Rapid three-dimensional functional magnetic resonance imaging of the initial negative BOLD response

Martin A. Lindquist [a,*], Cun-Hui Zhang [b], Gary Glover [c], Lawrence Shepp [b]

[a] *Department of Statistics, Columbia University, 1255 Amsterdam Avenue, 10th Floor, MC 4690, New York, NY 10027, United States*
[b] *Department of Statistics, Rutgers University, 501 Hill Center, Busch Campus, Piscataway, NJ 08854, United States*
[c] *Department of Radiology, Stanford University, Lucas Center, MC 5488, Palo Alto, CA 94305, United States*

Abstract

Functional MRI is most commonly used to study the local changes in blood flow that accompanies neuronal activity. In this work we introduce a new approach towards acquiring and analyzing fMRI data that instead provides the potential to study the initial oxygen consumption in the brain that accompanies activation. As the oxygen consumption is closer in timing to the underlying neuronal activity than the subsequent blood flow, this approach promises to provide more precise information about the location and timing of activity. Our approach is based on using a new single shot 3D echo-volumar imaging sequence which samples a small central region of 3D k-space every 100 ms, thereby giving a low spatial resolution snapshot of the brain with extremely high temporal resolution. Explicit and simple rules for implementing the trajectory are provided, together with a straightforward reconstruction algorithm. Using our approach allows us to effectively study the behavior of the brain in the time immediately following activation through the initial negative BOLD response, and we discuss new techniques for detecting the presence of the negative response across the brain. The feasibility and efficiency of the approach is confirmed using data from a visual–motor task and an auditory–motor–visual task. The results of these experiments provide a proof of concept of our methodology, and indicate that rapid imaging of the initial negative BOLD response can serve an important role in studying cognition tasks involving rapid mental processing in more than one region.
© 2007 Elsevier Inc. All rights reserved.

Keywords: Rapid fMRI; Echo-volumar imaging; Single shot; Initial negative dip

1. Introduction

Functional magnetic resonance imaging (fMRI) is most commonly performed using blood oxygenation level-dependent (BOLD) contrast [1–3] to study local changes in deoxyhemoglobin concentration in the brain. Neural activity leads to an increase in both the cerebral metabolic rate for oxygen ($CMRO_2$) and the supply of oxygen via the cerebral blood flow (CBF). The positive BOLD signal is believed to be the result of a transient uncoupling between $CMRO_2$ and the supply increase, causing a reduction in paramagnetic deoxyhemoglobin in the capillaries and ven-

ules. Most analyses of fMRI data use the positive BOLD response to study the underlying neural activity. However, BOLD imaging based on the positive response is limited by the sluggish nature of the underlying evoked hemodynamic response to the neural event (or the hemodynamic response function, HRF), which peaks 5–8 s after that neural activity has peaked. Therefore, all inference regarding where and when activation is taking place is based on oxygenation patterns that are far removed from the underlying event we are interested in studying (i.e., the neural activity).

Several studies have shown that $CMRO_2$ increases more rapidly than CBF in the time immediately following neural activity, giving rise to a decrease in the BOLD signal in the first 1–2 s following activation, called the initial negative BOLD response or the negative dip [4–8]. The amplitude

---

* Corresponding author. Fax: +1 212 851 2164.
E-mail address: martin@stat.columbia.edu (M.A. Lindquist).

of the dip is much smaller than that of the positive BOLD signal, and there is evidence that it is also more localized to areas of neural activity [9–12]. Due in part to these reasons, the negative response has so far not been reliably observed, but if it could these signals would be more appropriate to use for tracking rapid neural events as they occur in a time scale closer to the neural activity. Current implementations of fMRI are capable of visualizing changes in oxygenation patterns with high spatial resolution, but as they are focused on the positive rise they are generally unable to infer dynamic neural activity without confounding. The development of rapid imaging techniques that provide information about the initial negative BOLD signal could potentially allow for the possibility of obtaining more accurate measures of the location and timing of activity in the brain.

In this work we suggest a new approach towards acquiring rapid 3D fMRI data in a single shot, which is sensitive to the initial negative BOLD signal. The standard approach towards three-dimensional imaging is the acquisition of a stack of 2D slices. Using this approach, it usually takes up to 2 s to obtain a full scan of the brain, which does not provide sufficient temporal resolution to study the initial dip. As an alternative to multi-slice sampling, a more effective approach would be to directly sample in 3D k-space. Previous attempts at sampling in this manner includes echo-volumar imaging (EVI) [13–17], which extends echo-planar imaging (EPI) [18] to three dimensions and provides a single shot snapshot image of the brain. Irarrazabal and Nishimura [19] provide a nice overview of a variety of possible multi-shot trajectories, while Sabat et al. [20] and Mir et al. [21] use optimization algorithms to design trajectories that maximize the coverage in k-space for a fixed scan time. The trajectories described in these papers are based on the use of many interleaved shots, which leads to a relatively low temporal resolution, as each shot requires at least 60 ms. In this work we introduce a new type of EVI trajectory that can be used to sample the central cubic portion of 3D k-space in a single shot, which provides a low spatial resolution snapshot of the brain at a temporal resolution of 100 ms.

It should be noted that in this work we are primarily concerned with speeding up the temporal resolution, with less regard for spatial resolution. The trade-off between spatial and temporal resolution is often used to increase the sampling speed required in many applications. In techniques such as the keyhole [22,23], singular value decomposition [24,25] and generalized series reconstruction [26] methods, a priori information consisting of a high-resolution reference image is incorporated with the reduced k-space data in order to maintain the spatial resolution of the dynamic images. Multiple coil techniques such as SMASH [27] and SENSE [28], can also be used to achieve reduction of k-space sampling. With multi-coil techniques, prior knowledge about RF field distributions or the image sensitivity of the coils is utilized for constructing images from under-sampled k-space data. Though in its current implementation our methodology sacrifices spatial resolu-

tion, we hope to reclaim the lost resolution at a future stage using such multi-coil techniques. This issue is discussed in greater detail in Section 4.

In this work we provide explicit and simple rules for designing single shot three-dimensional k-space trajectories, together with a straightforward reconstruction algorithm. An implementation of the method is supplied together with step-by-step guidelines towards the statistical analysis of the resulting data, which is discussed in greater detail in a companion paper [29]. The feasibility and efficiency of the approach is confirmed using data from both a visual–motor task and an auditory–visual–motor task. While the existence of a negative dip in fMRI is still considered controversial [30], the data presented in this paper gives strong evidence for its existence in both experiments. Further the results suggest that the initial negative BOLD response contains important information regarding the timing of activation, information that may be confounded when studying the positive BOLD response.

2. Methods

In this section we introduce a new approach towards the acquisition of fMRI data that allows one to sample the central portion of 3D k-space in a single shot with a temporal resolution of 100 ms. Thereafter, we discuss an efficient algorithm for reconstructing the resulting k-space data. Finally, we deal with issues that arise in the statistical analysis of the resulting high-frequency time series data. These issues include the removal of seasonal components due to heart-rate and respiration, as well as new statistical techniques for detecting a significant negative BOLD response in an fMRI time course.

2.1. Single shot 3D k-space sampling

Our approach towards k-space sampling attempts to sample as large a portion of 3D k-space as possible in a single shot. In order to effectively sample the data, new ways of transversing 3D k-space must be developed. Our goal is to find a trajectory, $k(t)$, that moves through the central portion of 3D k-space and satisfies the necessary machine, time and space-filling constraints. The trajectory is defined as a continuous curve and along this curve, measurements will be made at uniform time intervals (e.g., once every 4 μs) determined by the sampling bandwidth of the scanner. The trajectory needs to satisfy the three constraints outlined in the following subsections.

2.1.1. Machine constraints

Let $g(t)$ represent the value of the gradient at time t, and $s(t)$ the slew rate. They are related to the trajectory as its first and second derivative, respectively, and must satisfy the following constraints:

$$|g(t)| \leqslant G_0, \qquad g(t) = \frac{1}{\gamma}\dot{k}(t) \qquad (2.1)$$

and

$$|s(t)| \leqslant S_0, \qquad s(t) = \frac{1}{\gamma}\ddot{k}(t), \tag{2.2}$$

where the parameter γ is the gyromagnetic ratio. Here G_0 and S_0 are the maximum gradient and slewrate, which are machine dependent constants. For a 3 T GE scanner (General Electric Medical Systems, Milwaukee, WI, USA) the value for the maximum slewrate is 15 G/cm/ms, while the maximum gradient amplitude is 4 G/cm.

2.1.2. Time constraint

While measuring the raw k-space data, there is a finite amount of time in which the signal can be measured before the nuclei need to be re-exited. This leads to the constraint, $t \leqslant T_{max}$. For our implementation on a 3 T scanner we worked under the assumption that T_{max} should be smaller than 60 ms. It is important to note that this constraint is self-imposed. Longer readout times are possible, but could potentially lead to increased susceptibility artifacts in the reconstructed images. Similarly, shorter readout times are also possible but would lead to a decrease in the number of k-space points that can be sampled in a single shot, and thus reduce the spatial resolution of our reconstructed images. To strike the appropriate balance between these two important issues (susceptibility and resolution), we ultimately decided to use a value of $T_{max} \approx 48$ ms in our final trajectory. We refer readers to the section describing the implementation of the sampling trajectory for further details on this issue.

2.1.3. Space-filling constraint

Finally, the trajectory needs to be space-filling, i.e., it needs to satisfy the Nyquist criteria [31]. To better understand this criteria, think of k-space as a lattice where the distance between each point is determined by $2\pi/FOV$, where FOV stands for the field-of-view of the reconstructed image. In order not to violate the Nyquist criteria we need to visit, long enough to make a measurement, each point in the lattice contained within some cubic or spherical region around the center of k-space.

In this paper we present a 3D analogue to EPI sampling [18], called echo-volumar imaging (EVI) [13–15]. Mansfield implemented an EVI trajectory that sampled $64 \times 64 \times 8$ points in k-space. Here we introduce an alternative version that samples a cubic volume in k-space. While volumes of a variety of other shapes could have been chosen (e.g., cylinders, spherical or rectangular volumes), we used a cube in this first implementation of our methodology for a number of reasons. First, it simplifies reconstruction as the fast Fourier transform can be applied to the resulting data, and secondly we were interested in maintaining a roughly equivalent spatial resolution in each direction of the brain.

2.2. An echo-volumar imaging trajectory

Our goal is to design an EVI trajectory that zigzags through 3D k-space with the goal of hitting each coordinate point on a 3D Cartesian grid. In our implementation, such a trajectory will travel from one end of k-space to the other in a straight line and thereafter move to the next line, by traveling along a half circle. This procedure is repeated until all coordinates in a cube of size N^3 are visited. The value of N is determined by the amount of k-space it is possible to cover while still satisfying the necessary constraints. The final trajectory will consist of a collection of straight lines and half circles. To ensure that each straight line consists of the same number of points, each line should begin at the same speed u, accelerate in the first half of the line and de-accelerate in the second half. The trajectory should then travel in a half circle with constant speed u before starting the process again on the next line. The length of the line should be $2m\Delta_k$, where $m = N/2$ and Δ_k is the appropriate spacing in k-space determined by $(2\pi)^{-1}FOV$. Hence our trajectory is designed to maintain a constant slewrate throughout, while the gradient will be non-constant during the acquisition of any given line of k-space. However, the gradient will take a constant value in the curve.

Consider a trajectory $k(t)$ which pieces together $N^2 = 4m^2$ sections such that each section is composed of a straight line of length $2m\Delta_k$ and a half circle with radius r. Suppose that the trajectory has a fixed slew rate, $|\ddot{k}(t)| = \ddot{k}_*$, and that it begins with speed u, accelerates in the first half of the line, de-accelerates in the second half, and then travels in the half circle with constant speed u. Since $r = u^2/\ddot{k}_*$ and the circumference of the half circle is πr, the amount of time the trajectory spends in the half circle is

$$\begin{aligned} t_1 &= \pi r/u \\ &= \pi u/\ddot{k}_*. \end{aligned} \tag{2.3}$$

To determine the time spent in the straight line we invoke the equations of motion for constant acceleration which states that the distance traveled, d, depends on the time t, through the equation

$$d = ut + \frac{at^2}{2}, \tag{2.4}$$

where u is the speed and a is the acceleration. Using this result and the fact that the constant acceleration is equal to $\ddot{k}_*$, the total time spent in the line is $2t_2$ where t_2 can be obtained by solving

$$ut_2 + \ddot{k}_* t_2^2/2 = m\Delta_k. \tag{2.5}$$

Putting these results together the total travel time for the trajectory is given by

$$\begin{aligned} T &= 4m^2(t_1 + 2t_2) \\ &= 4m^2\left(\frac{\pi u}{\ddot{k}_*} + \frac{2}{\ddot{k}_*}\{-u + (u^2 + 2m\Delta_k\ddot{k}_*)^{1/2}\}\right). \end{aligned} \tag{2.6}$$

Since Eq. (2.6) is increasing in u, it is minimized by choosing the smallest possible value of r which is given by Δ_k. Hence, it is optimal to sample adjacent lines as in the case of two-dimensional echo-planar schemes. Since $r = u^2/\ddot{k}_*$, this provides $u^2 = \Delta_k\ddot{k}_*$ and

$$T = 4m^2\{(\pi - 2) + 2(2m + 1)^{1/2}\}\left(\frac{\Delta_k}{\ddot{k}_*}\right)^{1/2}. \qquad (2.7)$$

The best strategy is therefore to make a small turn if we accelerate in the line. An example of such a trajectory starts at the point $(0, 0, k_{z,\min})$ and moves along the z-axis to the point $(0, 0, k_{z,\max})$. Upon reaching this point the trajectory makes a half circular loop over to the point $(1, 0, k_{z,\max})$ and then continues along the z-axis in the opposite direction until it reaches $(1, 0, k_{z,\min})$. The trajectory continues in a similar manner until it has completed a square spiral in the $k_x k_y$-plane (Fig. 1A and B).

Note that this version of the EVI trajectory differs from those currently being used. In conventional EVI, data are acquired in a sequence of N_z two-dimensional $k_x k_y$ planes of k-space, and the trajectory only moves to a new plane after the acquisition of a full $k_x k_y$ plane. In our trajectory we sample along the k_z axis for a fixed value of k_x and k_y. Once we have sampled every point between $k_{z,\min}$ and $k_{z,\max}$ the trajectory moves to a new coordinate point in the $k_x k_y$ plane. We continue this procedure until we have moved through a square spiral in the $k_x k_y$ plane.

2.3. Implementation of EVI trajectory

Consider the discrete case where a trajectory is given by a sequence $\{k(i), i = 1, 2, \ldots, T\}$, where each point is sampled at uniform intervals of length 4 μs regardless of the position of the trajectory. The gradient of a trajectory $k(i)$ is defined as

$$g(i) = |\dot{k}(i)|/\gamma, \qquad (2.8)$$

where $\gamma = 2\pi * 4.257 \times 10^3$ rad Hz/G and we approximate

$$\dot{k}(i) = \frac{k(i) - k(i-1)}{\Delta t}. \qquad (2.9)$$

Here the sampling interval is defined as $\Delta t = (250)^{-1}$ ms. Since $G_0 = 4$ G/cm, it must hold that

$$|k(i) - k(i-1)| = \dot{k}_* \times \Delta t$$
$$< G_0 \gamma \Delta t$$
$$= 4(2\pi * 4.257)/250 = 0.428 \text{ rad/cm}. \qquad (2.10)$$

Similarly, the slew rate is defined as

$$s(i) = |\ddot{k}(i)|/\gamma, \qquad (2.11)$$

where we approximate

$$\ddot{k}(i) = \frac{k(i) + k(i-2) - 2k(i-1)}{(\Delta t)^2}. \qquad (2.12)$$

Since the slew rate is limited to $S_0 = 15$ G/cm/ms, it must hold that

$$|k(i) + k(i-2) - 2k(i-1)| = \ddot{k}_* \times (\Delta t)^2$$
$$< S_0 \gamma (\Delta t)^2$$
$$= 15 * (2\pi * 4.257)/250^2$$
$$= 6.419 \times 10^{-3} \text{ rad/cm}. \qquad (2.13)$$

Assuming the field-of-view (FOV) is equal to 20 cm, the appropriate spacing in k-space must be

$$\Delta_k = 2\pi/\text{FOV}$$
$$= 2\pi/20 \qquad (2.14)$$
$$= 0.314 \text{ cm}^{-1}.$$

Since we sample at a rate of 250 points per ms, the total number of points sampled in a 60 ms time window is 15,000. For $T = 15,000$, $\Delta_k = 0.314$ and $\ddot{k}_* = 6.419^{-3}$, we find from Eq. (2.6) that $7 < m < 8$, where $m = N/2$. Thus, the largest possible m is 7 with $T = 12,187$, $u = 0.04491$, $t_2 = 20.1$, and $|\dot{k}| \leqslant u + \ddot{k}_* t_2 = 0.174$. Since $0.174 < \min(\Delta_k, \dot{k}_*) = 0.314$, the center of the sampling region is not sampled too thinly and the speed of the trajectory is within the limit. This choice of T implies that we are


Fig. 1. (A) An implementation of the echo-volumar imaging trajectory. (B) The echo-volumar imaging trajectory shown in (A) projected onto the $k_x k_y$ plane.

sampling during a time window that is approximately 48 ms long. This is short enough that the resulting image reconstructions will not suffer from severe susceptibility artifacts, while still maintaining a reasonable spatial resolution. The resulting trajectory is an EVI sampling scheme, where the in-and-out pattern starts from the center and works its way out according to the trajectory depicted in Fig. 1. An initial ramp is added to the beginning of the trajectory to assure that the trajectory reaches the appropriate starting position and speed. Fig. 2 shows plots of $(k_x(t), k_y(t), k_z(t))$, $|g(t)|$ and $|s(t)|$.

2.4. Reconstruction of EVI data

Once the central portion of 3D k-space has been collected, the data needs to be reconstructed for statistical analysis. A standard approach towards reconstructing non-uniformly sampled k-space data is to interpolate the data onto a Cartesian grid [32] and thereafter apply the fast Fourier transform (FFT). Our data is sampled on a Cartesian grid in the $k_x k_y$ plane, and it is relatively straightforward to use linear interpolation to get uniformly spaced measurements in the k_z direction as well. After interpolation, our k-space data consists of 2744 (e.g., $14 \times 14 \times 14$) uniformly sampled measurements in 3D k-space. As the data is sampled on a grid, reconstruction is straightforward using the FFT. The data is zero-filled to a resolution of $64 \times 64 \times 64$ prior to reconstruction, and a prolate spheroidal wave function filter (PSWF) [33–37] is applied to reduce truncation artifacts. The PSWF has proven to be an efficient filter to use for handling under-sampled k-space data. Its efficiency has been con-

firmed using both simulations and experimental data in previous work [34,36].

2.5. Statistical analysis

After reconstruction, the statistical analysis is performed voxel-wise using a two step procedure that is described in greater detail in a companion paper [29]. In the first step we detect regions in the brain where there is a significant positive BOLD signal. These are the regions that would typically be categorized as having task-induced neuronal activation in a standard fMRI analysis. We will ultimately be more concerned with detecting regions with significant negative BOLD signal. However, we feel it is a natural assumption that regions having a significant dip will also ultimately have a positive rise in BOLD signal if they are involved in the task at hand. This step therefore works as a simple screening process to remove uninteresting voxels, where there are no signs of task-induced activation, and allows for closer and more data-intensive inspection of voxels that are actually involved in the task. In the second step of the analysis we calculate bootstrap distributions for the amplitude of the negative dip, as well as for the time of the onset of the negative dip. Using these distributions we can perform statistical tests to determine whether there is a significant negative dip in a voxel, as well as compare the relative timing of the dips across various regions of the brain using an equivalent metric as that used when studying the relative timing of the rise [38–40].

This second step is important as there is evidence that the negative dip is more localized to areas of neural activity [9,10,12] than the subsequent rise which appears less

Fig. 2. (A–E) Plots of $k_x, k_y, k_z, |g|, |s|$ for the echo-volumar imaging trajectory.

spatially specific. Hence, we may be able to prune away voxels that are simply adjacent to regions involved in the neuronal activity, rather than being directly involved. In addition, since the negative dip occurs in a time scale closer to that of the neural activity, there would appear to be less confounding factors influencing the order of the dip in comparison to the order of the positive rise. Below follows a more detailed outline of our two-step analysis.

Step One: After reconstructing the data into image-space we are left with a sequence of T three-dimensional images, each of size $64 \times 64 \times 64$. We model the fMRI time course using a classical decomposition model. The fMRI signal from voxel i, $i = 1, \ldots, 64^3$ can be modeled as:

$$Y_i(t) = \mu_i(t) + s_i(t) + z_i(t) + \epsilon_i(t). \tag{2.15}$$

Here $\mu_i(t)$ is a drift term which we model using a quadratic polynomial function (three parameters). The term $s_i(t)$ is a seasonal component that is due to heart-rate and respiration, modeled using a finite impulse response (FIR) basis set containing one free parameter for every time point of the response we seek to model. Including this term adds d parameters to the model, where d is the periodicity of the seasonal component (in units of 0.1 s). The term $z_i(t)$ represents the BOLD response to the neuronal stimuli, which is modeled as the convolution of the stimulus function with a canonical hemodynamic response function (e.g., SPMs double gamma function) [41]. To increase its flexibility to handle slight temporal shifts in the onset of activation, we also include a term for the temporal derivative [41,42]. Finally, $\epsilon_i(t)$ is the noise present in the MR signal, which is modeled using an AR(2) process which adds three additional terms to our model. In summary our model has a total of $d + 5$ regression parameters and 3 variance component parameters.

For each voxel, the model is fit using an iterative generalized least-squares approach, where the variance compo-nents and regression coefficients are alternately calculated and updated. Thereafter, t-statistics corresponding to the HRF regressors are calculated, and statistical maps of the voxel-wise t-statistics are constructed. The maps are thresholded and corrected for multiple comparisons using the false detection rate (FDR) procedure [43,44]. This allows us to determine regions with statistically significant signal corresponding to the positive BOLD signal.

Step Two: Time courses from each of the voxels deemed active in Step One are decomposed into signal, trend and seasonality components (see Fig. 3), and the seasonal and trend components are removed from the time course. The remaining time series is averaged over the m repetitions of the stimuli to obtain estimates of the HRF. Using the estimated HRF from each voxel we can calculate the maximum amplitude of the positive rise, the minimum amplitude of the negative dip, as well as the time of onset of both the rise and dip.

Thereafter, for each time series, a statistical test based on the bootstrap procedure [45] is performed to test whether the dips are statistically significant, i.e., whether the amplitude of the dips are significantly different from zero. For voxels with significant dips, a bootstrap test for the pair-wise difference in the time of onset of the dip is performed to determine whether the onset is significantly different between voxels across the brain. This allows us to determine the order in which the dip occurs in various regions associated with the task. We refer interested readers to our statistical companion paper [29] for a more in-depth discussion on the statistical analysis.

2.6. Experimental design

Both the EVI trajectory and the reconstruction algorithm were implemented in Matlab (Mathworks Inc.). To

Fig. 3. (A) A typical time course decomposed into (B) quadratic drift, (C) periodic nuisance parameters and (D) fMRI signal. The length of the period for the nuisance parameters is approximately 3 s and represents artifacts due to respiration.

demonstrate the methods utility for dynamic studies, two high temporal resolution fMRI experiments were designed to track the hemodynamic signals in the brain while the subject undergoes a visual–motor activation paradigm and an auditory–motor–visual activation paradigm.

The first activation paradigm consisted of fifteen cycles of 20 s intervals. At the beginning of each interval a 100 ms light flash was presented. The subject was instructed to press a button with their right thumb immediately after sensing the flash, thereby leading to activation of the motor cortex. During the 20 s interval, images were acquired rapidly every 100 ms using our cubic EVI trajectory. The sequence was repeated 15 times, each time producing a dynamic data set of 200 temporal points.

The second activation paradigm also consisted of 15 cycles of 20 s intervals. At the beginning of each interval a tone was sounded through headphones which the subject was wearing. The subject was instructed to press a button with their right thumb immediately after hearing the tone. Upon pressing the button a 100 ms light flash was presented, leading to activation of the visual cortex. During the 20 s interval, images were acquired rapidly every 100 ms using our cubic EVI trajectory. The sequence was repeated 15 times, each time producing a dynamic data set of 200 temporal points.

A healthy male volunteer participated in the study after giving informed consent in accordance with a protocol approved by the Stanford Institutional Review Board. In both experiments the first cycle was thrown out and the resulting data consisted of 14 cycles with a total of 2800 time units. The resulting k-space data was reconstructed and statistical analysis was performed as outlined in the previous section. The data was acquired with an effective TE 30 ms, flip angle 20°, field-of-view 200×200 mm^2, slice thickness 185 mm and bandwidth 250 kHz. The experiment was performed on a 3.0 T whole body scanner (GE magnet, General Electric Medical Systems, Milwaukee, WI, USA). T_2-weighted FSE scans were obtained for anatomic reference (TR/TE/ETL = 3000 ms/68 ms/12, 5 mm interleaved contiguous slices, FOV = 24 cm, 256×128 matrix).

3. Results

The feasibility of our rapid 3D imaging approach was tested experimentally using a visual–motor and an auditory–motor–visual stimulation paradigm, both described in the previous section. After data collection, the raw k-space data was reconstructed into images of size $64 \times 64 \times 64$, and the time courses corresponding to the 64^3 voxels were analyzed for activation. Examples of the raw EVI images that were obtained can be seen in the top row of Fig. 4. These images correspond roughly to the anatomical images shown in the rows below.

In a first step statistical analysis was performed using the approach outlined above. The design matrix consisted of three columns corresponding to a quadratic trend model for the signal drift, as well as, d columns corresponding

to the periodicity of the seasonal component. As the respiration is the dominant source of seasonality in the data we used its periodicity, which was empirically determined to be 3 s (i.e., $d = 30$), to determine the number of parameters. In addition, two extra parameters corresponding to the canonical HRF and its temporal derivative were added, giving a design matrix consisting of a total of 35 columns. The middle row of Fig. 4 shows examples of statistical parametric maps for two slices of the brain indicating voxels with significant task-related activity (i.e., positive BOLD response) using a t-test (p-value $\leqslant 0.01$). A clear activation pattern is present both in the visual and motor cortices as would be expected. For voxels that were deemed active in the GLM analysis, their respective time courses were analyzed further. Each time course was decomposed into a trend component, a signal component and a seasonal component. Fig. 3 shows the results of the decomposition of a representative time course. The quadratic trend and the seasonal component were removed from each time course and only the signal component is brought forth to the next stage of the analysis.

For each voxel deemed to have a significant rise according to the GLM analysis, the timing of the onset of the rise was estimated. Results for two slices, one centered in the visual and the other in the motor cortex, are shown in the middle row of Fig. 4. It is clear that the onset appears in the visual cortex prior to the motor cortex. Among active voxels, a statistical test based on the use of the bootstrap procedure, was performed to test for significant dips (p-value < 0.05). The timing of the onset of the dip was estimated for active voxels and the results for the same two slices are shown in the bottom row of Fig. 4. Again it is clear that the activation appears to be occurring in the visual cortex prior to the motor cortex.

Fig. 5A shows the averages for two time courses extracted from the center of the visual and motor cortices, respectively. We clearly see that the HRF estimated from the visual cortex proceeds the one estimated from the motor cortex throughout the course of the 20 s run. Fig. 5B shows a close-up of the first 3 s following activation. A negative dip appears first in the visual cortex, as makes sense since the visual cortex is logically the first region of the brain that begins to work on the image. After a few hundred milliseconds delay we see a delayed negative response in the motor cortex. Bootstrap tests (Fig. 5C–E) confirm these results, and show that while both dips are significant, the dip in the visual cortex occurs at a significantly earlier time point compared to that of the motor cortex. In this experiment both the timing of the dip and rise give compelling evidence that neuronal activity is taking place in the visual cortex prior to the motor cortex as would be expected by the experimental paradigm.

The exact same statistical analysis was repeated for the auditory–motor–visual stimulation paradigm. For each voxel deemed to have a significant rise according to the initial analysis, the onset of the positive BOLD signal was estimated. Results for two slices, centered in the visual
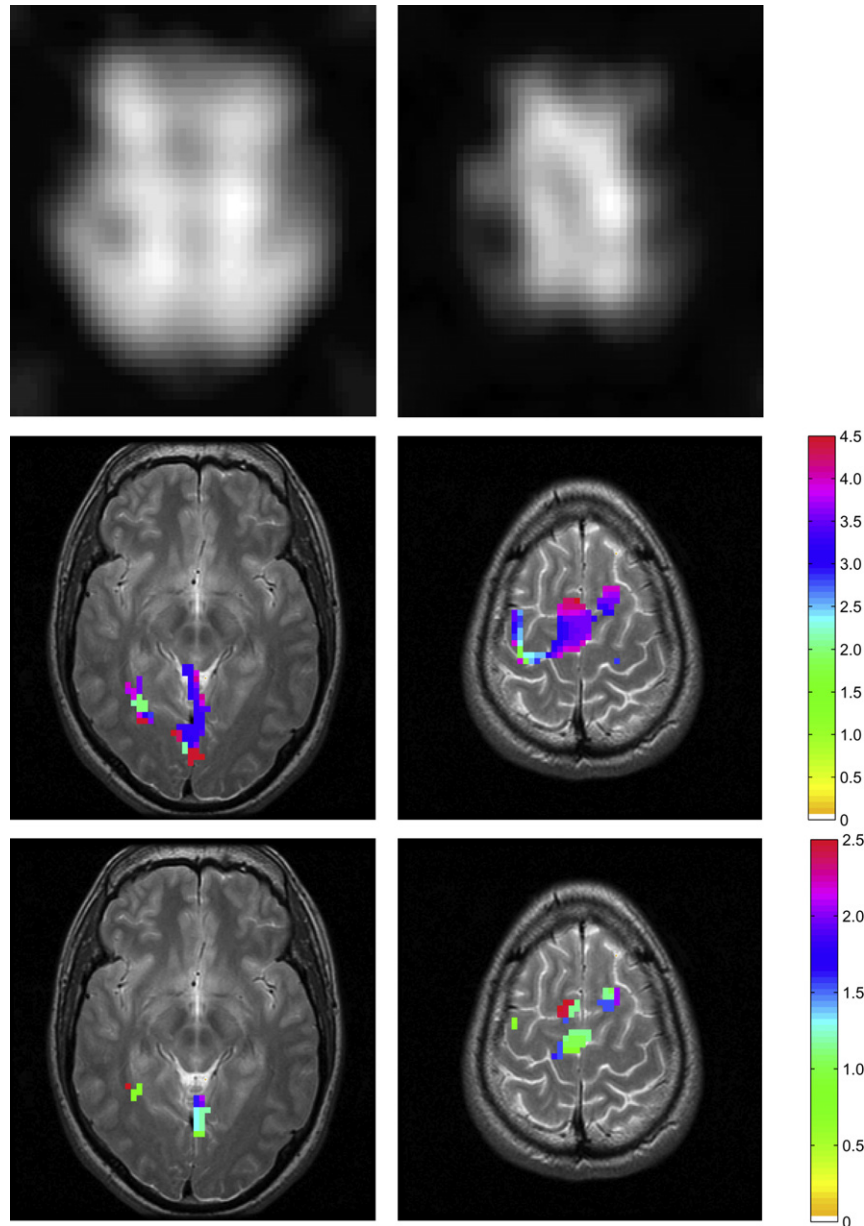
Fig. 4. (Top row) Examples of two raw EVI images that correspond roughly to the anatomical images shown in the rows below. (Middle row) Maps depicting the onset of the positive BOLD signal for voxels with significant activation are shown for two slices in the visual–motor experiment. The slices cover the visual (left) and motor (right) cortices. The results indicate that the rise appears earlier in the visual then the motor cortex. (Bottom row) Maps depicting the onset of the negative dip in voxels with significant dips for the same two slices. The dip appears earlier in the visual then the motor cortex.

and motor cortices, respectively, are shown in the top row of Fig. 6. It is important to note that in this experiment the order of activation should now be auditory, followed by motor, followed by visual. Fig. 7 shows time courses from these three regions. While, the rise in the auditory cortex does appear first, there appears to be confounding in the timing of the onset in the visual and motor cortices. Clearly the signal over the motor cortex starts rising after the visual cortex. However, studying the dip alleviates this confounding, which can be seen in both Figs. 6 and 7B. The difference in the onset of the dip between the visual and motor cortices is not statistically different from zero. However, this is hardly surprising as the visual stimulus appears immediately after the button press. Also note that both the dip and rise in the auditory cortex appears before that of the visual and motor cortices. This is to be expected as the tone is presented before the button is pressed and the visual stimulus is presented.

4. Discussion

This paper introduces a novel approach to the acquisition and analysis of rapid 3D fMRI data. There are several benefits to rapid imaging. First, it allows one to study the initial negative BOLD response, instead of solely depending on the positive BOLD signal, for purposes of
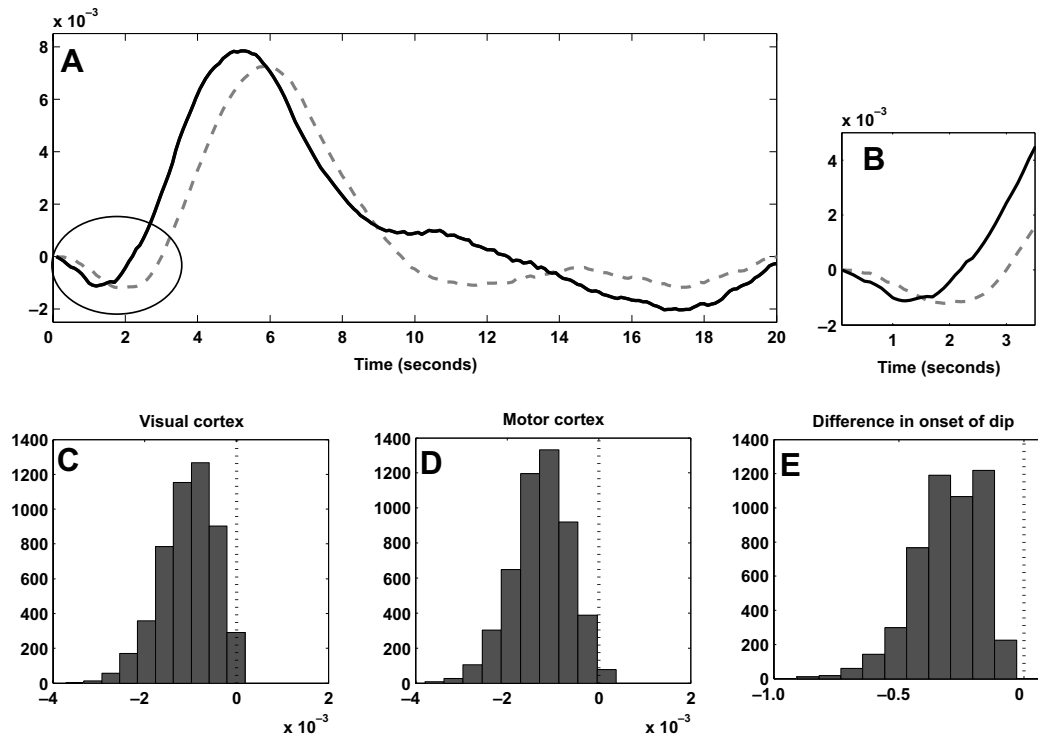
Fig. 5. (A) Time courses from the visual (bold) and motor (dashed) cortices averaged over the 14 cycles of the visual–motor stimulus. (B) The first 3 s following stimulation for the two time courses appearing in (A). The dip appears earlier in the visual cortex than the motor cortex, which is consistent with the experimental paradigm. (C–E) The results of bootstrap tests show significant dips in the visual and motor cortices, as well as a significant difference in the onset of the dip between the two regions.

determining active regions as well as the order of activation in multiple active regions. Our results indicate that the negative response may contain valuable information regarding the timing of activation, information that may be confounded when studying the positive BOLD response. However, it should be noted that in both experiments we performed there appears to be an increase in the number of false negatives when looking for voxels with significant dips compared to rise. This is made clear by comparing the results presented in the bottom two rows of both Figs. 4 and 6. The reason for this discrepancy may be due to the fact that the signal-to-noise ratio for the rise is on the order of 5 times larger than that of the dip. Alternatively, there is evidence that the initial negative BOLD response is more localized to areas of neural activity [9–12] than the subsequent rise. This is reasonable since if the CBF response is widespread, then the dip should provide better spatial localization than the positive rise as it is based solely on local vasculature. Hence, the decreased number of active voxels may in fact be giving a more accurate picture of the true activation patterns in the brain. However, this statement is hard to verify, in part because the low spatial resolution provided by the current implementation of the method.

The second benefit of rapid imaging is that it allows for the efficient removal of physiological noise due to cardiac and respiratory effects. In a typical fMRI analysis, the time resolution is on the order of 2 s. Since respiration gives rise to a periodic function, with a period length of approximately 3 s, the Nyquist criteria does not allow us to fully reconstruct the signal. Rapid fMRI (with a temporal resolution on the order of 100 ms) circumvents this issue and allows for efficient reconstruction of the underlying signal without the problem of aliasing. This is beneficial as it allows us to significantly clean up the fMRI signal prior to analysis and obtain more accurate estimates of the hemodynamic response function.

Finally, rapid imaging alleviates issues related to the fact that spatially separate regions of the brain are sampled at different times, thus negating the need for slice-time correction. In addition, it may also allow for more accurate correction of subject motion, as movement occurring during the acquisition of each individual volume will be reduced.

The main drawback to our approach is that we have sacrificed spatial resolution in order to increase temporal resolution. However, advances in multi-coil techniques [27,28] give an avenue to bridge this gap in the future. In these techniques, multiple k-space measurements can simultaneously be made and prior knowledge about RF field distributions or the image sensitivity of the coils can be utilized to construct images from under-sampled k-space data. We have recently implemented two new trajectories that allow us to obtain images with a temporal resolution of 100 ms and a spatial resolution on the order of $25 \times 25 \times 17$ and $46 \times 46 \times 17$, respectively. With the latter

Fig. 6. (Top row) Maps depicting the onset of the positive rise in voxels with significant activation for two slices in the auditory–motor–visual experiment. The slices that are included cover the visual (left) and motor cortices (right). (Bottom row) Maps depicting the onset of the dip in voxels with significant dips for the same two slices.
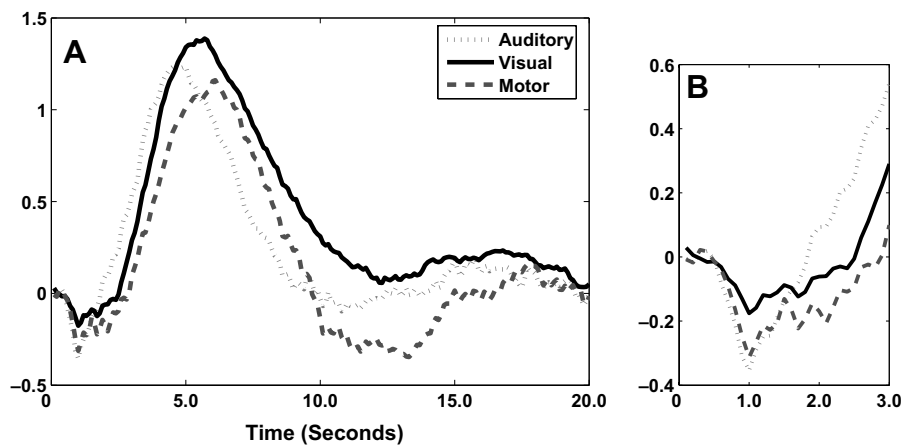

Fig. 7. (A) Time courses from the auditory (dotted), visual (bold) and motor (dashed) cortices averaged over the 14 cycles. (B) The first 3 s following stimulation for the three time courses appearing in (A). The dip appears first in the auditory cortex, followed by the motor and visual cortices which are not statistically differentiable. However, the rise in the motor signal appears after the rise in the visual signal which implies an incorrect temporal ordering of activation.

spatial resolution we are quickly approaching the resolution that is used in standard fMRI experiments. However, the temporal resolution is increased 10-fold. The results of these trajectories will be presented in future work.

To summarize our findings, we have shown reproducibly that significant dips are present in both the visual and motor cortices. In addition, there is a statistically signifi-cant difference in the time of onset of the dip between the visual and motor cortices in our first experiment, which is consistent with the experimental paradigm. In the second experiment there is also a significant dip in the auditory cortex (again consistent with the paradigm) followed by dips in the visual and motor cortices which are not statisti-cally differentiable. However, if we instead use the onset of

the rise as a metric, the order of activation between the visual and motor cortices is confounded. This strengthens our notion that studying the initial negative BOLD response can be a critical tool for determining the exact timing of activation in the brain. However, it is important to keep in mind that the results presented in this work are obtained from a single subject. Therefore, this work should be considered more of a proof of concept of our imaging technique, rather than a proof of our ability to reliably and reproducibly detect the initial negative BOLD signal in a population of subjects. Further experiments need to be performed to determine (i) whether the dip can be reliably detected over multiple subjects, and (ii) whether in these subjects it provides an accurate picture of the timing of activation. However, we feel that if these results were to hold up, the technique would be an extremely valuable tool in neuroimaging studies.

5. Conclusions

A new approach towards rapid fMRI is introduced where a small central region of 3D k-space is sampled every 100 ms and a low spatial resolution snapshot of the brain with extremely high temporal resolution is obtained. In addition we discuss a new approach towards the statistical analysis of the resulting high-frequency fMRI data. The feasibility and efficiency of the combined acquisition and analysis approach is confirmed using data from a visual–motor task and an auditory–motor–visual task. The increased temporal resolution allows us for the first time to perform a statistical analysis over the brain based solely on the initial negative BOLD response, rather than the sluggish positive BOLD response. In the visual–motor experiment there are coherent regions in both the visual and motor cortices with a significant initial negative BOLD signal. Further, the onset of the negative response is shown to be significantly earlier in the visual cortex which is consistent with the experimental paradigm. In the auditory–motor–visual experiment there is a significant dip in the auditory cortex followed by dips in the visual and motor cortices which are not statistically differentiable. However, if we instead use the onset of the peak positive BOLD response as a metric, the order of activation between the visual and motor cortices are confounded. This leads us to believe that studying the initial negative BOLD response can be an important tool for determining the timing of activation across different regions of the brain. This paper provides a proof of concept of our methodology. Further experiments need to be performed to verify its reproducibility.

References

[1] S. Ogawa, D. Tank, R. Menon, J. Ellerman, S. Kim, H. Merkle, K. Ugurbil, Intrinsic signal changes accompanying sensory simulation: functional brain mapping and magnetic resonance imaging, Proc. Natl. Acad. Sci. USA 89 (1992) 5951–5955.

[2] S. Ogawa, T. Lee, B. Barrere, The sensitivity of magnetic resonance image signals of a rat brain to changes in the cerebralvenous blood oxygenation, Magn. Reson. Med. 29 (1993) 205–210.

[3] S. Ogawa, R. Menon, D. Tank, S. Kim, H. Merkle, J. Ellerman, K. Ugurbil, Functional brain mapping by blood oxygenation level-dependent contrast magnetic resonance imaging: a comparison of signal characteristics with a biophysical model, Biophys. J. 64 (1993) 803–812.

[4] Z. Cho, Y. Ro, T. Lim, NMR venography using the susceptibility effect produced by deoxyhemoglobin, Magn. Reson. Med. 28 (1992) 25–38.

[5] T. Ernst, J. Hennig, Observation of a fast response in functional MR, Magn. Reson. Med. 32 (1994) 146–149.

[6] R. Menon, S. Ogawa, X. Hu, J. Strupp, P. Andersen, K. Ugurbil, Bold based functional MRI at 4 tesla includes a capillary bed contribution: echo-planar imaging mirrors previous optical imaging using intrinsic signals, Magn. Reson. Med. 33 (1995) 453–459.

[7] D. Malonek, A. Grinvald, The imaging spectroscopy reveals the interaction between electrical activity and cortical microcirculation: implication for optical PET and MR functional brain imaging, Science 272 (1996) 551–554.

[8] E. Yacoub, T. Le, X. Hu, Detecting the early response at 1.5 tesla, NeuroImage 7 (1998) S266.

[9] E. Yacoub, A. Shmuel, J. Pfeuffer, P. Van De Moortele, G. Adriany, K. Ugurbil, X. Hu, Investigation of the initial dip in fmri at 7 tesla, NMR Biomed. 14 (2001) 408–412.

[10] T. Duong, D. Kim, K. Ugurbil, S. Kim, Spatio-temporal dynamics of the BOLD fMRI signals: toward mapping columnar structures using the early negative response, Magn. Reson. Med. 44 (2000) 231–242.

[11] D. Kim, T. Duong, S. Kim, High-resolution mapping of iso-orientation columns by fMRI, Nat. Neurosci. 3 (2000) 164–169.

[12] J. Thompson, M. Peterson, R. Freeman, High-resolution neurometabolic coupling revealed by focal activation of visual neurons, Nat. Neurosci. 7 (2004) 919–920.

[13] P. Mansfield, A. Howseman, R. Ordidge, Volumar imaging using NMR spin echos: echo-volumar imaging (EVI) at 0.1 T, J. Phys. E 22 (1989) 324–330.

[14] P. Mansfield, R. Coxon, J. Hykin, Echo-volumar imaging (EVI) at 3.0 T: first normal volunteer and functional imaging results, J. Comput. Assist. Tomogr. 19 (1995) 847–852.

[15] P. Harvey, P. Mansfield, Echo-volumar imaging (EVI) at 0.5 T: first whole-body volunteer studies, Magn. Reson. Med. 35 (1996) 80–88.

[16] Y. Yang, V. Mattay, D. Weinberger, J. Frank, J. Duyn, Localized echo-volume imaging methods for functional MRI, J. Magn. Reson. Imaging 7 (1997) 371–375.

[17] W. van der Zwaag, S. Francis, R. Bowtell, Improved echo volumar imaging (EVI) for functional MRI, Magn. Reson. Med. 56 (2006) 1320–1327.

[18] P. Mansfield, Multi-planar image formation using NMR spin echoes, J. Phys. C10 (1977) L55–L58.

[19] P. Irarrazabal, D. Nishimura, Fast three dimensional magnetic resonance imaging, Magn. Reson. Med. 33 (1997) 656–662.

[20] S. Sabat, R. Mir, M. Guarini, A. Guesalaga, P. Irarrazaval, Three dimensional k-space trajectory design using genetic algorithms, Magn. Reson. Imaging 21 (2003) 755–764.

[21] R. Mir, A. Guesalaga, J. Spiniak, M. Guarini, P. Irarrazaval, Fast three-dimensional k-space trajectory design using missile guidance ideas, Magn. Reson. Med. 52 (2004) 329–336.

[22] J. van Vaals, M. Brummer, W. Dixon, H. Tuithof, H. Engels, R. Nelson, B. Gerety, J. Chezmar, J. den Boer, Keyhole method for accelerating imaging of contrast agent uptake, J. Magn. Reson. Imaging 3 (1993) 671–675.

[23] J. Gao, J. Xiong, S. Lai, E. Haacke, M. Woldorff, J. Li, P. Fox, Improving the temporal resolution of functional MR imaging using keyhole techniques, Magn. Reson. Med. 35 (1996) 854–860.

[24] G. Zientara, L. Panych, F. Jolesz, Dynamically adaptive MRI with encoding by singular value decomposition, Magn. Reson. Med. 32 (1994) 268–274.

[25] L. Panych, C. Oesterle, G. Zientara, J. Hennig, Implementation of a fast gradient-echo svd encoding technique for dynamic imaging, Magn. Reson. Med. 35 (1996) 554–562.

[26] Z. Liang, P. Lauterbur, An efficient method for dynamic magnetic resonance imaging, IEEE Trans. Med. Imaging 13 (1994) 677–686.

[27] D. Sodickson, W. Manning, Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays, Magn. Reson. Med. 38 (1997) 591–603.

[28] K. Pruessmann, M. Weiger, M. Scheidegger, P. Boesiger, SENSE: sensitivity encoding for fast mri, Magn. Reson. Med. 42 (1999) 952–956.

[29] M. Lindquist, C. Zhang, G. Glover, L. Shepp, Q. Yang, Acquisition and statistical analysis of rapid 3d fMRI data, Stat. Sin. (2008), in press.

[30] N. Logothetis, Can current fMRI techniques reveal the micro-architecture of cortex? Nat. Neurosci. 3 (2000) 413.

[31] E. Haacke, R. Brown, M. Thompson, R. Venkatesan, Magnetic Resonance Imaging: Physical Principles and Sequence Design, John Wiley & Sons, Inc., 1999.

[32] J. Jackson, C. Meyer, D. Nishimura, A. Macovski, Selection of a convolution function for fourier inversion using gridding, IEEE Trans. Med. Imaging 10 (1991) 473–478.

[33] L. Shepp, C.-H. Zhang, Fast functional magnetic resonance imaging via prolate wavelets, Appl. Comput. Harmonic Anal. 9 (2000) 99–119.

[34] Q. Yang, M. Lindquist, L. Shepp, C. Zhang, J. Wang, M. Smith, The two dimensional prolate spheroidal wave function for MRI, J. Magn. Reson. 58 (2002) 43–51.

[35] M. Lindquist, Optimal data acquisition in fMRI using prolate spheroidal wave functions, Int. J. Imaging Syst. Technol. 13 (2003) 803–812.

[36] M. Lindquist, C. Zhang, G. Glover, L. Shepp, Q. Yang, A generalization of the two dimensional prolate spheroidal wave function method for non-rectilinear MRI data acquisition methods, IEEE Trans. Image Process. 15 (2006) 2792–2804.

[37] M. Lindquist, T. Wager, Spatial smoothing in fMRI using prolate spheroidal wave functions, Hum. Brain Mapp. (2008), doi:10.1002/hbm.20475.

[38] W. Richter, K. Ugurbil, A. Georgopoulos, S. Kim, Time-resolved fMRI of mental rotation, Neuroreport 8 (1997) 3697–3702.

[39] S. Kim, W. Richter, K. Ugurbil, Limitations of temporal resolution in functional MRI, Magn. Reson. Med. 37 (1997) 631–636.

[40] R. Menon, D. Luknowsky, J. Gati, Mental chronometry using latency-resolved functional MRI, Proc. Natl. Acad. Sci. USA 95 (1998) 10902–10907.

[41] K. Friston, O. Josephs, G. Rees, R. Turner, Observation of a fast response in functional MR, Magn. Reson. Med. 39 (1998) 41–52.

[42] R. Henson, C. Price, M. Rugg, R. Turner, K. Friston, Detecting latency differences in event-related bold responses: application to words versus nonwords and initial versus repeated face presentations, NeuroImage 15 (2002) 83–97.

[43] Y. Benjamini, Y. Hochberg, Controlling the false discovery rate: a practical and powerful approach to multiple testing, J. R. Stat. Soc. Ser. B 57 (1995) 289–300.

[44] C. Genovese, N. Lazar, T. Nichols, Thresholding of statistical maps in functional neuroimaging using the false discovery rate, NeuroImage 15 (2002) 870–878.

[45] B. Efron, R. Tibshirani, An Introduction to the Bootstrap, Chapman & Hall, CRC, 1998.