhunen-Loéve-like expansion for $1/f$ processes via wavelets. *IEEE Trans. Info. Theory*, 36:859–861, Mar 1998.