14th European Signal Processing Conference (EUSIPCO 2006), Florence, Italy, September 4-8, 2006, copyright by EURASIP

EXTRACTING THE HAEMODYNAMIC RESPONSE FUNCTION FROM FMRI TIME SERIES USING FOURIER-WAVELET REGULARISED DECONVOLUTION WITH ORTHOGONAL SPLINE WAVELETS

Alle Meije Wink[1] *and Jos B. T. M. Roerdink*[2]

[1] Brain Mapping Unit, Department of Psychiatry, University of Cambridge, United Kingdom
Tel: +44-1223-764676, fax: +44-1223-336581
[2] Institute for Mathematics and Computing Science / BCN Neuroimaging Center
University of Groningen, The Netherlands
Tel: +31-50-3633931, fax: +31-50-3633800
E-mail: amw71@cam.ac.uk, j.b.t.m.roerdink@rug.nl

ABSTRACT

We describe a method to extract the haemodynamic response function (HRF) from functional magnetic resonance imaging (fMRI) time series based on Fourier-wavelet regularised deconvolution (ForWaRD), and introduce a simple model for the HRF. The HRF extraction algorithm extends the ForWaRD algorithm by introducing a more efficient computation in the case of very long wavelet filters. We compute shift-invariant discrete wavelet transforms (SI-DWT) in the frequency domain, and apply ForWaRD using orthogonal spline wavelets. Extraction and modelling of subject-specific HRFs is demonstrated, as well as the use of these HRFs in a subsequent brain activation analysis. Temporal responses are predicted by using the extracted HRF coefficients. The resulting activation maps show the effectiveness of the proposed method.

1. INTRODUCTION

Functional magnetic resonance imaging (fMRI) is a versatile tool for functional neuroimaging. Regional brain activation induces changes in blood oxygenation, generating a blood oxygenation level dependent (BOLD) contrast in MR images [16]. An important tool for fMRI analysis is statistical hypothesis testing, where the fMRI signal is predicted using the stimulus pattern and a response model. The general linear model (GLM) assumes that the BOLD response to each stimulus is linear and time-invariant (LTI) and that the total response is the sum of all modelled responses. Statistical parametric mapping [6] uses a model of the noise (Gaussian), and hypothesis tests are based on the parameters of the model. Within the GLM framework and given the stimulus pattern, a response can be computed with a fixed impulse response function, by convolving the time pattern of stimuli with the impulse response.

This paper presents a method to extract the haemodynamic response function (HRF) from fMRI data. The method is based on Fourier-wavelet regularised deconvolution, ForWaRD [15], using orthogonal spline wavelets [19]. Extraction of the HRF is a difficult task. Most methods in use today are either based on selective averaging, which requires long inter-stimulus times to prevent overlapping responses [1, 3], or they combine HRF extraction and modelling, restricting the HRF to predefined functions [4, 9, 14]. Nonlinear stimulus-response relations have been found [7, 20], *e.g.*, for varying stimulus durations and amplitudes. The experiments presented here have fixed stimulus durations and amplitudes, and the GLM is assumed to be valid.

The advantage of deconvolution in the frequency domain is that overlapping responses can be separated [10], which is not possible with selective averaging. ForWaRD combines frequency-domain deconvolution with regularisation in the frequency and wavelet domains, to reduce noise as much as possible without introducing large errors. An orthogonal spline wavelet basis is used; spline wavelets have many favourable properties, such as smoothness, optimal approximation properties, and good localisation in both the time and frequency domains [18].

Prior to the HRF extraction, low-frequency trends are removed from the time signals using a standard wavelet-based technique [13]. The input of the extraction routine is an fMRI time series and a file containing the stimulus times, and the output is given as a time series of image volumes containing the HRF at each voxel location. Compared to other HRF extraction methods, the method requires only few assumptions. The stimulus-response relation is assumed to be linear and time invariant (LTI), and the signal is assumed to be separable from noise in the Fourier and/or wavelet representation. Our ForWaRD algorithm uses a novel frequency-domain implementation of the shift-invariant discrete wavelet transform (SI-DWT) that we have recently developed [23]. It is efficient to compute the wavelet transform in the frequency domain if the basis functions do not have compact support [21]. This is the case for orthogonal spline wavelets, which have exponential decay.

To obtain a precise temporal representation, we use a very simple model for the HRF, based on a damped oscillator. Sampling the HRF only at the start of each scan yields a coarse temporal resolution, while a model function can be computed at arbitrary time points. The HRF is found by fitting the model function to the extracted HRF coefficients. The model is easily combined with a standard time series analysis and is more flexible than the canonical HRF as used in the SPM program [6], which is the the difference of two specific gamma density functions.

2. THEORETICAL BACKGROUND

2.1 SPM

Statistical parametric mapping (SPM) is the standard fMRI analysis tool. It uses the general linear model (GLM): the response to a stimulus pattern is modelled as the output of a linear, time invariant (LTI) system, see Fig. 1(a). It assumes the temporal noise to be independent, identically distributed and Gaussian. SPM consists of the following steps: (*i*) estimate the parameters of the noise, (*ii*) compute a statistic at every voxel location, (*iii*) threshold the statistic values using the noise parameters and a multiple testing correction method. Assuming an LTI stimulus-response relation, an fMRI data set of T time samples in N voxels is modelled as:

$$Y_{[T\times N]} = X_{[T\times M]}\beta_{[M\times N]} + e_{[T\times N]}. \qquad (1)$$

Y represents the fMRI data, X is the *design matrix* of M explanatory variables (modelled effects), β contains the weights of each of these variables in each voxel, and e contains the residuals, *i.e.*, the unmodelled part of the signals. Each column of X contains the modelled response to the stimuli of one type. An LTI response to one type of stimulus is given by a convolution of the time pattern of the stimuli and the appropriate HRF. A good HRF model is critical to the success of the estimation based on (1), because an inaccurate model may lead to a non-Gaussian distribution of the values in e.

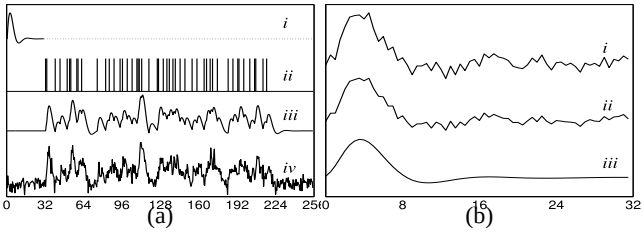

Figure 1: (a) Inputs and output of the GLM: (*i*) the impulse response function, (*ii*) the stimulus pattern, (*iii*) the total response without noise, (*iv*) the noisy response. (b) Various stages of the ForWaRD algorithm: (*i*) extracted HRF after frequency domain inversion, (*ii*) after frequency domain shrinkage, (*iii*) after subsequent wavelet domain Wiener filtering.

2.2 Modelling the HRF

A simple way to model a delayed response like the BOLD signal (the delivery of extra oxygen follows the increased oxygen consumption) is by describing it as a damped oscillator. The function

$$f_{H,D,P,L}(t) = \begin{cases} H\ \sin(\frac{t-L}{P})\ e^{\frac{-t+L}{D}}, & \text{if } t > L \\ 0, & \text{otherwise} \end{cases} \qquad (2)$$

models a damped oscillator with H(eight), L(ag), P(eriod), and D(ilation) parameters. Function (2) can be used to model the BOLD response by fitting it to the extracted HRF coefficients, and using the resulting function as the HRF. It is demonstrated in section 4 that if there are enough time points to model the complete undershoot of the HRF as well as the initial peak, a sum of two such functions can also be used to describe the HRF, yielding a more accurate model of the undershoot.

2.3 ForWaRD

Using the LTI model described in section 2.1, an HRF can be extracted from an fMRI time series by deconvolving the measured time signals with the stimulus pattern. In the frequency domain, deconvolution can be done via pointwise division; this is called Fourier inversion. The great advantage of Fourier inversion is the fact that it can separate overlapping responses. However, noise is amplified at frequencies where the signal is small, introducing instablitity: small changes in the inputs induce large changes in the output. Regularisation suppresses the destabilising effects. If the destabilising factor is the noise, regularisation is tantamount to denoising.

A common regularisation technique for frequency domain deconvolution is shrinking the frequency coefficients after inversion. The signal of interest is usually smooth (low-frequency) and noise is usually erratic (high-frequency). Two familiar shrinkage methods are Wiener shrinkage and Tikhonov shrinkage [15]. Non-smooth parts of signals (such as steep edges) are not efficiently represented in the frequency domain, because they contain much high-frequency energy. As a result, noise at those frequencies is not shrunk. Using more shrinkage to remove noise introduces artifacts, such as ringing.

The ForWaRD deconvolution scheme performs noise regularisation via scalar shrinkage in the Fourier and wavelet domains. Fourier shrinkage exploits the economical representation of the noise in the Fourier domain, whereas wavelet shrinkage exploits the fact that piecewise smooth signals and images with singularities, such as step edges, have a sparse representation in the wavelet domain. The main steps of this approach are depicted in Fig. 2, see also Fig. 1(b). Wiener shrinkage reduces the magnitude of wavelet coefficients at indices where the true signal is weak, and preserves those coefficients where the true signal is strong. The true signal is unknown, so ForWaRD uses two wavelet transforms of a signal:

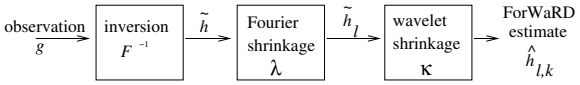

Figure 2: ForWaRD HRF extraction scheme. Fourier shrinkage (determined by shrinkage factor λ) is applied to partially attenuate noise amplified during the inversion step. Subsequent wavelet shrinkage (determined by shrinkage factor κ) effectively attenuates the residual noise.

one transform with basis functions (ϕ_1, ψ_1) to estimate the true signal by thresholding detail coefficients, and another transform with basis functions (ϕ_2, ψ_2), whose detail coefficients are shrunk. ForWaRD uses the SI-DWT to ensure shift-invariance.

3. COMPUTING THE SI-DWT IN THE FREQUENCY DOMAIN

ForWaRD requires the SI-DWT, which was implemented in the time domain [12]. Spline wavelets [19] can be computed most efficiently in the frequency domain.

Starting from the definition of the SI-DWT of Shensa [17] (resulting in the *à trous* algorithm), we recently derived an efficient implementation of the SI-DWT in the frequency domain [23]. Let h and g define an orthogonal wavelet basis, with *dual* filters defined as $\widetilde{h}(n) = \overline{h(-n)}$ and $\widetilde{g}(n) = \overline{g(-n)}$, with $\bar{x}$ denoting the complex conjugate of x. Then the forward SI-DWT in the Fourier domain is given by:

$$C^{j+1} = H^Q \bullet C^j, \quad D^{j+1} = G^Q \bullet C^j, \qquad (3)$$

for $j = 0, 1, \cdots L-1$, $Q = 2^j$, where L denotes the number of levels in the decomposition. Input is a vector C^0; output are vectors $D^1, D^2, \cdots, D^L, C^L$. Here $x \bullet y$ denotes pointwise multiplication of vectors x and y, and C^j, D^j, H^Q and G^Q denote the DFT vectors of C^j, D^j, $\uparrow_Q h$ and $\uparrow_Q g$, respectively. Here, $\uparrow_Q x$ denotes upsampling vector x with a factor Q, and the length-N vectors H^Q and G^Q contain the Fourier coefficients of the upsampled filters. In particular, if $j = 0$, *i.e.*, $Q = 1$, H^1 equals the DFT vector H of h.

It can be shown that [23],

$$H^Q = \overbrace{[(\downarrow_{2^j} H) \quad (\downarrow_{2^j} H) \cdots (\downarrow_{2^j} H)]}^{Q \text{ times}}$$

So, in iteration j of the decomposition, the DFT vector H^Q is obtained by downsampling the DFT vector H by a factor of $Q = 2^j$, and then Q times repeating this reduced vector of length N/Q to again get a filter of length N. Alternatively, two copies of the even-numbered samples of the filter values in the previous iteration $j-1$ (*i.e.*, $H^{Q/2}$) are concatenated to obtain H^Q. The case of G^Q is analogous.

Reconstruction in the Fourier domain is given by

$$C^{j-1} = H'^Q \bullet C^j + G'^Q \bullet D^j, \qquad (4)$$

where H'^Q and G'^Q are obtained in the same way from h' and g' as described above for H^Q.

3.1 Computation times: spline wavelets

The algorithm for computing the SI-DWT in the frequency domain is more efficient than a time-domain computation when the wavelet basis functions have wide support. We compared the computation times of the time and frequency-domain versions of the SI-DWT, by computing the three-level SI-DWT of signals varying in length, using the symmetric orthogonal cubic spline wavelet basis [11]. For each length, 100 signals were transformed and reconstructed, using the time-domain implementation, a naive frequency-domain implementation, and our fast implementation of the *à trous* algorithm.

Figure 3 shows the results: the naive frequency-domain version is faster than the time-domain version for signals of more than 64 points. The optimised frequency-domain version is always fastest. Therefore, the frequency-domain SI-DWT is preferred for long filters (like orthogonal spline wavelets), and is used in this paper.

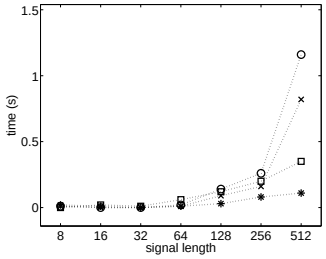

Figure 3: Computation times of the time-domain SI-DWT and SI-IDWT and the frequency-domain SI-DWT and SI-IDWT, respectively, with symmetric orthogonal cubic spline wavelet basis functions. Time domain: $\diamond$: SI-DWT, $+$: SI-IDWT. Frequency domain (naive) $\times$: SI-DWT, $\circ$: SI-IDWT, (fast) $\square$: SI-DWT, $*$: SI-IDWT.

3.2 ForWaRD using spline wavelets

Our new version of ForWaRD with efficient frequency-domain SI-DWT and spline wavelets was used in the HRF extraction program. We use orthonormal splines to preserve the signals' energy during the transform. Unser et al. have proposed fractional spline wavelets because of their favourable properties [2, 19]. Their implementation of fractional splines was used to generate the wavelet basis functions, cf. Fig. 4. In this paper, we used causal splines with degree $\alpha = 4$ for (ϕ_1, ψ_1) and anticausal splines with degree $\alpha = 3$ for (ϕ_2, ψ_2).

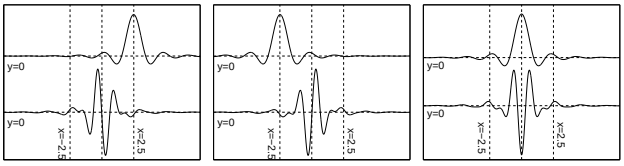

(a) (b) (c)

Figure 4: Orthonormal quartic (α=4) spline wavelet basis functions: (a) causal, (b) anticausal, (c) symmetric. Top: scaling function ϕ, bottom: wavelet ψ.

3.3 HRF extraction based on ForWaRD

The total routine for extracting an HRF from an fMRI time series, using only the stimulus times, consists of the following steps:

1: load the time series g and the stimulus pattern f;
2: subtract the time series mean;
3: remove low-frequency trends;
this is done with a wavelet-based method described in the literature [13];
4: apply ForWaRD to g, estimate the HRF h_κ to the stimuli with pattern f.

Time signals at different voxel locations can be processed independently, so that it is possible to process multiple time signals at once. The number of time signals processed simultaneously can be adjusted according to the amount of available computer memory.

This routine was implemented in MatLab (The Mathworks, USA) and was used in combination with the SPM program [6] to analyse fMRI data.

4. EVENT-RELATED FMRI EXPERIMENTS

The HRF extraction routine was used in the analysis of two event-related fMRI experiments of one subject, measured on different days. The subject had to make a fist upon presentation of a visual stimulus, and then immediately relax. Stimuli were presented on a white screen inside the MRI scanner. A white disc was shown as the default, a red disc was the cue to make a fist. One experiment was performed with a fixed inter-stimulus interval (ISI) and one with a randomised ISI. Realignment, normalisation, and statistical analysis were done with the SPM program. Denoising was done with a wavelet-based technique [22]. We computed HRFs for the whole brain and in a region of interest, respectively, which were then used in covariance analyses to test for activation.

4.1 Fixed-ISI Experiment

The fixed-ISI data set consisted of 156 volumes of $64\times64\times46$ voxels with size $3.5 \times 3.5 \times 3.5$ mm^3. Cues were given every 24 s (8 scans $\times$ 3 s) starting at scan 2. HRFs were extracted by our method, and also by selective averaging [5], a simple and robust extraction method that works for long ISIs. A first statistical analysis was done to detect activation synchronous to the stimuli. We used a design matrix with a set of 6 Fourier basis functions, modulated by a Hanning window, in the time interval of 8 scans after each stimulus, so as not to impose shape assumptions on the HRF. An SPM$\{F\}$ resulting from an F-test was computed, using false discovery rate (FDR) control [8] with $q = 0.05$ for multiple hypothesis testing. With both ForWaRD (using 128 of the 156 scans) and selective averaging, we computed a whole-volume HRF and a regional HRF in a $7 \times 7 \times 7$-voxel region with high activity (see the region indicated by a '$<$' in Fig. 5a). The post-stimulus volumes were multiplied by the F-values in the map, after thresholding with an FDR-parameter $q = 0.0001$, and averaged over the volume/region. Figures 6a-b show the HRFs. Selective averaging almost returns to baseline within the ISI, whereas ForWaRD remains below baseline. This may be because the real HRF does not return to baseline within the measured interval, so in the LTI model the response decreases at every next stimulus. This results in an HRF with a lower baseline. Selective averaging forces each response to begin at baseline, even when the response does not return to baseline within the sampled time interval. In general, ForWaRD-based methods are expected to work better with a varying ISI, because a fixed ISI leads to a badly conditioned Fourier inversion. The next experiment therefore used a randomised ISI.

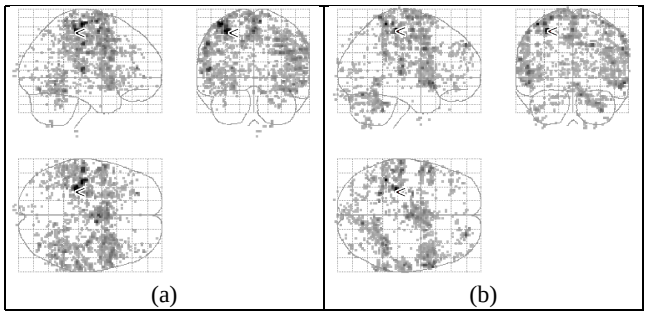

(a) (b)

Figure 5: SPM $\{F\}$ of the fixed-ISI experiment (a), SPM $\{F\}$ of the random-ISI experiment (b), thresholded using FDR correction with q=0.05

4.2 Random-ISI Experiment

In a second fMRI experiment, stimulus times for this experiment were random and the length of the random-ISI data set was 256 scans. The other parameters were unchanged. The stimulus signal was created by thresholding a vector of uniformly distributed

random numbers. The number of stimuli was 39, so the average ISI was between 6 and 7 scans. Post-stimulus image volumes were produced by ForWaRD. Thanks to the spreading of the stimulus, the HRF could be sampled on a much larger time interval. Due to the response overlap, neither selective averaging nor the Fourier basis set could be used. The design matrix X was made by convolving the stimulus signal with the 'canonical' HRF from the SPM´99 program [6] and its time and dilation derivatives. HRFs were made from the SPM$\{F\}$ (see Fig. 5b) and the post-stimulus volumes, see Fig. 6c. The regional HRF corresponds most to the previously extracted HRFs. Both HRFs return to baseline within the post-stimulus interval. This does indicate that with random inter-stimulus times, where selective averaging cannot be used, ForWaRD-based HRF extraction performs *better* than with fixed inter-stimulus times.

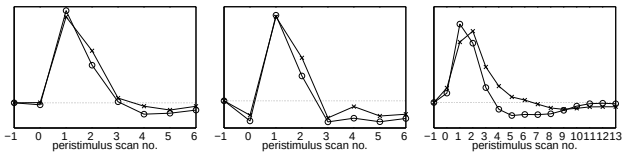

(a) (b) (c)

Figure 6: HRFs extracted from the fixed-ISI data set by selective averaging (a) and by ForWaRD (b), and from the random-ISI time series by ForWaRD (c). ×: whole-volume, ○, region-specific.

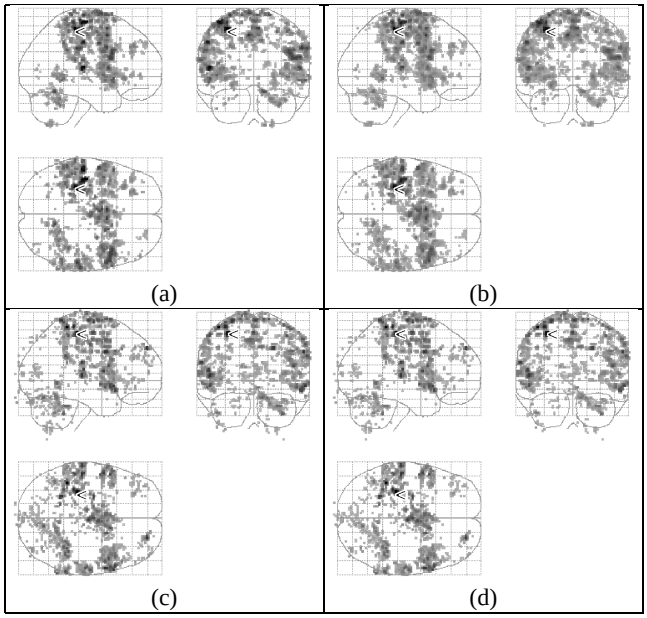

(a) (b) (c) (d)

Figure 8: SPM$\{T\}$s of the activation found by using the modelled HRFs: (a) fixed-ISI data, random-ISI whole-volume HRF, (b) fixed-ISI data, random-ISI regional HRF, (c) random-ISI data, fixed-ISI whole-volume HRF, (d) random-ISI data, fixed-ISI regional HRF.

4.3 Using the extracted HRFs in activation tests

A covariance test was done on the random-ISI data using the fixed-ISI HRF coefficients in the model, and *vice versa*. HRFs extracted from one data set cannot be used for covariance tests on that same data set: a model must be specified *a priori*, and inferences cannot be made from models that are determined by the data. We modelled the fixed-ISI HRFs by fitting one model function (2) to the HRFs extracted from the fixed-ISI data. Two such functions (one

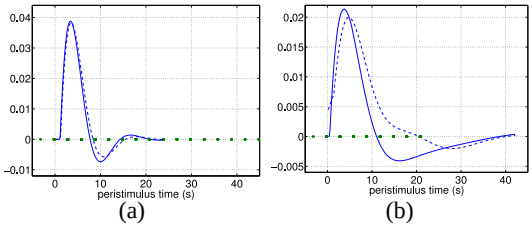

(a) (b)

Figure 7: The modelled HRFS for the covariance test, with the coefficients from the fixed-ISI experiment (a) and the random-ISI experiment (b). Solid lines: regionally determined HRFs, dashed lines: whole-volume HRFs.

to model the peak and one to model the undershoot) were used for the HRFs from the random-ISI data. The fixed-ISI signals did not have enough coefficients to accurately model the undershoot. The fitted functions were then used to build the design matrices for subsequent covariance tests. The maps in Fig. 8 resulting from a t-test show very similar shapes as those in Fig. 5, but here the detected activations are stronger. This indicates that the model used here captures all variance captured by those methods. The difference between this analysis and the previous is that only one basis function is used here, enabling a covariance test with stronger responses.

Table 1 shows the maximum variance ratio values found in the tests with the modelled HRFs. A high variance ratio indicates that much of the variance in the signal is explained by the model, and that the residual noise in the GLM (see Eq. (1)) is small. It shows that ForWaRD works as well on the random-ISI dataset as it does on the fixed-ISI data set. Its performance is similar to that of averaging on the fixed-ISI dataset. The modelled HRFs, acquired with ForWaRD as well as with selective averaging, outperform the canonical HRF in terms of explained variance. The modelled region-specific HRFs generally perform better than whole-volume HRFs. The maps of detected activation indicate that the modelled HRFs do not only detect activation in the region from which they were extracted, but that they are general enough also to detect activation in other areas.

Table 1: A comparison of maximum variance ratio values. The activation test for the fixed-ISI experiment used the HRFs computed from the random-ISI data, and *vice versa*.

	ForWaRD		selective averaging		
	volume	region	volume	region	HRF_{spm}
fixed-ISI	120	163	–	–	117
random-ISI	103	101	102	104	74

5. CONCLUSION

We have presented an HRF extraction method for fMRI time series based on ForWaRD (Fourier-wavelet regularised deconvolution). The extraction method removes the time series mean, removes low-pass trends with a wavelet-based method, and applies ForWaRD to the resulting signal to extract the HRF. The output of is given as a post-stimulus time series of image volumes, representing the HRF in every voxel.

The existing ForWaRD method has been extended by introducing a novel frequency-domain implementation of the SI-DWT. Timings show that for signals longer than 64 points, the speed gain of the frequency-domain transform is considerable. This enabled us to efficiently use orthogonal spline wavelets. We also presented a model for the HRF that can be used in combination with the extracted coefficients to predict event-related fMRI responses.

In combination with HRF data extracted from fMRI time series, this model yielded an accurate an temporal description of subject-specific or group-specific BOLD responses.

In the event-related fMRI experiments, the ForWaRD-based method was able to extract the HRF from both fixed-ISI fMRI time series and random-ISI fMRI time series. The HRF extracted from the fixed-ISI data by ForWaRD looked more noisy than the HRF extracted by selective averaging, which is most likely due to the response not completely returning to baseline within 24 seconds, as well as to the badly conditioned inversion (deconvolution) problem with a fixed ISI. The ForWaRD-based method worked particularly well in the latter case. This suggests that additional long-ISI studies for HRF extraction are not necessary: they may just be extracted from other studies (with the same stimuli) of the same subject. Despite the fact that the average inter-stimulus distance was smaller in the random-ISI study, the HRF could be sampled in a much longer interval. In the fixed-ISI case, the size of the sampling interval was bounded by the ISI itself. This is another good reason to use random ISIs and deconvolution methods to extract the HRF, rather than fixed ISIs and selective averaging.

The modelled HRFs, using the model function and extracted coefficients, have shown to capture the same amount of variance in one basis function as traditional methods that require multiple basis functions. The SPM $\{F\}$ maps of the traditional models and the SPM $\{T\}$ maps show the same regions. The SPM $\{T\}$ maps resulting from the 'canonical' HRF show much smaller detected regions and the amount of detected activity (indicated by the maximum values for the variance ratio) is much smaller. This suggests that the performance of standard HRF functions may be significantly improved by re-evaluating them for each subject. The generality of the results need not depend on the parameters of the model function, but only on the acceptance of the HRF model.

6. ACKNOWLEDGEMENTS

This research is part of the project "Wavelets and their applications", funded by the Dutch National Science Foundation (NWO), project no. 613.006.570.

References

[1] P. A. Bandettini and R. W. Cox. Event-related fMRI contrast when using constant interstimulus interval: Theory and experiment. *Magnetic Resonance in Medicine*, 43:540–548, 2000.

[2] T. Blu and M. Unser. The fractional spline wavelet transform: Definition and implementation. In *Proc. IEEE: International Conference on Acoustics, Speech, and Signal Processing.*, pages 512–515, 2000. http://bigwww.epfl.ch/blu/fractsplinewavelets.

[3] R. L. Buckner, P. A. Bandettini, K. M. O. Craven, R. L. Savoy, S. E. Petersen, M. E. Raichle, and B. R. Rosen. Detection of cortical activation during averaged single trials of a cognitive task using functional magnetic resonance imaging. *Proceedings of the National Academy of Sciences*, 93:14878–14883, 1996.

[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 Transactions on Medical Imaging*, 22(10):1224– 1234, 2003.

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

[6] K. J. Friston, A. P. Holmes, K. J. Worsley, J. P. Poline, C. D. Frith, and R. S. J. Frackowiak. Statistical parametric maps in functional imaging: A general linear approach. *Human Brain Mapping*, 2:189–210, 1995. http://www.fil.ion.ucl.ac.uk/spm.

[7] K. J. Friston, A. Michelli, R. Turner, and C. J. Price. Nonlinear responses in fMRI: The balloon model, Volterra kernels and other hemodynamics. *NeuroImage*, 12:466–477, 2000.

[8] C. R. Genovese, N. A. Lazar, and T. E. Nichols. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *NeuroImage*, 15:772–786, 2002. http://www.sph.umich.edu/-nichols/FDR.

[9] G. H. Glover. Deconvolution of impulse response in event-related BOLD fMRI. *NeuroImage*, 9:416–429, 1999.

[10] J. C. Hansen. Separation of overlapping waveforms having known temporal distributions. *Journal of Neuroscience Methods*, 9:127–139, 1983.

[11] S. G. Mallat. A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(7):674–693, 1989.

[12] S. G. Mallat. Zero-crossings of a wavelet transform. *IEEE Transactions on Information Theory*, 37(4):1019–1033, 1991.

[13] F. G. Meyer. Wavelet-based estimation of a semiparametric generalized linear model of fMRI time-series. *IEEE Transactions on Medical Imaging*, 22(3):315–322, 2003.

[14] F. M. Miezin, L. Macotta, J. M. Ollinger, S. E. Petersen, and R. L. Buckner. Characterizing the hemodynamic response: Effects of presentation rate, sampling procedure and the possibility of ordering brain activity based on relative timing. *NeuroImage*, 11:735–759, 2000.

[15] R. Neelamani, H. Choi, and R. G. Baraniuk. ForWaRD: Fourier-wavelet regularized deconvolution for ill-conditioned systems. *IEEE Transactions on Signal Processing*, 52(2):418– 433, 2004.

[16] S. Ogawa, T. M. Lee, A. R. Kay, and D. W. Tank. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences*, 87:9868–9872, 1990.

[17] M. J. Shensa. The discrete wavelet transform: wedding the *à trous* and Mallat algorithms. *IEEE Transactions on Signal Processing*, 40(10):2464–2482, 1992.

[18] M. Unser. Splines: A perfect fit for signal and image processing. *IEEE Signal Processing Magazine*, 16(6):22–38, 1999.

[19] M. Unser and T. Blu. Fractional splines and wavelets. *SIAM Review*, 42(1):43–67, 2000.

[20] A. L. Vazquez and D. C. Noll. Nonlinear aspects of the bold response in functional MRI. *NeuroImage*, 7(2):108–118, 1998.

[21] M. A. Westenberg and J. B. T. M. Roerdink. Frequency domain volume rendering by the wavelet X-ray transform. *IEEE Transactions on Image Processing*, 9(7):1249–1261, 2000.

[22] A. M. Wink and J. B. T. M. Roerdink. Denoising functional MR images: a comparison of wavelet denoising and Gaussian smoothing. *IEEE Transactions on Medical Imaging*, 23(3):374–387, 2004.

[23] A. M. Wink and J. B. T. M. Roerdink. Polyphase decompositions and shift-invariant discrete wavelet transforms in the frequency domain. Technical report, Institute for Mathematics and Computing Science, University of Groningen, Nov. 2005. Submitted.