STRUCTURED SPARSE DECONVOLUTION FOR PARADIGM FREE MAPPING OF FUNCTIONAL MRI DATA

César Caballero Gaudes[a,c], Fikret Işık Karahanoğlu[a,b], François Lazeyras[a], Dimitri Van De Ville[a,b]

[a] Department of Radiology and Medical Informatics, University of Geneva, CH-1211, Switzerland
[b]Medical Image Processing Lab, École Polytechnique Fédérale de Lausanne, CH-1015, Switzerland
[c]Basque Center on Cognition, Brain and Language, Donostia, 20009, Spain

ABSTRACT

Paradigm-free mapping enables to map the hæmodynamic response in space and time without prior knowledge of the timing of the underlying neuronal events (i.e., no stimulation paradigm). Such deconvolution approach can take advantage of modern sparsity-promoting regularization. Here we extend this concept using structured sparsity approaches in order to gain robustnesss against model mismatch. Specifically, we extend the hæmodynamic dictionary with the informed basis set (i.e., canonical HRF, and its temporal and dispersion derivatives) and we deploy state-of-the art structured sparsity functionals. In addition, we propose the group-weighted fusion penalty. We demonstrate the feasibility of the proposed approach for both synthetic and experimental data, showing superior abilities to characterize the single-trial BOLD response with no timing information.

Index Terms— Structured sparsity, brain imaging, functional MRI, paradigm free mapping.

1. INTRODUCTION

Functional magnetic resonance imaging (fMRI) enables to noninvasively map in space and time the hæmodynamic response following neuronal activations through the blood-oxygenation level dependent (BOLD) effect. Typical fMRI data analysis is performed with either confirmatory approaches to reveal voxels whose time series shows statistical evidence for a hypothetical task-related BOLD response, or exploratory methods, such as independent component analysis or clustering techniques, which explore fMRI data with no (or partial) information regarding the experimental conditions or the shape of the hæmodynamic response (see [1] for review).

There is an increasing interest in model-based methods that aim to identify neuronal events in BOLD fMRI time courses, but when no or insufficient information is available regarding the events' timings. Such approach becomes relevant, for instance, for the identification of interictal epileptic discharges or transient hæmodynamic events in resting state data. In essence, these methods attempt to deconvolve the neuronal-related signal underlying the BOLD response either assuming a linear model [2–6] or formulating a more complex, nonlinear dynamic representation of the BOLD effect [7–9]. This work is in line with the first group of linear deconvolution methods, where an inverse and ill-posed problem is formulated using a dictionary with shifted hæmodynamic response functions. Initially, deconvolution was done via L_2-norm regularization, such as ridge regression [3] or empirical bayesian estimators with Gaussian priors [6]. Recently, sparse-promoting regularization techniques were evaluated in sparse paradigm free mapping [2,4], using majorization-minimization techniques [5] or building new wavelet bases, termed activelets, that sparsify the neuronal-related hæmodynamic signal [10].

These methods make linear-system assumptions with a fixed hæmodynamic response function (HRF). Hence, any mismatch between the actual and modelled HRF could deteriorate the performance both in terms of prediction error and localization of the timing of the events [2]. Yet, the hæmodynamic response is known to vary across subjects, cortical regions and events [11], and this is compensated in model-based fMRI data analysis by explaining the BOLD response as a linear combination of temporal basis functions (e.g., the canonical HRF, and its temporal derivative and partial derivative towards the "dispersion" parameter [12]).

The aim of this work is to evaluate the use of structured sparsity to gain robustness against HRF mismatches in the deconvolution of the fMRI signal. To that end, we evaluate the performance of several recently proposed group-structured sparsity regularization functionals [13–15], which we solve using fast proximal gradient-based methods [16–18]. Note that the use of structured sparsity has already been proposed for voxel classification in fMRI brain decoding (e.g., see [19, 20]). The novelty of this work is to propose "structured paradigm free mapping"; i.e., to deconvolve the neuronal-related components of the fMRI signal without prior timing information.

The paper is organized as follows: In Section 2 we introduce the signal model and problem setting. Then, in Section 3 we describe the different algorithms investigated to solve our problem, whereas the results of our evaluations in synthetic and experimental data are presented in Section 4. Finally, we draw some conclusions.

2. PROBLEM FORMULATION

Let us consider that the fMRI signal of a voxel can be decomposed as $y(t) = x(t) + e(t)$, where $x(t)$ and $e(t)$ represent the neuronal-related hæmodynamic and noise components of the signal, respectively. The hæmodynamic signal $x(t)$ is commonly modelled with a linear time invariant system, $x(t) = h(t) * s(t)$, characterized by the HRF $h(t)$ and whose input signal $s(t)$ is related (but not equal) to the underlying neuronal signal. Here, we further assume that $s(t)$ can be modelled as a train of Dirac impulses at the fMRI timescale such that $x(t) = \sum_i s_i h(t - t_i) = \sum_i s_i h_i(t)$, where s_i is the amplitude of the hæmodynamic response with onset t_i earlier, and we define $h_i(t) = h(t - t_i)$. This event-related model is commonly adopted in fMRI experiments [12]. Sampling every TR seconds, the continuous-domain model can be written as $\mathbf{y} = \mathbf{x} + \mathbf{e} = \mathbf{H}\mathbf{s} + \mathbf{e}$, where N is the number of observations of the fMRI signal; $\mathbf{y}$, $\mathbf{s}$, $\mathbf{e}$

This work was supported in part by the Swiss National Science Foundation under grant PP00P2-123438, in part by Center for Biomedical Imaging (CIBM) of the Geneva-Lausanne Universities and the EPFL, and the Leenaards and Louis-Jeantet foundations, and in part by the grant CONSOLIDER -INGENIO2010 CSD2008-00048 from the Spanish Government.

$\in \mathbb{R}^N$; and $\mathbf{H} \in \mathbb{R}^{N\times N}$ is the convolution matrix (dictionary) with shifted HRFs. Note that the support of $\mathbf{s}$ (i.e., the set of non-zero coefficients) corresponds to those time points where the neuronal-related signal $s(t)$ exhibits non-zero amplitude at the fMRI resolution. Subsampling rates could be easily adopted with this formulation so that events can take place between two sampling times.

Contrary to previous deconvolution approaches that only consider a particular HRF to define $\mathbf{H}$ [2–6], we propose to describe the HRF as a linear combination of three temporal basis functions: the canonical HRF $h_c(t)$, its temporal derivative $h_t(t)$ and its dispersion derivative $h_d(t)$ [12], such that $h_i(t) = a_{c,i}h_{c,i}(t) + a_{t,i}h_{t,i}(t) + a_{d,i}h_{d,i}(t)$. The expanded model can be formulated as

$$\mathbf{y} = \tilde{\mathbf{H}}\tilde{\mathbf{s}} + \mathbf{e}, \tag{1}$$

where $\tilde{\mathbf{H}} = [\tilde{\mathbf{H}}_1, \ldots, \tilde{\mathbf{H}}_N] \in \mathbb{R}^{N\times 3N}$, and each submatrix $\tilde{\mathbf{H}}_i \in \mathbb{R}^{N\times 3}$, $i = 1, \ldots, N$, is defined as $\tilde{\mathbf{H}}_i = [\mathbf{h}_{c,i}\ \mathbf{h}_{t,i}\ \mathbf{h}_{d,i}]$; i.e., its columns are shifted replications of the canonical HRF, the temporal and dispersion derivatives. Equivalently, $\tilde{\mathbf{s}} \in \mathbb{R}^{3N}$ can be partitioned into N sub-vectors $\tilde{\mathbf{s}} = (\tilde{\mathbf{s}}_1, \ldots, \tilde{\mathbf{s}}_N)$, and each of the sub-vectors $\tilde{\mathbf{s}}_i = (\tilde{s}_{c,i}\ \tilde{s}_{t,i}\ \tilde{s}_{d,i})$ includes coefficients defined as $\tilde{s}_{\cdot,i} = a_{\cdot,i}s_i$. Finally, we consider that each $\tilde{\mathbf{H}}_i$ is orthonormalized; i.e., $\tilde{\mathbf{H}}_i^T\tilde{\mathbf{H}}_i = \mathbf{I}_3$, $i = 1, \ldots, N$, however, $\tilde{\mathbf{H}}_i^T\tilde{\mathbf{H}}_j$ is not the identity matrix for $i \neq j$ (sub-matrices at different time lags are not orthogonal to each other).

3. STRUCTURED SPARSE DECONVOLUTION

We will simplify our notation and remove the tilde from now on, but always refering to the expanded model in (1). Specifically, our aim is to deconvolve $\mathbf{s}$ by solving the following optimization problem

$$\mathbf{s}^* = \arg\min_{\mathbf{s}} J(\mathbf{s}) = \frac{1}{2}\|\mathbf{y} - \mathbf{H}\mathbf{s}\|_2^2 + \Omega(\mathbf{s}), \tag{2}$$

where $\Omega(\mathbf{s})$ is a regularization or penalty term that helps to reduce multicollinearity problems of the dictionary $\mathbf{H}$.

Our first choice for $\Omega(\mathbf{s})$ is the l_1-norm or LASSO penalty to encourage sparse estimates with few non-zero coefficients [21]:

$$\text{LASSO:}\quad \Omega(\mathbf{s}) = \lambda_1\|\mathbf{s}\|_1 = \lambda_1\sum_{i=1}^{3N}|s_i|, \tag{3}$$

where the regularization parameter λ_1 provides a tradeoff between data fidelity and sparsity. The LASSO tends to select only a few variables among a group of highly correlated variables, and disregards structural information in the signal model. Clearly, when a hæmodynamic event occurs at time i, the coefficients within the subvector $\mathbf{s}_i$ can be non-zero; otherwise, all of them should vanish. Motivated by this fact, the $l_{2,1}$ mixed-norm or Group LASSO penalty (G-LASSO) makes a reasonable choice:

$$\text{G-LASSO:}\quad \Omega(\mathbf{s}) = \lambda_1\|\mathbf{s}\|_{2,1} = \lambda_1\sum_{i=1}^{N}\|\mathbf{s}_i\|_2, \tag{4}$$

and $\|\mathbf{s}_i\|_2$ is the l_2-norm of each subvector $\mathbf{s}_i$. The G-LASSO penalty tends to promote sparsity across groups, while retaining l_2-norm regularization between the group coefficients [13]. Yet, one can also consider to penalize pairwise differences between highly correlated coefficients via correlation-driven weights. Let $\rho_{ij} = \mathbf{h}_i^T\mathbf{h}_j$ be the pairwise correlation between the columns of $\mathbf{H}$, then the Weighted Fusion (W-FUSION) penalty is defined as [14,15]

$$\text{W-FUSION: } \Omega(\mathbf{s}) = \lambda_1\|\mathbf{s}\|_1 + \lambda_2\sum_{i<j}\omega_{ij}(s_i - \alpha_{ij}s_j)^2, \tag{5}$$

where ω_{ij} are non-negative weights non-decreasing in $|\rho_{ij}|$; and $\alpha_{ij} = sgn(\rho_{ij})$ is the sign of ρ_{ij}. Note that when $\omega_{ij} = 0$, for all i, j, this penalty term reduces to the LASSO penalty in (3), and becomes the smooth LASSO [22] if only positive successive pairwise correlations are considered; i.e., $\omega_{ij} = 1$, for $j = i+1$, and $\omega_{ij} = 0$, otherwise. The use of weighted fusion encourages highly correlated variables to be jointly selected together, which is relevant in our case to deconvolve prolonged activations. Furthermore, the sign weight α_{ij} coordinates the direction in which correlated coefficients influence the fitted signal [14,15]. Yet, the original weighted fusion problem does not consider the grouped structure of our model. Therefore, in this work we propose the Group Weighted Fusion (GW-FUSION) penalty as

$$\text{GW-FUSION: } \Omega(\mathbf{s}) = \lambda_1\|\mathbf{s}\|_{2,1} + \lambda_2\sum_{i<j}\omega_{ij}(s_i - \alpha_{ij}s_j)^2, \tag{6}$$

which tends to form groups of highly correlated coefficients, but reduces them to zero when they are irrelevant to fit the fMRI signal.

Note that the common term of the weighted fusion penalties (5) and (6) admits a simple quadratic representation as $\mathbf{s}^T\mathbf{Q}\mathbf{s}$ where the entries of the matrix $\mathbf{Q} \in \mathbb{R}^{3N\times 3N}$ are given by [14,15]

$$q_{i,j} = \begin{cases} \sum_{i\neq j} w_{ij}, & \text{if } i = j, \\ -\alpha_{ij}w_{ij}, & \text{otherwise.} \end{cases} \tag{7}$$

The matrix $\mathbf{Q}$ is positive semi-definite and it admits a Cholesky decomposition $\mathbf{Q} = \mathbf{D}^T\mathbf{D}$. Hence, the functionals to solve are

$$\mathbf{s}^* = \arg\min_{\mathbf{s}} J(\mathbf{s}) = \frac{1}{2}\|\bar{\mathbf{y}} - \bar{\mathbf{H}}\mathbf{s}\|_2^2 + \lambda_1\|\mathbf{s}\|_1, \tag{8.1}$$

$$\mathbf{s}^* = \arg\min_{\mathbf{s}} J(\mathbf{s}) = \frac{1}{2}\|\bar{\mathbf{y}} - \bar{\mathbf{H}}\mathbf{s}\|_2^2 + \lambda_1\|\mathbf{s}\|_{2,1}, \tag{8.2}$$

where (8.1) corresponds to W-FUSION and LASSO for λ_2=0; (8.2) corresponds to the GW-FUSION and G-LASSO for λ_2=0; and we defined the augmented variables

$$\bar{\mathbf{y}} = \begin{pmatrix} \mathbf{y} \\ \mathbf{0} \end{pmatrix}, \qquad \bar{\mathbf{H}} = \begin{pmatrix} \mathbf{H} \\ \sqrt{\lambda_2}\mathbf{D} \end{pmatrix}. \tag{9}$$

3.1. Forward-backward splitting and monotone FISTA

We employ a forward-backward splitting approach [16] along with the fast monotone iterative shrinkage thresholding algorithm (M-FISTA) [17]. Forward-backward splitting allows the optimization of functionals with the sum of two convex functions, $J(\mathbf{s}) = f(\mathbf{s}) + g(\mathbf{s})$, where f is smooth and g can be non-smooth (see, e.g., [16–18]). In our case, $f(\mathbf{s}) = \|\bar{\mathbf{y}} - \bar{\mathbf{H}}\mathbf{s}\|_2^2/2$ and $g(\mathbf{s}) = \Omega(\mathbf{s})$, either the LASSO or G-LASSO penalties, so the functional $J(\mathbf{s})$ is convex. In the special case of f being the least squares term, the solution can be found with the popular Iterative Shrinkage Thresholding algorithm (ISTA) with two steps: 1) the forward step: $\mathbf{z}_k = \mathbf{s}_k + t\mathbf{H}^T(\mathbf{y} - \mathbf{H}\mathbf{s}_k)$, where t must be larger than the Lipschitz constant of f (i.e. $L = \rho(\mathbf{H}^T\mathbf{H})$, the spectral norm of $\mathbf{H}$); 2) the backward step: $\mathbf{s}_{k+1} = \text{prox}_t(g)(\mathbf{z}_k)$, where where $\text{prox}_t(g)(\mathbf{s})$ denotes the proximal map of g [16]. The proximal maps for the l_1-norm and the $l_{2,1}$-norm are the following thresholding operators [16, 18]

$$\text{prox}_t(\lambda\|\cdot\|_1)(\mathbf{s}) = s_i\max(0, 1 - t\lambda/|s_i|), \tag{10.2}$$

$$\text{prox}_t(\lambda\|\cdot\|_{2,1})(\mathbf{s}) = \mathbf{s}_i\max(0, 1 - t\lambda/\|\mathbf{s}_i\|_2). \tag{10.1}$$

Besides, we employ the M-FISTA algorithm to achieve faster and non-increasing convergence to the solution [17] instead of the simple ISTA algorithm. Algorithm 1 shows the pseudo-code of M-FISTA used in our approach.

Algorithm 1: Monotone FISTA algorithm

Input: $\mathbf{\Lambda} = \bar{\mathbf{H}}^T\bar{\mathbf{H}}$, $\mathbf{v} = \bar{\mathbf{H}}^T\bar{\mathbf{y}}$, $t \geq 1/\rho(\bar{\mathbf{H}}^T\bar{\mathbf{H}})$;
Initialization: $\mathbf{s}_0 = \mathbf{0}$, $\mathbf{u}_0 = \mathbf{0}$, $k = 0$, $t_0 = 1$;
repeat
  $\mathbf{z}_k \leftarrow \text{prox}_t(g)(\mathbf{u}_k + \gamma(\mathbf{v} - \mathbf{\Lambda}\mathbf{u}_k)))$;
  $t_{k+1} \leftarrow (1 + \sqrt{1 + 4t_k^2})/2$;
  if $J(\mathbf{z}_k) < J(\mathbf{z}_{k+1})$ **then**
    $\mathbf{s}_{k+1} \leftarrow \mathbf{z}_k$;
    $\mathbf{u}_{k+1} \leftarrow \mathbf{s}_k + \frac{t_k - 1}{t_{k+1}}(\mathbf{u}_k - \mathbf{s}_k)$;
  else
    $\mathbf{s}_{k+1} \leftarrow \mathbf{s}_k$;
    $\mathbf{u}_{k+1} \leftarrow \mathbf{s}_k + \frac{t_k}{t_{k+1}}(\mathbf{u}_k - \mathbf{s}_k)$
until *stopping criterion* ;
Output: $\mathbf{s}_k$

4. EXPERIMENTAL RESULTS

We assess the feasibility of using structured sparsity for the deconvolution of the hæmodynamic response in BOLD fMRI with both synthetic and experimental data. In all experiments, we used the canonical HRF, and its temporal and dispersion derivative with standard parameters in SPM8 [12]. For the weighted fusion penalties, the weights are defined as $\omega_{i,j} = |\rho_{i,j}|^{0.5}/(1 - |\rho_{i,j}|)$, which tends to equate the magnitudes of highly correlated variables as $\omega_{i,j} \to \infty$ when $|\rho_{i,j}| \to 1$ [14].

4.1. Synthetic data

One hundred fMRI simulated time series were created with a duration of 256 s at a temporal resolution (TR) of 1 s (N = 256). The hæmodynamic component $x(t)$ was created as a neuronal-related signal $s(t)$ including 6 ON periods of duration 0.2 s, 3 s or 6 s with onsets 10, 40, 100, 120, 190, and 230 s, convolved with a hæmodynamic response function $h(t)$ created by the weighted sum of the canonical HRF ($a_c = 1$), the temporal derivative ($a_t = 1.5$), and the dispersion derivative ($a_c = 0.5$). The amplitude of the hæmodynamic signal was then normalized for a maximum signal change of 6%. Finally, these time courses were corrupted with additive white gaussian noise with standard deviation $\sigma = 1/\text{SNR}$ according to three different temporal SNR conditions: low SNR = 30, middle SNR = 55, and high SNR = 80. These SNR values are typically observed at 3T and 7T fMRI experiments [2]. In case of correlated noise, morphological component analysis approaches for paradigm free mapping can be easily incorporated in our formulation [4].

The evaluation is done in terms of the mean squared error (MSE) of the estimates of the neuronal related signal $\mathbf{s}$ (MSE_s) and the neuronal-related hæmodynamic signal $\mathbf{x}$ (MSE_x). We evaluated two dictionaries with the LASSO: a dictionary based only on the canonical HRF (LASSO-1) and considering the three basis functions (LASSO-3). The regularization parameters were chosen according to an Oracle procedure minimizing MSE_s. For that, we defined a logarithmic grid with 128 values and range between $[0.1\hat{\sigma}_\text{MAD}, 50\hat{\sigma}_\text{MAD}]$, where $\hat{\sigma}_\text{MAD}$ is the median absolute deviance of the noise standard deviation after decomposing the fMRI time series with the Daubechies wavelets with 4 vanishing moments.

It can be seen in Table 1 that GW-FUSION outperforms the rest of penalty terms in all scenarios, except for the estimation of the hæmodynamic signal at SNR = 80 and event duration of 6 s where W-FUSION results in improved performance. LASSO-1 yields unsatisfactory operation proving its lack of sufficient degrees of freedom (provided here by the temporal and dispersion derivatives) to fit the shape of the HRF. Comparing LASSO-3 with G-LASSO, it is clear that these extra degrees of freedom must be included in a structured way as groups of coefficients. In general, incorporating additional structural information via the weighted fusion penalty further improves the deconvolution.

		SNR = 30		SNR = 55		SNR = 80	
		MSE_s	MSE_x	MSE_s	MSE_x	MSE_s	MSE_x
0.2 s	LA1	1.003	0.982	1.003	0.981	1.003	0.973
	LA3	1.000	0.980	1.000	0.988	1.252	0.210
	GLA	1.000	0.911	0.853	0.422	0.720	0.199
	WFU	1.000	0.980	1.000	0.986	1.160	0.193
	GWF	**0.977**	**0.803**	**0.827**	**0.361**	**0.706**	**0.192**
3 s	LA1	0.987	0.965	0.958	0.938	0.960	0.844
	LA3	0.981	0.896	0.939	0.701	0.918	0.585
	GLA	0.946	0.781	0.730	0.356	0.598	0.172
	WFU	0.973	0.845	0.927	0.672	0.906	0.549
	GWF	**0.882**	**0.688**	**0.641**	**0.305**	**0.523**	**0.169**
6 s	LA1	0.965	0.938	0.942	0.888	0.943	0.690
	LA3	0.987	0.938	0.960	0.706	0.975	0.343
	GLA	0.975	0.904	0.909	0.581	0.826	0.334
	WFU	0.974	0.793	0.934	0.631	0.944	**0.132**
	GWF	**0.942**	**0.720**	**0.845**	**0.404**	**0.771**	0.291

Table 1. Mean square error of the neuronal-related signal (MSE_s) and neural-related hæmodynamic signal (MSE_x). LA1: LASSO-1; LA3: LASSO-3; GLA: Group LASSO; WFU: Weighted Fusion; GWF: Group Weighted Fusion. Highest values per row in bold.

4.2. Experimental data

To evaluate our work in real conditions, we used fMRI data acquired from one subject performing a visual task in a Siemens Trio 3T MR scanner with a 32-channel head coil. The task involved 10 events of visual flickering checkerboard (duration 1 s, random onsets). At rest (no stimulus projected), the subject fixated eyes on a cross in the center of the screen. The fMRI data comprised N = 140 T2*-weighted gradient echo-planar images (TR/TE/FA= 2s/30ms/85°, voxel size = 3.25x3.25x3.5 mm^3). As for preprocessing, fMRI data were corrected for head motion, high pass filtered with a cutoff period of 128 s, and finally smoothed spatially with a 3D isotropic Gaussian filter (FWHM= 5 mm). Based on the results with synthetic data, we describe the results of GW-FUSION, GLASSO and LASSO-1. In the absence of Oracle information, the regularization parameters were set to $\lambda_1 = 4\hat{\sigma}_\text{MAD}$ for the three methods to achieve high specificity in the detection of the BOLD events, whereas $\lambda_2 = 5\hat{\sigma}_\text{MAD}$ for GW-FUSION to promote the grouping of correlated coefficients.

Figure 1 illustrates the results of the deconvolution in a voxel located in the primary visual cortex (see activation map). Interestingly, the fitted hæmodynamic signal obtained by the three methods are nearly identical (top). However, we can see that the GW-FUSION estimates of the neuronal-related coefficients (middle) delimit the onset of the hæmodynamic events (or stimuli) better than those obtained by LASSO-1 (bottom). The GLASSO coefficients, not shown in Figure 1, were nearly identical to those of GW-FUSION due to the high contrast to noise ratio of the BOLD events in the primary visual cortex and the use of a high value for λ_1. However, we observed that the specificity obtained with GLASSO rapidly deteriorates with lower values of λ_1 which is relevant to detect BOLD events in cortical areas with lower contrast-to-noise ratios. Note that the GW-FUSION coefficients of the temporal derivative are negative

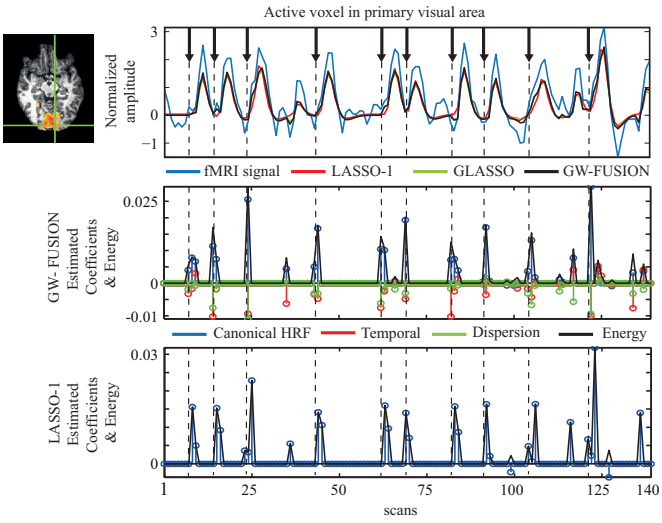


Fig. 1. Deconvolution obtained by GW-FUSION, GLASSO and LASSO. Vertical bars indicate the onset of the visual stimuli. Top: Preprocessed fMRI time series and hæmodynamic estimates. Middle and bottom: Coefficient estimates and the energy time series (i.e. the L_2-norm of the coefficients at each time point).

for most of the events, suggesting that the actual hæmodynamic response is slower (i.e., longer time-to-peak) than the canonical HRF. This type of characterization is not available if the model only includes the canonical HRF, and thus causes the LASSO-1 estimates to be slightly delayed with respect to the onset of the stimuli. Access to this type of information enables a more accurate characterization of the single-trial BOLD response, even without information about the timing of the events.

5. CONCLUSION

We showed that structured sparsity is a promising regularization for paradigm free mapping deconvolution of the fMRI signal. Structural information was defined in terms of groups of coefficients corresponding to basis functions (canonical HRF, temporal and dispersion derivatives) describing the BOLD response via the group LASSO, and their pairwise correlation via a weighted fusion. We proposed the group-weighted fusion, which resulted into the best performance among the functionals investigated in simulated data. Our results in real fMRI data demonstrated that structured sparsity enables better single-trial fMRI modelling. Future work will focus on developing efficient, data-driven algorithms to choose the regularization parameters and on investigating alternative structured sparsity penalties for paradigm free mapping.

6. REFERENCES

[1] G. E. Sarty, *Computing brain activity maps from fMRI time-series images*, Cambridge University Press, 2007.

[2] C. Caballero-Gaudes, N. Petridou, S. T. Francis, I. L. Dryden, and P. A. Gowland, "Paradigm free mapping with sparse regression automatically detects single-trial fMRI BOLD responses," *Hum. Brain Mapp.*, in press.

[3] C. Caballero-Gaudes, N. Petridou, I. L. Dryden, L. Bai, S. T. Francis, and P. A. Gowland, "Detection and characterization of single-trial fMRI BOLD responses: Paradigm free mapping," *Hum. Brain Mapp.*, vol. 9, pp. 1400–1418, 2011.

[4] C. Caballero-Gaudes, D. Van de Ville, N. Petridou, F. Lazeyras, and P. A. Gowland, "Paradigm free mapping with morphological component analysis: Getting most out of fMRI data," in *Proc. SPIE Vol 8138: Wavelets and Sparsity XIV*, 2011.

[5] L. Hernandez-Garcia and M. O. Ulfarsson, "Neuronal event detection in fMRI time series using iterative deconvolution techniques," *Magn. Reson. Imaging*, vol. 29, pp. 353–364, 2011.

[6] D. R. Gitelman, W. D. Penny, J. Ashburner, and K. J. Friston, "Modeling regional and psychophysiologic interactions in fMRI: The importance of hemodynamic deconvolution," *Neuroimage*, vol. 19, pp. 200–207, 2003.

[7] J. Riera, J. Watanabe, I. Kazuki, M. Naoki, E. Aubert, T. Ozaki, and R. Kawashima, "A state-space model of the hemodynamic approach: Nonlinear filtering of BOLD signals," *Neuroimage*, vol. 21, pp. 547–567, 2004.

[8] K. J. Friston, N. Trujillo-Barreto, and J. Daunizeau, "DEM: A variational treatment of dynamic systems," *Neuroimage*, vol. 41, pp. 849–885, 2008.

[9] M. Havlicek, K. J. Friston, J. Jiri, M. Brazdil, and V. D. Calhoun, "Dynamic modeling of neuronal responses in fMRI using cubature Kalman filtering," *Neuroimage*, vol. 56, pp. 2109–2128, 2011.

[10] I. Khalidov, M. J. Fadili, F. Lazeyras, D. Van De Ville, and M. Unser, "Activelets: Wavelets for sparse representation of hemodynamic responses," *Signal Processing*, vol. 91, pp. 2810–2821, 2011.

[11] J. R. Duann, T. P. Jung, W. J. Kuo, T. C. Yeh, S. Makeig, J. C. Hsieh, and T. J. Sejnowski, "Single-trial variability in event-related BOLD signals," *Neuroimage*, vol. 15, no. 4, pp. 823–835, 2002.

[12] K. J. Friston, P. Fletcher, O. Josephs, A. Holmes, M. D. Rugg, and R. Turner, "Event-related fMRI: Characterizing differential responses," *Neuroimage*, vol. 7, pp. 30–40, 1998.

[13] M. Yuan and Y. Lin, "Model selection and estimation in regression with grouped variables," *J. R. Statist. Soc. B*, vol. 68, pp. 49–67, 2006.

[14] Z. J. Daye and X. J. Jeng, "Shrinkage and model selection with correlated variables via weighted fusion," *Comput. Stat. Data Anal.*, vol. 53, pp. 1284–1298, 2009.

[15] M. Slawski, W. Z. Castell, and G. Tutz, "Feature selection guided by structural information," *The Annals of Applied Statistics*, vol. 4, pp. 1056–1080, 2010.

[16] P. L. Combettes and V. R. Wajs, "Signal recovery by proximal forward-backward splitting," *Multiscale Model. Simulat.*, vol. 4, pp. 1168–1200, 2005.

[17] A. Beck and M. Teboulle, "Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems," *IEEE Trans. Image Processing*, vol. 18, pp. 2419–2434, 2009.

[18] J. C. Baritaux, K. Hassler, S. Sanyal, and M. Unser, "Sparsity-driven reconstruction for FDOT with anatomical priors," *IEEE Trans. Medical Imaging*, vol. 30, pp. 1143–1153, 2011.

[19] R. Jenatton, A. Gramfort, V. Michel, G. Obozinski, F. Bach, and B. Thirion, "Multi-scale mining of fMRI data with hierarchical structured sparsity," *IEEE International Workshop on Pattern Recognition in NeuroImaging*, pp. 69–72, 2011.

[20] B. Ng and R. Abugharbieh, "Modeling spatiotemporal structure in fMRI brain decoding using generalized sparse classifiers," *IEEE International Workshop on Pattern Recognition in NeuroImaging*, pp. 65–68, 2011.

[21] R. Tibshirani, "Regression shrinkage and variable selection via the lasso," *J. R. Statist. Soc. B*, vol. 58, pp. 671–686, 1996.

[22] M. Hebiri and S. van de Geer, "The smooth lasso and other l_1+l_2-penalized methods," *Electron. J. Statist.*, vol. 5, pp. 1184–1226, 2011.