Paradigm-Free Mapping with Morphological Component Analysis: Getting Most Out of fMRI Data

César Caballero Gaudes[a], Dimitri Van de Ville[a,b], Natalia Petridou[c,d], François Lazeyras[a], Penny Gowland[c]

[a]Department of Radiology and Medical Informatics, University of Geneva and University Hospital of Geneva, Switzerland;
[b]Institute of Bioengineering, Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland;
[c]Sir Peter Mansfield Magnetic Resonance Centre, University of Nottingham, United Kingdom;
[d]Rudolf Magnus Institute/Radiology, University Medical Centre Utrecht, Netherlands;

ABSTRACT

Functional magnetic resonance imaging (fMRI) is a non-invasive imaging technique that maps the brain's response to neuronal activity based on the blood oxygenation level dependent (BOLD) effect. This work proposes a novel method for fMRI data analysis that enables the decomposition of the fMRI signal in its sources based on morphological descriptors. Beyond traditional fMRI hypothesis-based or blind data-driven exploratory approaches, this method allows the detection of BOLD responses without prior timing information. It is based on the deconvolution of the neuronal-related haemodynamic component of the fMRI signal with paradigm free mapping and also furnishes estimates of the movement-related effects, instrumental drifts and physiological fluctuations. Our algorithm is based on an overcomplete representation of the fMRI voxel time series with an additive linear model that is recovered by means of a L_1-norm regularized least-squares estimators and an adapted block coordinate relaxation procedure. The performance of the technique is evaluated with simulated data and real experimental data acquired at 3T.

Keywords: Functional Magnetic Resonance Imaging, Paradigm Free Mapping, Block Coordinate Relaxation, Morphological Component Analysis

1. INTRODUCTION

Since its introduction in the early 1990s, functional magnetic resonance imaging (fMRI) has become a widespread tool for neuroscientists and clinicians to study brain function.[1] fMRI is a non-invasive functional imaging technique that allows spatio-temporal mapping of local changes in blood magnetic susceptibility related to changes in cerebral blood flow, cerebral blood volume, and oxygen metabolism that occur due to increased neuronal activity, the blood oxygenation level dependent (BOLD) effect.

In general, the fMRI signal is modeled as the linear convolution of the haemodynamic response function (HRF) with an input signal related to the neuronal activity (e.g. by means of the known stimulus timecourse). Unfortunately, changes in the fMRI signal originating from neuronal activity are only of the order of few percent of the signal's amplitude. Besides, the fMRI signal also comprises instrumental and physiological confounds that hamper the detection of the BOLD responses. Hence, assuming a priori knowledge of the timing of the events and looking for evidence of stimulus-related activity is nowadays the prevalent approach for fMRI analyses. However, this information can be difficult to obtain under certain clinical, behavioural or experimental conditions. For instance, the study of brain activations associated to interictal epileptic discharges with concurrent EEG-fMRI where the onset of the activations in the fMRI GLM model is obtained from the discharges detected on the scalp EEG by an experimented electroencephalographer. Therefore, there is a growing need for fMRI data analysis techniques that avoid a-priori specification of the onsets of BOLD responses, and which allow the use of more unconstrained experimental paradigms.

If there is no information about when cortical activations occur or about the shape of the HRF, model free techniques, such as principal component analysis (PCA), independent component analysis (ICA), or clustering

Wavelets and Sparsity XIV, edited by Manos Papadakis, Dimitri Van De Ville, Vivek K. Goyal,
Proc. of SPIE Vol. 8138, 81381K · © 2011 SPIE · CCC code: 0277-786X/11/$18
doi: 10.1117/12.893920

methods, can be used to identify voxels showing consistent and common temporal or spatial patterns of activation.[2] Despite the common belief, model-based techniques have also been developed to study the brains response with no timing assumptions. To cope with certain flexibility in the onset and duration of the response, temporal and dispersion derivatives of the canonical HRF can be used.[3] Bayesian deconvolution of an asynchronous HRF model in which the image acquisition rate is jittered has also been proposed.[4,5] Nevertheless, these are semi-blind methods in that they still use the knowledge of the stimulus timing. Probabilistic frameworks based on Hidden Markov models[6,7] or change point theory methods[8–10] also have similar goals. More recently, dynamic deconvolution of nonlinear stochastic haemodynamic models[11,12] has attempted to provide an estimate of the neuronal-related signal without using timing information, along with the model parameters and physiological variables underlying the BOLD effect.[13–16] Nevertheless, these approaches are designed for the deconvolution of regional time series exhibiting high signal-to-noise ratios and thus their practical use for single voxel analysis is challenging. Finally, the combination of a new HRF-shaped wavelet basis, termed activelets, which sparsifies the BOLD signal, and nonlinear sparse search algorithms has yielded satisfactory results for the analysis of slow event-related fMRI data without prior information of the event timings.[17]

Recently, we have presented two methods that aim to map the cortical response to single-trial BOLD events with no prior information about their timing or location: Paradigm Free Mapping (PFM)[18] and Sparse Paradigm Free Mapping (SPFM).[19] These methods are based on the voxelwise deconvolution of the HRF assuming a linear haemodynamic model.[20] For PFM, the ridge regression estimator was used for the deconvolution and time points exhibiting significant changes in the deconvolved signal were detected based on a temporal t-statistic against a baseline state.[18] In contrast, Sparse Paradigm Free Mapping used the Dantzig Selector[21] to deconvolve the haemodynamic signal assuming a sparse activation model, enabling an automatic detection of the events and avoiding the need to define a baseline period (free from neuronal activity). Due to its sparsifying features, the Dantzig Selector reduces to zero those time points of the deconvolved signal where no BOLD events occur.[19]

In this paper, we propose a new extension to the family of Paradigm Free Mapping methods. This approach decomposes the fMRI signal into its multiple sources according to their morphological signatures. Similar to morphological component analysis (MCA),[22] the components are estimated using an adaptation of the Block Coordinate Relaxation (BCR) algorithm[23] where the fMRI signal at each voxel is recursively projected onto a set of dictionaries. Without prior timing models, our method extracts the neuronal-related haemodynamic component of the signal assuming it admits a sparse representation into a haemodynamic dictionary. Furthermore, it also furnishes estimates of the baseline fluctuations related to scanner drifts, physiological confounds and motion-related effects. Therefore, the algorithm can be understood as a generalization of partially linear models or semiparametric generalized linear models for fMRI time series.[24–26] As a proof of concept, the method is evaluated using realistic simulations of fMRI data and we demonstrate its usefulness with experimental data acquired at 3T.

2. PARADIGM FREE MAPPING WITH MORPHOLOGICAL COMPONENT ANALYSIS

2.1 Signal model

Let us consider that the fMRI signal of a voxel can be decomposed as $y(t) = x(t) + g(t) + e(t)$, where $x(t)$, $g(t)$ and $e(t)$ represent the neuronal-related haemodynamic component, the baseline fluctuations and random noise of the signal, respectively. The haemodynamic signal $x(t)$ is commonly assumed to be the output signal of a linear time invariant system, $x(t) = h(t) * s(t)$, characterized by the haemodynamic response function $h(t)$ and whose input signal $s(t)$ is related (but not equal) to the underlying neuronal signal. Sampling every TR seconds, the continuous-domain model outlined above can be written as

$$y_n = \sum_{d=0}^{D-1} h_d s_{d-n} + g_n + e_n, \qquad n = 1, \ldots, N, \tag{1}$$

where N is the number of observations of the signal, and D is the discrete duration (at TR resolution) of the haemodynamic response function. Typical values of TR range between 0.5 s and 3 s for whole brain coverage. Equivalently, the model in (1) can be written as $\mathbf{y} = \mathbf{H}\mathbf{s} + \mathbf{g} + \mathbf{e}$, where $\mathbf{y}$, $\mathbf{s}$, $\mathbf{g}$ and $\mathbf{e}$ are $N \times 1$ vectors, and

the matrix $\mathbf{H}$ of size $N \times N$ is the Toeplitz convolution matrix with shifted HRFs such that $\mathbf{x} = \mathbf{Hs}$. Note that traditional model-based techniques for fMRI data analysis test for the presence of a given haemodynamic regressor, which is modeled as the convolution of the HRF with a stimulus signal based on the stimulus or event times,[3,27] whereas our model just assumes the shape of the HRF without specifying a stimulus function.

On the other hand, baseline fluctuations consist of multiple confounding sources, such as scanner drifts, physiological cardiac and respiratory fluctuations or movement-related effects.[28] If each component were to be described by a unique morphological dictionary, we envision that the baseline term could be decomposed in P additive components as $\mathbf{g} = \sum_{i=1}^{P} \mathbf{g}_i = \sum_{i=1}^{P} \mathbf{\Phi}_i \boldsymbol{\alpha}_i$, where each component $\mathbf{g}_i = \mathbf{\Phi}_i \boldsymbol{\alpha}_i$ is expanded as a weighted sum of the parametric functions in the dictionary matrix $\mathbf{\Phi}_i$. Henceforth, we will assume P=3 components for the baseline.

The first component describes very low-frequency scanner drifts (below 0.01 Hz) due to hardware instabilities that can be described by a set of L_s pairwise uncorrelated Legendre polynomials.[29] We denote this first component as $\mathbf{g}_s = \mathbf{\Phi}_s \boldsymbol{\alpha}_s$, where $\mathbf{\Phi}_s$ is a $N \times L_s$ matrix (dictionary) comprising the Legendre polynomials, and $\boldsymbol{\alpha}_s$ is the $L_s \times 1$ vector with the unknown coefficients.

The second component describes residual movement-related artefacts that can be present in the fMRI signal even after perfect rigid-body realignment. In case of smooth motion, it is normally sufficient to model the motion-related effects by a linear combination of the 6 translation and rotation time series estimated during realignment. Similar to the scanner drifts, the motion-related component can be described as $\mathbf{g}_m = \mathbf{\Phi}_m \boldsymbol{\alpha}_m$, where $\mathbf{\Phi}_m$ is the $N \times 6$ matrix (dictionary) with the realignement parameters.[28,30]

The third component of the fMRI signal considered here describes physiological fluctuations (denoted as $\mathbf{g}_p$) with noise-like properties, caused by changes in respiration and the heartbeat, which explain a substantial percentage of the BOLD signal variance mainly in voxels located in cerebrospinal fluid (CSF) and gray matter (GM).[31–33] It is assumed that the physiological processes that give rise to the respiratory and cardiac fluctuations are quasi-periodic signals at the fundamental respiratory and cardiac frequencies and their harmonics. RETROICOR[31] is one of the methods most commonly used to remove physiological artefacts from fMRI data if the cardiac and respiratory phases are monitored during the scan acquisition. In RETROICOR, the physiological component to be removed is expressed in continous time as:

$$g_p(t) = \sum_{m=1}^{M} a_m^c \cos(m\phi_c(t)) + b_m^c \sin(m\phi_c(t)) + a_m^r \cos(m\phi_r(t)) + b_m^c \sin(m\phi_r(t)), \tag{2}$$

where M is the number of harmonics including the fundamental frequency, and $\phi_c(t)$ and $\phi_r(t)$ are the time-dependent cardiac and respiratory phases, respectively. Providing no abrupt changes in the cardiac and respiratory phases occurred during the scan acquisition, the physiological component in (2) can be written in discrete time as a linear combination of cosine and sine functions:

$$g_p(n) = \sum_{k=0}^{N-1} a_k \cos(\omega_k n) + b_k \sin(\omega_k n), \tag{3}$$

for $n = 0, \ldots, N-1$, and $\omega_k = 2\pi k/N$, and where only those coefficients corresponding to the respiratory and cardiac fundamental frequencies and their harmonics are nonzero. Following this assumption, in this work we propose modelling the physiological noise component of the signal as $\mathbf{g}_p = \mathbf{\Phi}_p \boldsymbol{\alpha}_p$, where $\mathbf{\Phi}_p$ is the $N \times N$ Discrete Fourier Transform (DFT) matrix and $\boldsymbol{\alpha}_p$ is the $N \times 1$ complex vector of the unknown DFT coefficients such that the physiological component would admit a sparse representation in the DFT dictionary.

In summary, our model for the fMRI signal is given as $\mathbf{y} = \mathbf{Hs} + \mathbf{\Phi}_s \boldsymbol{\alpha}_s + \mathbf{\Phi}_m \boldsymbol{\alpha}_m + \mathbf{\Phi}_p \boldsymbol{\alpha}_p + \mathbf{e}$. Our goal is estimating these four components of the fMRI signal ($\mathbf{x}$, $\mathbf{g}_s$, $\mathbf{g}_m$, and $\mathbf{g}_p$) from N observations with no prior information about the haemodynamic events, no physiological monitoring, and only using the dictionaries $\mathbf{H}$, $\mathbf{\Phi}_s$, $\mathbf{\Phi}_m$, and $\mathbf{\Phi}_p$. Note that this problem is underdetermined since the total number of unknowns (N for the haemodynamic component, L_s for the scanner drifts, 6 for the motion-related component and N for the physiological component) largely exceeds the number of observations N.

2.2 Estimation algorithm

In this work, we propose to estimate the components by solving the following constrained optimization problem:

$$\{\hat{\mathbf{s}}, \hat{\boldsymbol{\alpha}}_s, \hat{\boldsymbol{\alpha}}_m, \hat{\boldsymbol{\alpha}}_p\} = \min_{\mathbf{s},\boldsymbol{\alpha}_s,\boldsymbol{\alpha}_m,\boldsymbol{\alpha}_p} \|\mathbf{y} - \mathbf{H}\mathbf{s} - \boldsymbol{\Phi}_s\boldsymbol{\alpha}_s - \boldsymbol{\Phi}_m\boldsymbol{\alpha}_m - \boldsymbol{\Phi}_p\boldsymbol{\alpha}_p\|_2^2 + \lambda\|\mathbf{s}\|_1 + \lambda\|\boldsymbol{\alpha}_\mathbf{p}\|_1, \tag{4}$$

where the non-negative real regularization parameter λ balances the model imperfection measured in terms of the residual sum of squares, and the sparsity of the representation of the neuronal-related haemodynamic and physiological components in the haemodynamic and DFT dictionaries.

Since no constraints have been imposed on the vectors characterizing the scanner drifts $\boldsymbol{\alpha}_s$ and motion-related effects $\boldsymbol{\alpha}_m$, these components can be combined as $\mathbf{g}_{sm} = \mathbf{g}_m + \mathbf{g}_m = \boldsymbol{\Phi}_s\boldsymbol{\alpha}_s + \boldsymbol{\Phi}_m\boldsymbol{\alpha}_m = \boldsymbol{\Phi}_{sm}\boldsymbol{\alpha}_{sm}$, where $\boldsymbol{\Phi}_{sm} = [\boldsymbol{\Phi}_s\ \boldsymbol{\Phi}_m]$ and $\boldsymbol{\alpha}_{sm}^T = [\boldsymbol{\alpha}_s^T\ \boldsymbol{\alpha}_m^T]$. In contrast, penalty terms in terms the L_1-norm are used for the haemodynamic component $\mathbf{s}$ and the physiological component $\boldsymbol{\alpha}_p$ to favour sparse estimates of these vectors. In first place, our model considers that the physiological source has a sparse representation in the DFT dictionary. Secondly, we assume that in slow event-related paradigms or single-trial BOLD responses, the neuronal-related signal $\mathbf{s}$ is a sparse vector at the fMRI timescale where few coefficients have amplitudes significantly different from zero. In other words, the haemodynamic component of the fMRI signal admits a sparse representation in the dictionary $\mathbf{H}$ characterized by the HRF shape.

To solve the optimization problem in (4), we use an adaptation of the Block Coordinate Relaxation (BCR) algorithm[23] and morphological component analysis,[22,34] similar to the procedure proposed in (Fadili and Bullmore, 2005).[24] The algorithm proceeds as follows. Let $x^{(i)}$ denote the value of an arbitrary variable x at iteration i. The vector estimates are initialized (at iteration i=0) as $\hat{\mathbf{s}}^{(0)} = \mathbf{0}$, $\hat{\boldsymbol{\alpha}}_p^{(0)} = \mathbf{0}$, and $\hat{\boldsymbol{\alpha}}_{sm}^{(0)} = (\boldsymbol{\Phi}_{sm}^T\boldsymbol{\Phi}_{sm})^{-1}\boldsymbol{\Phi}_{sm}^T\mathbf{y}$ (i.e. the orthogonal least squares (OLS) estimator of $\boldsymbol{\alpha}_{sm}$ given the original fMRI voxel time-series). Next, the residuals of the OLS projection onto the subspace spanned by $\boldsymbol{\Phi}_{sm}$ (i.e. $\mathbf{r}^{(0)} = \mathbf{y} - \boldsymbol{\Phi}_{sm}\hat{\boldsymbol{\alpha}}_{sm}^{(0)}$) are decomposed into the haemodynamic and DFT dictionaries such that the initial value of the regularization parameter is set to the maximum absolute coefficient considering both decompositions, i.e. $\lambda^{(0)} = \max\{\|\mathbf{H}^T\mathbf{r}^{(0)}\|_\infty, \|\boldsymbol{\Phi}_p^T\mathbf{r}^{(0)}\|_\infty\}$, which enforces the estimates of $\hat{\mathbf{s}}^{(0)}$ and $\hat{\boldsymbol{\alpha}}_p^{(0)}$ to be zero as initialized.

After initialization, the algorithm proceeds recursively by updating the solutions to the optimization problem in (4) with decreasing values of λ. At each iteration, an inner loop with $\log_2 N$ iterations is repeated to guarantee convergence of the estimates for the corresponding λ. First, update $\hat{\boldsymbol{\alpha}}_p$ by solving the following Basis Pursuit Denoising (BPDN) problem[35] assuming that $\hat{\mathbf{s}}$ and $\hat{\boldsymbol{\alpha}}_{sm}$ are fixed:

$$\hat{\boldsymbol{\alpha}}_p^{(i+1)} = \min_{\boldsymbol{\alpha}_p} \|\mathbf{y} - \mathbf{H}\hat{\mathbf{s}}^{(i)} - \boldsymbol{\Phi}_{sm}\hat{\boldsymbol{\alpha}}_{sm}^{(i)} - \boldsymbol{\Phi}_p\boldsymbol{\alpha}_p\|_2^2 + \lambda^{(i)}\|\boldsymbol{\alpha}_p\|_1, \tag{5}$$

which can be efficiently computed via soft-thresholding since the DFT is an orthonormal basis. Since the rest of components have been removed from the signal, the BPDN estimator in (5) will choose the most salient coefficients describing sinusoidal fluctuations. Second, the vector $\hat{\mathbf{s}}$ is then updated by solving the following BPDN problem assuming that $\hat{\boldsymbol{\alpha}}_p$ and $\hat{\boldsymbol{\alpha}}_{sm}$ are fixed:

$$\hat{\mathbf{s}}_p^{(i+1)} = \min_{\mathbf{s}} \|\mathbf{y} - \mathbf{H}\hat{\mathbf{s}} - \boldsymbol{\Phi}_{sm}\hat{\boldsymbol{\alpha}}_{sm}^{(i)} - \boldsymbol{\Phi}_p\hat{\boldsymbol{\alpha}}_p^{(i+1)}\|_2^2 + \lambda^{(i)}\|\mathbf{s}\|_1. \tag{6}$$

Since the haemodynamic dictionary $\mathbf{H}$ is not orthonormal, solving (6) cannot be done via soft-thresholding. Here, we use a primal-dual pursuit homotopy procedure[36] to solve (6), but note that other convergent algorithms could also be used without loss of generality of the procedure. Third, the estimate of $\boldsymbol{\alpha}_{sm}$ is updated via OLS assuming that $\hat{\mathbf{s}}$ and $\hat{\boldsymbol{\alpha}}_p$ are fixed:

$$\hat{\boldsymbol{\alpha}}_{sm}^{(i+1)} = (\boldsymbol{\Phi}_{sm}^T\boldsymbol{\Phi}_{sm})^{-1}\boldsymbol{\Phi}_{sm}^T(\mathbf{y} - \mathbf{H}\hat{\mathbf{s}}^{(i+1)} - \boldsymbol{\Phi}_p\hat{\boldsymbol{\alpha}}_p^{(i+1)}). \tag{7}$$

A convergence criterion is evaluated at the end of each loop (see section 2.3) and if not fulfilled, the regularization parameter λ is reduced and steps (5-7) are repeated for the new value of λ. Here, we choose to exponentially decrease the value of λ from its initial value $\lambda^{(0)}$ to a minimum value, λ_{min}, in N_{max} iterations, such that $\lambda^{(i+1)} = \delta\lambda^{(i)}$ where $\delta = (\lambda^{(0)} - \lambda_{min})/(N_{max} - 1)$. Therefore, the accuracy of the estimates and the speed of the algorithm depend on N_{max} and the range of λ considered. The smaller the step factor δ is

(either using larger N_{max} or reducing the difference between $\lambda^{(0)}$ and λ_{min}), the finer the minimization is, but the slower the algorithm converges. In addition, the value of λ_{min} must be chosen such that the condition of the convergence criterion is met. Here, λ_{min} is set to half of the Maximum Absolute Deviance estimate of the noise standard deviation[37] after decomposing the fMRI time series with a Daubechies wavelet with 2 vanishing moments.[38]

2.3 Convergence criterion and debiasing

Since the aim of our approach is to decompose the fMRI signal into its morphological components, the residuals after convergence would ideally become a stochastic, white noise process with an approximately flat power spectrum and an approximately linear cumulative power spectrum. Hence, at the end of each iteration, under the null hypothesis of white residuals, a test based on the linearity of the cumulative power spectrum of the best linear unbiased scalar covariance residuals (BLUS)[39,40] is used to decide when to stop the convergence. A detailed description of the test statistic and the corresponding approximate p-value can be found elsewhere;[40] we use the software implementation of the test that is available in the SPMd (Statistical Parametric Mapping Diagnosis) extension of SPM (FIL/UCL, UK). We propose stopping the convergence of the algorithm when the p-value for the test statistic was above 0.01.

Finally, after convergence is achieved, a final debiasing step is performed in order to overcome the tendency of estimators with L_1-norm regularization terms to underestimate the true value of the nonzero coefficients of $\hat{\mathbf{s}}$ and $\hat{\boldsymbol{\alpha}}_p$.[21] In other words, we rely on the iterative adapted BCR algorithm described above to select those coefficients that better explain each component according to its morphological dictionary, and then update the weights of these coefficients according to a least squares criterion to minimize estimation bias, keeping the non-selected coefficients as zero.

In computing the BLUS residuals and debiasing, the model matrix is defined as $\tilde{\mathbf{X}}^{(i)} = [\mathbf{H}^{(i)}\ \boldsymbol{\Phi}_s\ \boldsymbol{\Phi}_m\ \hat{\mathbf{g}}^{(i)}]$, where $\mathbf{H}^{(i)}$ is the matrix including the subset of columns of $\mathbf{H}$ corresponding to these nonzero coefficients of $\hat{\mathbf{s}}^{(i)}$. In general, using $\hat{\mathbf{g}}^{(i)}$ results in better operation of the convergence criterion instead of using the corresponding matrix $\boldsymbol{\Phi}_p^{(i)}$ (i.e. the subset of columns of the dictionary $\boldsymbol{\Phi}_p$ corresponding to these nonzero coefficients of $\hat{\boldsymbol{\alpha}}_p^{(i)}$).

3. RESULTS: SYNTHETIC DATA

The performance of the method is first evaluated on fMRI simulated time series of duration at 256 s with a temporal resolution (TR) of 1 s (i.e. $N = 256$). Each time series was created as the sum of a neuronal-related haemodynanamic component, a motion-related component, a physiological component, a signal with unit constant amplitude, and uncorrelated Normal noise representing thermal noise. The haemodynamic component was created as a neuronal-related signal including 6 events of duration of 4 s at onsets 10, 40, 100, 120, 190, and 230 s, convolved with the two gamma-variate canonical HRF:[3,27]

$$h(t) = g(t, \tau_1, l_1) - \frac{1}{6} g(t, \tau_2, l_2); \quad \text{where} \quad g(t, \tau, l) = \frac{l^\tau t_+^{\tau-1} e^{-l\tau}}{\Gamma(\tau)}, \tag{8}$$

with standard SPM parameters (i.e. the time-to-peak (τ/l) and dispersion (τ/l^2) are 6 s and 1 s for the initial peak, and 16 s and 1 s for the undershoot, respectively, and the total duration of the HRF is 32 s corresponding to a discrete length of D=33 time points). The neural-related input signal and the HRF were initially simulated at intervals of 100 ms, then downsampled to TR of 1 s. The motion-related component consisted of a weighted sum of the 6 timecourses of the translation and rotation parameters estimated in the realignment of an experimental fMRI dataset, where the weights where randomly generated following a Normal distribution with zero mean and unit variance. Following the RETROICOR model for the physiological fluctuations in (2), the physiological component was simulated as

$$g_p(t) = \sum_{m=1}^{2} \frac{1}{2^{i-1}} (\sin(2\pi f_{r,i} t + \phi_{r,i}) + \sin(2\pi f_{c,i} t + \phi_{c,i})). \tag{9}$$

The frequencies of each sinusoid were randomly generated following Normal distributions as $f_{r,i} \sim \mathcal{N}(if_r, 0.04)$ and $f_{c,i} \sim \mathcal{N}(if_c, 0.04)$, where the fundamental frequencies were $f_{r,1}$=0.3 Hz for the respiratory component,[32]

and $f_{c,1}$=1.1 Hz for the cardiac component,[33] and the phases were randomly selected from an uniform distribution between 0 and 2π radians. Prior to adding the physiological, motion-related and thermal components to the haemodynamic component, they were normalized to unit standard deviation and weighted according to a comprehensive model of the noise in fMRI experiments.[41,42]

In fMRI studies, the temporal signal-to-noise ratio (tSNR) is defined as tSNR=S/σ where S is the mean of the signal in the absence of events and physiological or systematic trends, and σ is the standard deviation of the time series. In our simulation model, the mean of the fMRI signal was set as $S = 1$, whereas the signal variability is $\sigma^2 = \sigma_p^2 + \sigma_m^2 + \sigma_0^2$, where σ_p^2 is the variance of the physiological component, σ_m^2 of the motion component, and σ_0^2 of the thermal noise.[42] We simulated a tSNR = 60 (i.e. $\sigma = 0.0167$), and considered the following distribution for the variance of the components: $\sigma_p^2 = 0.0105$, $\sigma_m^2 = 0.0024$, $\sigma_0^2 = 0.0127$, based on realistic relationships between the components.[42,43] In addition to the tSNR, the detection of the BOLD responses also depends on the contrast-to-noise ratio of the BOLD response defined as CNR = tSNR($\Delta S/S$), where ΔS is the BOLD signal change due to the haemodynamic events.[41,42] After normalization of the haemodynamic signal to a maximum amplitude of 1, the BOLD events were scaled to have a signal change ΔS of 4%, typically observed at 3T MR scanners,[44] resulting in a CNR of 2.4.

The results for one simulated fMRI time series are illustrated in Figure 1. It can be seen that the method is able to provide accurate estimates of the neuronal-related haemodynamic (second row), physiological (third row), motion-related signals (fourth row) and the constant signal which, when summed, provide a good denoised version of the fMRI signal (first row). Using sparse regression for the recovery of the haemodynamic and physiological components provides well-tuned estimates of the haemodynamic events without prior information of their timing and amplitude, and the fundamental harmonics of physiological component (the component at 0.3 Hz corresponds to the respiratory signal, whereas the one at 0.1 Hz corresponds to the aliased cardiac component) without prior information on their frequencies. In contrast, the harmonic at 0.2 Hz is only captured with very low amplitude, whereas the harmonic at 0.4 Hz is not detected, probably due to the fact that the percentage of both harmonics on the total fMRI signal variance is not sufficient to be accurately estimated. Furthermore, the algorithm is able to recover the motion-related signal despite nearly exhibiting a flat spectrum, benefiting from the morphological information about this component that is available in the realignment parameters.

Figure 2 shows the results averaged over 1000 realisations of the previous fMRI simulated time series with varying physiological component and thermal noise. Each plot displays the true simulated signals, and the median and confidence interval of the estimates defined by the 50% , and 10% and 90% percentiles respectively. It can be observed that debiasing the estimates (right column) helps to recover the amplitude of the haemodynamic events and the sinusoidal harmonics better, and adjust the trajectory of the motion-related signal better than not performing the final debiasing step. Nevertheless, if the initial estimates are incorrect (i.e. the algorithm detects false events or sinusoids), the estimation error increases with debiasing. This drawback is mainly observed in the form of errors in the estimate of the haemodynamic component (top figure, right column) even though, on average (median), the recovered haemodynamic signal closely follows the simulated one. Interestingly, errors in the estimation of the simulated haemodynamic signal are located in the neighbourhood of the real events, which might reflect an interaction between all components that leads to offsets in the recovery of the event onset. Finally, the confidence interval for the motion-related component is larger than that for the other components. This can be explained due to the small percentage of the fMRI signal attributed to the motion-effects in our simulations since the variance of the motion-related component is 5% of the variance of the physiological component. One can expect that the larger the ratio of signal variance explained by a component, the more accurate its estimate, providing it can be adequately described by its morphological dictionary and with negligible collinearity between the dictionaries.[34]

4. RESULTS: FMRI EXPERIMENTAL DATA

As a proof of concept, the usefulness of the algorithm was tested in an fMRI experimental dataset obtained after scanning one subject (27 years, male) in a Siemens TIM Trio 3T MR system with a 32-channel head coil. The paradigm consisted of 10 events of visual flickering checkerboard (8Hz frequency) with duration 1 s and random onsets. Visual stimuli were projected from the back of the scanner room onto a screen that the subject observed with a mirror. When no visual stimuli were presented, the subject was instructed to maintain visual fixation

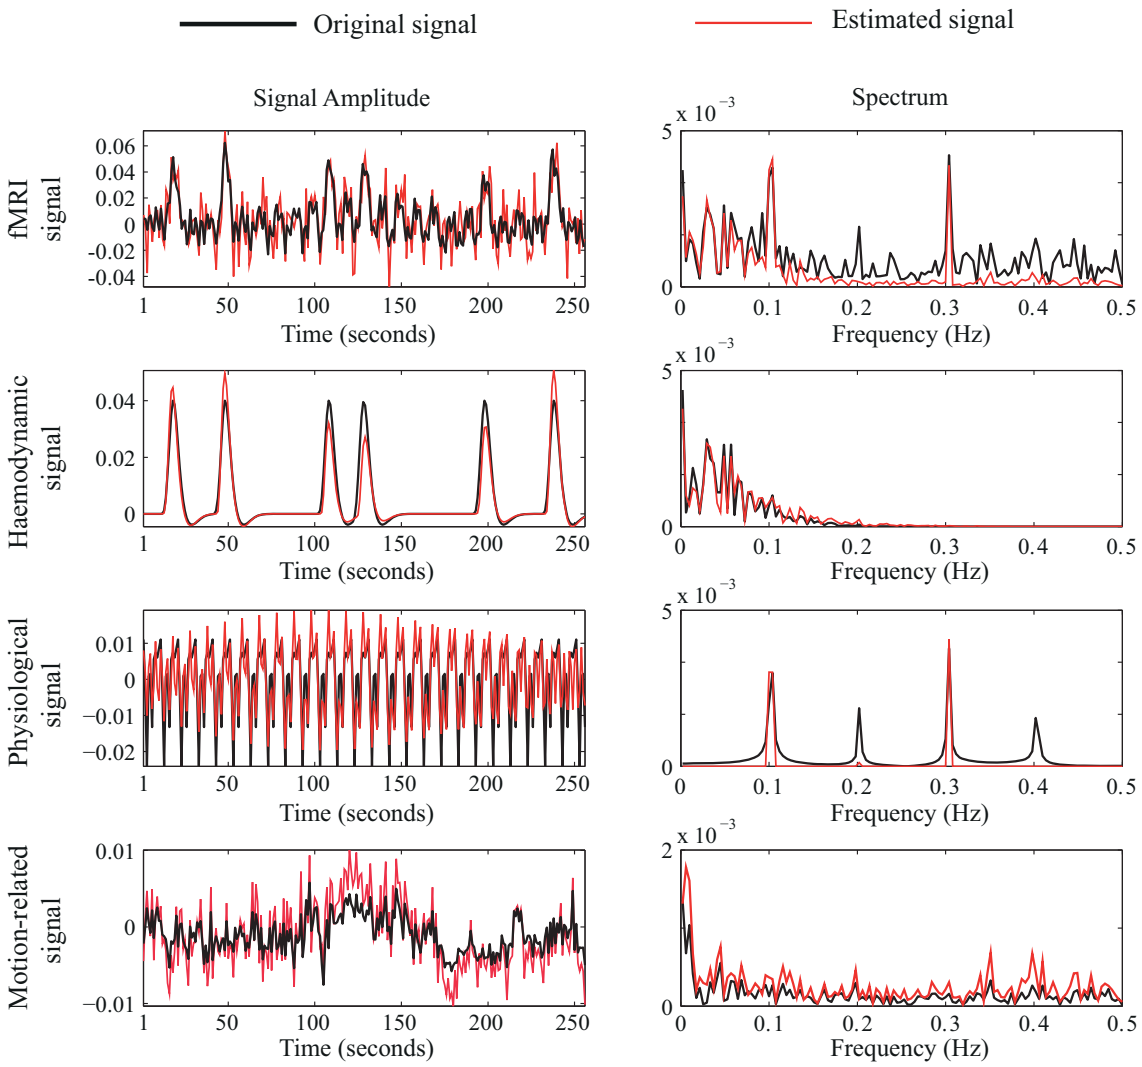

Figure 1. Decomposition of a simulated fMRI time series into the neuronal-related haemodynamic (6 BOLD events of duration 4 s), physiological and motion-related signals.

on a cross in the center of the screen. The fMRI data comprised $N = 140$ T2*-weighted gradient echo images acquired with an echo planar imaging (EPI) trajectory (TR = 2 s, echo time (TE) = 30 ms, flip angle = 85°, voxel size = 3.25x3.25x3.5 mm^3). As for preprocessing, fMRI data were first corrected for head motion by realigning all images to the first image, and then smoothed spatially with a 3D isotropic Gaussian filter with full-width at half-maximum of 5 mm, and both steps were performed with the SPM8 software (Wellcome Trust Centre for Neuroimaging, UCL, UK).[27] A T1-weighted MPRAGE anatomical image was also acquired, and registered to the first fMRI image using AFNI (NIMH/NIH)[45] to help with the anatomical localization of the functional activations.

Figure 3 illustrates the operation of the algorithm in a voxel located in the primary visual area exhibiting a clear response to the visual events. To select this voxel, we performed a conventional GLM analysis using SPM8 with 27 regressors: two regressors for each individual event (20 regressors) resulting from the convolution of delta functions at the onset of the events with the canonical haemodynamic response function and its first temporal derivative,[3] the 6 translation and rotation time series estimated during realignment, and a constant term to model the signal mean. The activation map shown on the right reveal the common response to the

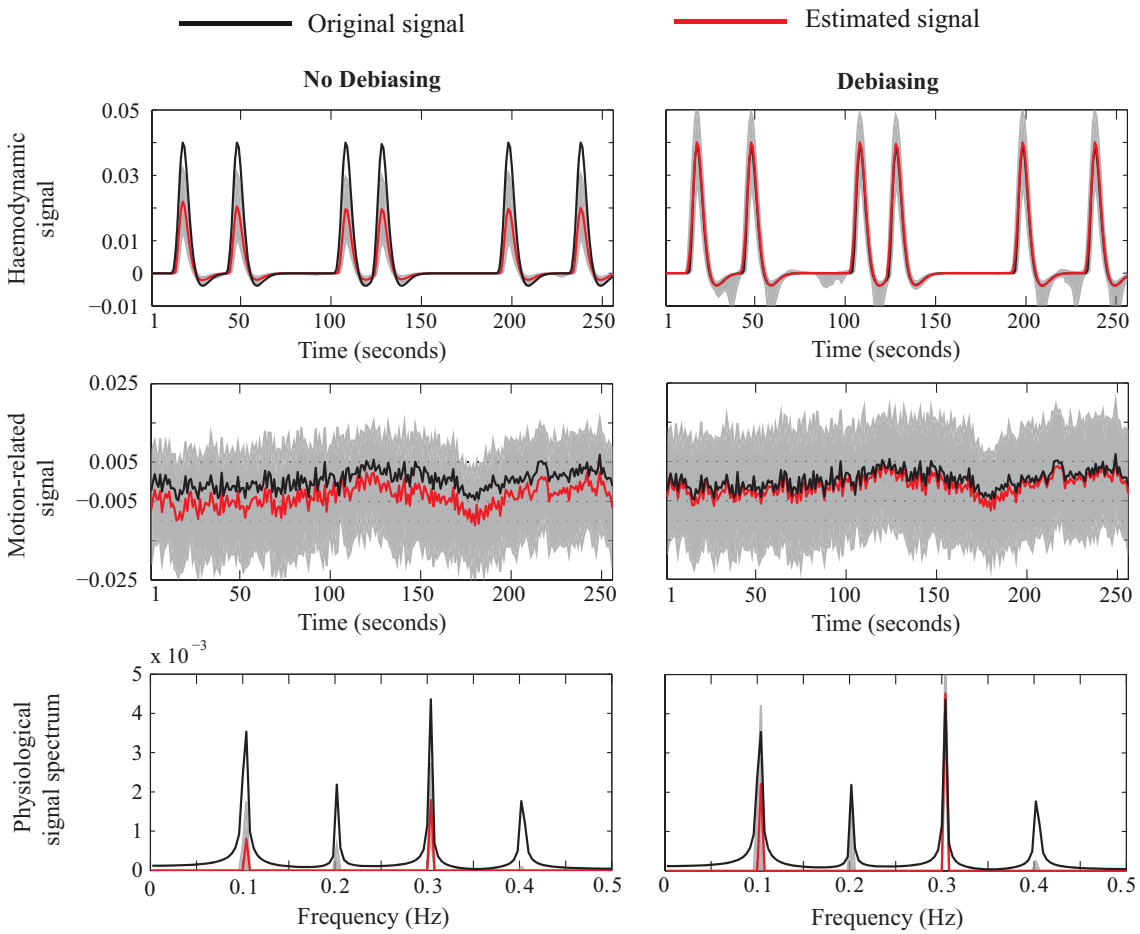

Figure 2. Averaged estimates with (left) and without debiasing (right) of the neuronal-related haemodynamic, physiological and motion-related signals over 1000 realisations with varying physiological and thermal noise. Grey shadow areas indicate the confidence interval of the estimates, defined by the 10% and 90% percentiles.

10 events (F-test) thresholded at uncorrected $p \leq 0.001$ (F-statistic = 2.51) at one slice covering the occipital cortex, overlaid over the corresponding T1-weighted image slices. The F-statistic of the selected voxel is $F_{20,128}$ = 9.276. The maps obtained with our approach also display this voxel as active. As it can be seen in the figure, the BOLD responses to the 10 visual events (whose onsets are marked with arrows) are clearly deconvolved in the haemodynamic signal. In addition, the algorithm is able to separate the part of the signal due to sinusoidal fluctuations, motion-related effects and scanner drifts. For the sake of simplicity, the signals modeling scanner drifts and motion-related changes are combined into a single timecourse.

5. CONCLUSIONS

We have presented an extension of paradigm free mapping which aims to decompose the fMRI data into its sources based on an overcomplete representation the fMRI voxel time series with multiple morphological dictionaries. Unlike conventional fMRI data processing methods, our algorithm does not need the knowledge of the stimulus events (i.e. it is a form of paradigm free mapping). Without loss of generalization, our model considered four sources: scanner drifts, motion-related effects, physiological sinusoidal fluctuations and the neuronal-related haemodynamic signal, that were estimated by means of regularized least-squares estimators with L_1-norm penalties for the neuronal-related haemodynamic and Fourier coefficients modeling the physiological fluctuations. Similar to morphological component analysis, an adapted block coordinate relaxation algorithm

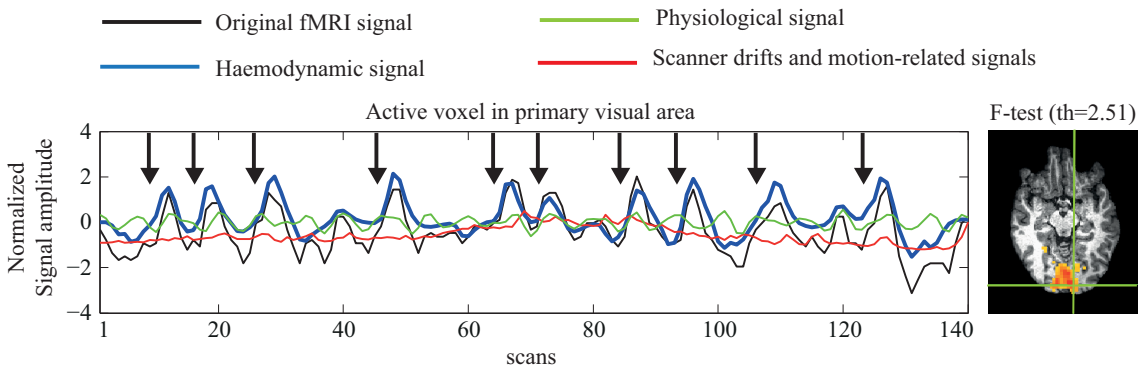

Figure 3. Decomposition of the fMRI time series of a voxel located in the primary visual area into the neuronal-related haemodynamic, physiological and motion-related signals. This voxel showed a significant response to the 10 visual stimuli whose onsets are marked with arrows, which are clearly deconvolved in the estimate of the haemodynamic signal.

was used to solve this problem and the criterion to assess convergence of the algorithm was based on a statistical test on the whiteness of the residuals.

We envisage that this way of decomposing fMRI data (based on a-priori expertise regarding the shape of the fMRI signal sources) will become particularly useful for ultrahigh-field MR systems where the large contrast and signal-to-noise ratios makes feasible the deconvolution of the components of the fMRI signal with overcomplete representations and using sparsity constraints. This technique provides a useful new method for exploratory analysis of fMRI data situated between conventional general-linear-model fitting (with a fixed model) and relatively unconstrained approaches such as independent component analysis.

ACKNOWLEDGMENTS

This work is supported in part by the Swiss National Science Foundation (grant PP00P2-123438, DVDV), and in part by the Faculty of Medicine of the University of Geneva (MIMOSA, CCG and DVDV), and in part by the Center for Biomedical Imaging (CIBM) of the Geneva-Lausanne Universities and Hospitals and the EPFL.

REFERENCES

1. S. Ogawa, T. M. Lee, A. R. Kay, and D. Tank, "Brain magnetic resonance imaging with constrant dependent on blood oxygenation," *Proc. Nat. Acad. Sci. USA* **87**, pp. 5951–5955, 1990.
2. G. E. Sarty, *Computing brain activity maps from fMRI time-series images*, Cambridge University Press, 2007.
3. K. J. Friston, P. Fletcher, O. Josephs, A. Holmes, M. D. Rugg, and R. Turner, "Event-related fMRI: Characterizing differential responses," *Neuroimage* **7**, pp. 30–40, 1998.
4. P. Ciuciu, J. B. Poline, G. Marrelec, J. Idier, C. Pallier, and H. Benali, "Unsupervised robust non-parametric estimation of the hemodynamic response function for any fMRI experiment," *IEEE Trans. Med. Imag.* **22**, pp. 1235–1251, 2003.
5. S. Makni, P. Ciuciu, J. Idier, and J. B. Poline, "Joint detection-estimation of brain activity in functional MRI: A multichannel deconvolution solution," *IEEE Trans. Signal Processing* **53**, pp. 3488–3502, 2005.
6. S. Faisan, L. Thoroval, J. P. Armspach, and F. Heitz, "Hidden Markov Multiple Event Sequence Models: A paradigm for the spatio-temporal analysis of fMRI data," *Medical Image Analysis* **11**, pp. 1–20, 2007.
7. R. A. Hutchinson, R. S. Niculescu, T. A. Keller, I. Rustandi, and T. M. Mitchell, "Modeling fMRI data generated by overlapping cognitive processes with unknown onsets using Hidden Process Models," *Neuroimage* **46**, pp. 87–104, 2009.
8. J. A. D. Aston and C. Kirch, "Estimation of the distribution of change-points with appliction to fMRI data." Working Paper. University of Warwick, 2011.

9. M. A. Lindquist, C. Waugh, and T. D. Wager, "Modeling state-related fMRI activity with change-point theory," *Neuroimage* **35**, pp. 1125–1141, 2007.
10. L. F. Robinson, T. D. Wager, and M. A. Lindquist, "Change point estimation in multi-subject fMRI studies," *Neuroimage* **49**, pp. 1581–1592, 2010.
11. R. B. Buxton, E. C. Wong, and L. R. Frank, "Dynamics of blood flow and oxygenation changes during brain activation: The Balloon model," *Magn. Reson. Med.* **39**, pp. 855–864, 1998.
12. K. J. Friston, A. Mechelli, R. T. R, and C. J. Price, "Nonlinear responses in fMRI: The balloon model, volterra kernels, and other hemodynamics," *Neuroimage* **12**, pp. 466–477, 2000.
13. 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* **21**, pp. 547–567, 2004.
14. K. J. Friston, N. Trujillo-Barreto, and J. Daunizeau, "DEM: A variational treatment of dynamic systems," *Neuroimage* **41**, pp. 849–885, 2008.
15. K. J. Friston, K. E. Stephan, and J. Daunizeau, "Generalized filtering," *Math. Prob. Eng.* **621670**, 2010.
16. M. Havlicek, K. J. F. J. Jiri, M. Brazdil, and V. D. Calhoun, "Dynamic modeling of neuronal responses in fMRI using cubature Kalman filtering," *Neuroimage* **56**, pp. 2109–2128, 2011.
17. I. Khalidov, M. J. Fadili, F. Lazeyras, D. V. D. Ville, and M. Unser, "Activelets: Wavelets for sparse representation of hemodynamic responses," *Signal Processing* **91**, pp. 2810–2821, 2011.
18. 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.* **9**, pp. 1400–1418, 2011.
19. 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.
20. 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* **19**, pp. 200–207, 2003.
21. E. Candes and T. Tao, "The Dantzig Selector: Statistical estimation when p is much larger than n," *Ann. Statist.* **35**, pp. 2313–2351, 2007.
22. J. L. Starck, M. Elad, and D. l. Donoho, "Image decomposition via the combination of sparse representation and a variational approach," *IEEE Trans. Image Processing* **14**, pp. 1570–1582, 2005.
23. S. Sardy, A. Bruce, and P. Tseng, "Block coordinate relaxation methods for nonparametric wavelet denoising," *J. Comp. Graph. Stat.* **9**, pp. 361–379, 2000.
24. J. Fadili and E. T. Bullmore, "Penalized partially linear models using sparse representations with application to fMRI time series," *IEEE Trans. Signal Processing* **53**, pp. 3436–3448, 2005.
25. F. G. Meyer, "Wavelet-based estimation of a semiparametric generalized linear model of fMRI time-series," *IEEE Trans. Med. Imag.* **22**, pp. 315–322, 2003.
26. H. Luo and S. Puthusserypady, "Analysis of fMRI data with drift: Modified general linear model and Bayesian estimator," *IEEE Trans. Biomed. Eng.* **55**, pp. 1504–1511, 2008.
27. K. J. Friston, J. T. Ashburner, S. J. Kiebel, T. E. Nichols, and W. D. Penny, *Statistical Parametric Mapping: The analysis of functional brain images*, Elsevier, 2007.
28. T. E. Lund, K. H. Madsen, K. Sidaros, W. L. Luo, and T. E. Nichols, "Non-white noise in fMRI: Does modelling have an impact?," *Neuroimage* **29**, pp. 54–66, 2006.
29. K. N. Kay, S. V. David, R. J. Prenger, K. A. Hansen, and J. L. Gallant, "Modeling low-frequency fluctuations and haemodynamic response timecourse in event-related fMRI," *Hum. Brain Mapp.* **29**, pp. 142–156, 2008.
30. K. J. Friston, S. Williams, R. Howard, R. S. J. Frackowiak, and R. Turner, "Movement-related effects in fMRI time-series," *Magn. Reson. Med.* **35**, pp. 346–355, 1996.
31. G. H. Glover, T. Q. Li, and D. Ress, "Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR," *Magn. Reson. Med.* **44**, pp. 162–167, 2000.
32. R. M. Birn, J. B. Diamond, M. A. Smith, and P. A. Bandettini, "Separating respiratory-variation-related fluctuations from neuronal-activity-related fluctuations in fMRI," *Neuroimage* **31**, pp. 1536–1548, 2006.

33. K. Shmueli, P. van Gelderen, J. A. de Zwart, S. G. Horovitz, M. Fukunaga, J. M. Jansma, and J. H. Duyn, "Low frequency fluctuations in the cardiac rate as a source of variance in the resting-state fMRI BOLD signal," *Neuroimage* **38**, pp. 306–320, 2007.
34. J. Bobin, J. L. Starck, and J. Fadili, "Morphological component analysis: An adaptive thresholding strategy," *IEEE Trans. Image Processing* **16**, pp. 2675–2680, 2007.
35. S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Sci. Comput.* **20**, pp. 33–61, 1998.
36. M. S. Asif and J. Romberg, "Dantzig selector homotopy with dynamic measurements," in *Proc. SPIE Computational Imaging VII*, **7246**, 2009.
37. D. L. Donoho and I. M. Johnstone, "Ideal spatial adaptation by wavelet shrinkage," *Biometrika* **81**, pp. 425–455, 1994.
38. I. Daubechies, "Orthonormal bases of compactly supported wavelets," *Comm. Pure Appl. Math.* **41**, pp. 909–996, 1988.
39. P. J. Diggle, *Time Series. A biostatistical introduction*, Oxford University Press, 1990.
40. W. L. Luo and T. E. Nichols, "Diagnosis & exploration of massively univariate fMRI models." Technical Report, Dept. of Biostatistics, University of Michigan, 2002.
41. G. Krüger and G. H. Glover, "Physiological noise in oxygenation-sensitive magnetic resonance imaging," *Magn. Reson. Med* **46**, pp. 631–637, 2001.
42. C. Triantafyllou, R. D. Hoge, G. Krueger, D. J. Wiggins, A. Potthast, G. C. Wiggins, and L. L. Wald, "Comparison of physiological noise at 1.5T, 3T and 7T and optimization of fMRI acquisition parameters," *Neuroimage* **26**, pp. 243–250, 2005.
43. M. Bianciardi, M. Fukunaga, P. van Gelderen, S. G. Horovitz, J. A. de Zwart, K. Shmueli, and J. H. Duyn, "Sources of functional magnetic resonance imaging signal fluctuations in the human brain at rest: A 7T study," *Magn. Reson. Imaging* **27**, pp. 1019–1029, 2009.
44. W. van der Zwaag, S. Francis, K. Head, A. Peters, P. Gowland, P. Morris, and R. Bowtell, "fMRI at 1.5, 3 and 7T: Characterising BOLD signal changes," *Neuroimage* **47**, pp. 1425–1434, 2009.
45. R. W. Cox, "AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages," *Comput. Biomed. Res.* **29**, pp. 162–173, 1996.