ELSEVIER

NeuroImage

www.elsevier.com/locate/ynimg
NeuroImage 24 (2005) 207–215

A signal processing model for arterial spin labeling functional MRI

Thomas T. Liu[a,*] and Eric C. Wong[a,b]

[a]*Center for Functional Magnetic Resonance Imaging and Department of Radiology, University of California-San Diego, La Jolla, CA 92093-0677, USA*
[b]*Department of Psychiatry, University of California-San Diego, La Jolla, CA 92093-0677, USA*

Received 8 April 2004; revised 7 September 2004; accepted 28 September 2004

A model of the signal path in arterial spin labeling (ASL)-based functional magnetic resonance imaging (fMRI) is presented. Three subtraction-based methods for forming a perfusion estimate are considered and shown to be specific cases of a generalized estimate consisting of a modulator followed by a low pass filter. The performance of the methods is evaluated using the signal model. Contamination of the perfusion estimate by blood oxygenation level dependent contrast (BOLD) is minimized by using either sinc subtraction or surround subtraction for block design experiments and by using pair-wise subtraction for randomized event-related experiments. The subtraction methods all tend to decorrelate the 1/*f* type low frequency noise often observed in fMRI experiments. Sinc subtraction provides the flattest noise power spectrum at low frequencies, while pair-wise subtraction yields the narrowest autocorrelation function. The formation of BOLD estimates from the ASL data is also considered and perfusion weighting of the estimates is examined using the signal model.

Keywords: fMRI; Perfusion; Arterial spin labeling; BOLD

Introduction

Perfusion-based functional magnetic resonance imaging (fMRI) using arterial spin labeling (ASL) methods has the potential to provide better localization of the functional signal to sites of neural activity as compared to blood oxygenation level dependent (BOLD) contrast fMRI (Duong et al., 2001; Golay et al., 2004; Luh et al., 2000). In an ASL experiment, a series of control images and tag images in which arterial blood is either fully relaxed or magnetically inverted, respectively, is acquired. Typically, the control and tag images are acquired in an interleaved fashion, and a perfusion time series can be formed by subtracting the tag images from the control images. The three subtraction methods that are most widely used in practice are (a) a simple pair-wise subtraction of control and tag images, (b) a surround subtraction in which the difference between each image and the average of its two nearest neighbors is formed, and (c) a subtraction of sinc-interpolated control and tag images. Surround subtraction was introduced to reduce transient artifacts due to blood oxygenation level dependent (BOLD) weighting of the acquired images (Wong et al., 1997). Aguirre et al. (2002) later showed using simulations of block designs that sinc subtraction performed slightly better than surround subtraction and significantly better than pair-wise subtraction with respect to BOLD contamination effects.

An advantage of the subtraction process in ASL is that it tends to whiten the low frequency 1/*f* noise typically observed in fMRI experiments (Aguirre et al., 2002). Using experimental data acquired at 1.5 and 4 T, Wang et al. (2003) found that sinc subtraction and pair-wise subtraction provided relatively flat power spectra in the frequency range below 0.10 Hz while surround subtraction showed a small increase in power at lower frequencies. A qualitative explanation of the noise whitening process was provided based upon treating the pair-wise subtraction process as a temporal derivative operator.

Despite the widespread use of subtraction-based methods for ASL, a theoretical model that would allow for the quantitative assessment of the various methods has been lacking. Here we present a model of the ASL signal processing chain that allows for an analytical evaluation of the performance of the various subtraction methods. We show that pair-wise, surround, and sinc subtraction are all specific cases of a general subtraction method that differ only in the selection of a low pass filter. We then apply the model to derive expressions for the BOLD contamination of the perfusion estimate and for the perfusion weighting of the BOLD estimate formed from the running average of the tag and control images. Using these expressions, we describe the selection of an optimal filter to minimize BOLD contamination and perfusion weighting. We also derive an expression for the temporal autocorrelation of the noise in the perfusion signal. A preliminary version of this work was presented in Liu and Wong (2004).

* Corresponding author. UCSD Center for Functional MRI, 9500 Gilman Drive, MC 0677, La Jolla, CA 92093-0677. Fax: +1 858 822 0605.
E-mail address: ttliu@ucsd.edu (T.T. Liu).
Available online on ScienceDirect (www.sciencedirect.com.)

1053-8119/$ - see front matter © 2004 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2004.09.047

Theory

Signal processing model

There are currently three major classes of arterial spin labeling methods: (a) pulsed ASL in which a brief radio frequency pulse is used to invert the magnetization of arterial blood within a thick slab proximal to the imaging region, (b) continuous ASL in which a long radio frequency pulse is used to invert spins as they flow through a thin plane proximal to the imaging region, and (c) velocity-selective ASL in which velocity-selective radio frequency pulses are used to saturate arterial spins flowing above a specified cut-off velocity (Duhamel et al., 2003; Wong et al., 1997). At present, the majority of ASL perfusion experiments utilize pulsed ASL, in part because of its ease of implementation. Although the model we present is generally applicable to all three classes of ASL methods, it does include details, such as the use of a presaturation pulse, that are specific to pulsed ASL experiments.

A signal processing model that captures the essential features of a pulsed ASL experiment is shown in Fig. 1. The measured time series $y[n]$ is the sum of the interleaved BOLD-weighted tag and control images plus an additive noise term $e[n]$ with autocorrelation function $\rho[n]$. Multiplicative BOLD weighting is represented by the time series $b[n] = \exp(-TE(R^*_{2,0} + \Delta R^*_2[n]))$ where TE is the echo time of the image, $R^*_{2,0}$ is the apparent transverse relaxation rate at rest, and $\Delta R^*_2[n]$ is the time-varying change in relaxation rate with functional activation. Because the percent change in the signal is typically on the order of a few percent, it is useful to approximate the time series as the sum $b[n] \approx b_0 + \Delta b[n]$ of a constant term $b_0 = \exp(-R^*_{2,0}TE)$ representing baseline BOLD weighting and a time-varying term $\Delta b[n] = -b_0 \Delta R^*_2[n]TE$ representing the stimulus-related changes.

The term $M[n]$ represents the magnetization of static tissue in the imaging volume of interest. If a presaturation pulse is applied at time TI_p prior to the image acquisition, the static tissue term is multiplied by a saturation recovery term $(1 - \beta e^{-TI_p/T_1})$ with $\beta = 1$ and tissue longitudinal time constant T_1. If a presaturation pulse is not used and the tagging process does not perturb the static tissue magnetization, then $\beta = 0$. If the tagging process leaves static tissue inverted, then $\beta = 2$ and $TI_p = TI$.

The term $q[n](1 - \alpha(1 + (-1)^n)e^{-TI/T_{1B}})$ represents the interleaved control and tag perfusion terms where $q[n] = A_{eff}\,CBF[n]$ is proportional to the amount of blood that perfuses the tissue (with CBF denoting cerebral blood flow) and A_{eff} reflects the details of the pulsed ASL pulse sequence used. In the control state (n odd), the arterial blood is assumed to be fully relaxed. In the tag state (n even), the magnetization of arterial blood is assumed to be inverted with inversion efficiency α at inversion time TI prior to image acquisition, and recovers with time constant T_{1B}.

The measured time series may be written as the sum $y[n] = y_b[n] + y_q[n] + e[n]$ of an unmodulated term

$$y_b[n] = b[n]\left(s_M M[n] + s_q q[n]\right) \tag{1}$$

where $s_M = 1 - \beta e^{-TI_p/T_1}$ and $s_q = 1 - \alpha e^{-TI/T_{1B}}$, a modulated term

$$y_q[n] = (-1)^{n+1} b[n] q[n] \alpha e^{-TI/T_{1B}}, \tag{2}$$

and the additive noise term. The unmodulated term is the sum of BOLD-weighted static tissue and perfusion components, while the modulated term is a BOLD-weighted perfusion signal. Perfusion estimates attempt to attenuate the unmodulated term $y_b[n]$ while preserving the modulated term $y_q[n]$. In contrast, BOLD estimates strive to maximize the unmodulated term and minimize the modulated term.

Perfusion and BOLD estimates

Subtraction-based perfusion estimates are based upon differences of control and tag images. In pair-wise subtraction, the estimates are the simple differences of adjacent control and tag images, for example, $y[1] - y[0]$, $y[1] - y[2]$, $y[3] - y[2]$. In surround subtraction, the estimates are the differences between each image and the average of its nearest neighbors, for example, $y[1] - (y[0] + y[2])/2$, $(y[1] + y[3])/2 - y[2]$. Both of these estimates may be written in the form

$$\hat{q}[n] = \left[(-1)^{n+1} y[n]\right] * g[n] \tag{3}$$

where $g[n]$ is a low pass interpolation filter that depends on the subtraction method used and $*$ denotes convolution (defined as $y[n] * g[n] = \sum_{k=-\infty}^{\infty} y[k]g[n-k]$; Oppenheim and Schafer, 1989). For pair-wise subtraction $g[n] = [1\ 1]$ and for surround subtraction $g[n] = [1\ 2\ 1]/2$. A useful generalization of the subtraction process can be gained by dividing the measured time series into a control time series $y_{con}[n] = y[2n+1]$ and a tag time series $y_{tag}[n] = y[2n]$ (Liu et al., 2002). Interpolated versions of

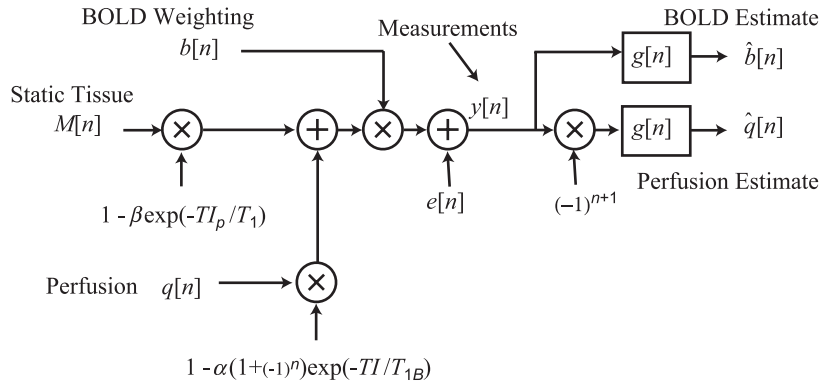

Fig. 1. Block diagram of the arterial spin labeling signal path.

the control and tag time series, denoted as $\hat{y}_{\text{con}}[n]$ and $\hat{y}_{\text{tag}}[n]$, respectively, are formed by upsampling and convolving with the filter $g[n]$. The perfusion estimate is then obtained from the difference

$$\hat{q}[n] = \hat{y}_{con}[n] - \hat{y}_{tag}[n] \tag{4}$$

of the interpolated time series. As shown in the Appendix, this estimate is identical to that presented in Eq. (3). A block diagram of the interpolation process is shown in Fig. 2b. Sinc subtraction corresponds to the choice of $g[n] = \text{sinc}[n/2]$ for the interpolation filter where $\text{sinc}[n] = \sin(\pi n)/(\pi n)$ as defined in Bracewell (1965).

In a manner analogous to the formation of perfusion estimates, estimates of the BOLD-weighted signal are formed from the sums of the tag and control images, with pair-wise addition having the form $y[1] + y[0]$, $y[1] + y[2]$, $y[3] + y[2]$, and surround addition having the form $y[1] + (y[0] + y[2])/2$, $(y[1] + y[3])/2 + y[2]$. The general form of these estimates is

$$\hat{b}[n] = y[n] * g[n], \tag{5}$$

which may also be written as the sum

$$\hat{b}[n] = \hat{y}_{con}[n] + \hat{y}_{tag}[n] \tag{6}$$

of the interpolated control and tag time series.

Additional interpolation of the estimates

The temporal resolution of the estimates discussed so far is equal to the repetition time (TR) of the image acquisition. For conventional ASL, the *TR* is typically 2 s or greater, while for turbo-ASL experiments the *TR* can be as short as 1 s (Wong et al., 2000). With the increasing popularity of event-related fMRI, it is not uncommon for the stimuli to be presented with a temporal resolution T_S that is finer than that of the image acquisition. For example, an event-related experiment might present stimuli on a 1-s grid, while acquiring images with a 2-s TR. In some cases, the experimenter may wish to compute an estimate at the temporal resolution of the stimuli. If *TR* is an integer multiple of T_S, the upsampling factor of 2 shown in the block diagram of Fig. 2b can be increased to $M = 2TR/T_s$ and the low pass filter $g[n]$ can be modified to attenuate the additional spectral images introduced by the higher upsampling rate (Oppenheim and Schafer, 1989). A generalized block diagram is shown in Fig. 2c, and a matrix equivalent was presented in Liu et al. (2002). An important special case is shown in Fig. 2d, where the perfusion time series from Fig. 2a is upsampled by a factor of $N = M/2$ and low pass filtered with $g_2[n]$. It is straightforward to show that the block diagram in Fig. 2d is equivalent to the generalized block diagram in Fig. 2c when $g[n] = \tilde{g}_1[n] * g_2[n]$, where $\tilde{g}_1[n]$ is equal to $g_1[n]$ upsampled by a factor of N. For example, if $g_1[n] = g_2[n] = [1\ 1]$ and $N = 2$, then $g[n] = [1\ 1\ 1\ 1]$. The criteria for the selection of the optimal overall filter are discussed in detail below. The first stage $g_1[n]$ of the filter has the greatest effect on BOLD contamination and the shape of the noise spectrum at low frequencies, while the second stage $g_2[n]$, which is used to attenuate spectral images due to upsampling, has relatively little effect on BOLD contamination, but does affect the shape of the noise spectrum at higher frequencies.

Applications

In the following sections, we show how the signal processing model can be used to (a) quantify BOLD contamination of the perfusion time series, (b) quantify perfusion weighting of the BOLD time series, (c) select an optimal filter to minimize BOLD contamination and perfusion weighting, and (d) characterize the noise in the perfusion time series. In most ASL fMRI experiments, simultaneous estimates of the perfusion and BOLD signals are of interest and can be obtained with Eqs. (3) and (5), respectively. BOLD contamination of the perfusion signal occurs when changes in the BOLD signal appear as changes in the perfusion estimate. Similarly, perfusion weighting of the BOLD signal occurs when changes in perfusion appear as changes in the BOLD signal. An understanding of the magnitude of these errors is critical for experiments that seek quantitative measures of perfusion and BOLD. For example, quantitative measures are important for furthering our understanding of the physiology and dynamics of functional perfusion and BOLD responses (Buxton et al., 1998; Hoge et al., 1999).

Knowledge of the temporal autocorrelation of noise in the perfusion estimate is important for proper statistical analysis of the experimental data (Burock and Dale, 2000). The analysis is greatly simplified if it can be assumed that the noise is uncorrelated. To complement the previously reported simulation and experimental results (Aguirre et al., 2002; Wang et al.,

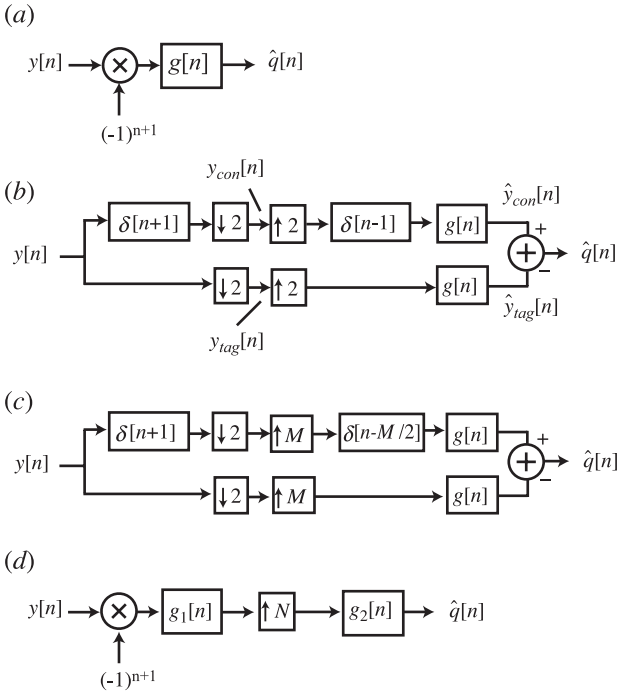

Fig. 2. (a) Block diagram of the perfusion estimate. (b) Equivalent form of the perfusion estimate (see Appendix for proof of equivalence). Downsampling and upsampling operators are denoted by ↓2 and ↑2, respectively. Advance and delay operators are denoted by the Kronecker delta functions $\delta[n + 1]$ and $\delta[n - 1]$, respectively. (c) Generalized form of the perfusion estimate to obtain increased output sampling rate. (d) Special case of the generalized form. This is equivalent to the generalized form when $N = M/2$ and $g[n] = \tilde{g}_1[n] * g_2[n]$, where $\tilde{g}_1[n]$ is obtained by upsampling $g_1[n]$ by a factor of N.

2003), we use the model to obtain an analytical expression for the autocorrelation of the perfusion noise.

BOLD contamination of perfusion time series

Images in an ASL fMRI experiment are typically acquired with a rapid single-shot imaging sequence such as an echo-planar imaging (EPI) or spiral readout sequence. The nonzero echo time (TE) and acquisition time of the readout sequence lead to BOLD weighting of the images. As we show below, this BOLD weighting results in two components of the BOLD contamination, one of which can be significantly reduced with the choice of an appropriate filter. In addition, BOLD weighting can be reduced by minimizing the echo time and, if possible, the acquisition time of the readout sequence. For example, with the use of a dual echo spiral sequence, perfusion estimates can be obtained from the first echo with minimum TE and BOLD estimates can be obtained from the second echo at a TE optimized for BOLD contrast (Liu et al., 2002; Yongbi et al., 2001). Another approach to reducing the effect of BOLD weighting is to reduce the static tissue component $M[n]$ using a background suppression technique (Alsop et al., 2000; Mani et al., 1997; Ye et al., 2000). This approach, however, reduces the sensitivity of the BOLD estimates. Nonsubtractive methods using background suppression have also been reported (Blamire and Styles, 2000; Duyn et al., 2001), but the quantification of the perfusion estimates from these methods is difficult.

In order to determine the degree of BOLD contamination of the perfusion estimate, it is useful to expand the perfusion estimate in Eq. (3) as the sum $\hat{q}[n] = q_q[n] + q_b[n] + q_e[n]$, where

$$q_q[n] = \left(\alpha b[n] q[n] e^{-TI/T_{1B}}\right) * g[n] \tag{7}$$

is a BOLD-weighted, low-pass-filtered perfusion component,

$$q_b[n] = \left[b[n]\left(s_M M[n] + s_q q[n]\right)\right](-1)^{n+1} * g[n] \tag{8}$$

is the sum of BOLD-weighted static tissue and perfusion components that are modulated and low pass filtered, and

$$q_e[n] = \left[(-1)^{n+1} e[n]\right] * g[n] \tag{9}$$

is the output noise component. For most perfusion fMRI experiments, it is reasonable to approximate the in-slice magnetization as a constant term $M[n] \approx M_0$. With this approximation, Eq. (8) may be further simplified by recalling that the BOLD weighting may be written as the sum $b[n] = b_0 + \Delta b[n]$ of a constant term and a time-varying term. The low pass filter will eliminate any constant terms that undergo modulation, yielding

$$q_b[n] \approx (-1)^{n+1}\left[s_M \Delta b[n] M_0 + b[n] s_q q[n]\right] * g[n] \tag{10}$$

Eqs. (7) and (10) show that there are two components of the BOLD contamination of the perfusion time series. The first component appears in the BOLD weighting of the $q_q[n]$ term in Eq. (7). This BOLD weighting is inherent in the measurement and cannot be reduced by filtering. However, it can be minimized by the use of acquisition schemes, such as the short echo-time single shot spiral acquisition discussed above. The second BOLD contamination component is the modulated BOLD-weighted static tissue term $(-1)^{n+1} s_M \Delta b[n] M_0$ in Eq. (10). This is a spurious component that can be attenuated by the low pass filter $g[n]$. In addition, the modulated and BOLD-weighted perfusion term $(-1)^{n+1} b[n] s_q q[n]$ in Eq. (10) forms a second spurious component. A key goal of the low pass filter $g[n]$ is to minimize both these spurious components.

To gain a sense of the filter attenuation that is required, it is useful to consider the magnitudes of the spurious components prior to filtering. These are defined as $M_b = |s_M \Delta b[n] M_0|$ and $M_{qm} = |s_q b[n] q[n]|$ for the BOLD and perfusion spurious components, respectively. Normalizing these terms by the magnitude $M_q = |\alpha b[n] q[n] e^{-TI/T_{1B}}|$ of the unfiltered perfusion component in Eq. (7) yields the relative magnitudes

$$M_b/M_q \approx s_M \exp(TI/T_{1B}) M_0 TE |\Delta R_2^*[n]| / (\alpha |q[n]|) \tag{11}$$

$$M_{qm}/M_q = s_q \exp(TI/T_{1B})/\alpha, \tag{12}$$

where the first order approximation $\Delta b[n]/b[n] \approx -\Delta R_2^*[n] TE$ has been used. With typical experimental parameters of $\alpha = \beta = 1$, $TI = TI_p = 1400$ ms, and representative physiological parameters $T_{1B} = 1300$ ms, $T_1 = 1000$ ms, a cerebral blood flow (CBF) of 60 ml/(100 g min) corresponding to $q[n] \approx 0.01 M_0$, and a 1% BOLD change corresponding to $TE|\Delta R_2^*[n]| = 0.01$ (i.e., ΔBOLD (%) $= 100|\Delta R_2^*[n]|TE$), the relative magnitudes are $M_b/M_q = 2.1$ and $M_{qm}/M_q = 1.9$. Thus, with these assumed parameters, the relative magnitudes of the spurious components are almost the same. An upper bound on the relative magnitude of the sum of spurious components can be obtained from the sum of the relative magnitudes, with equality of the sums if the spurious components have the same sign. For the example given, the sum of relative magnitudes is 4. This means that if the spurious components are attenuated to 1% of their value by the low pass filter, the filtered sum of spurious components will be roughly 4% of the desired perfusion component.

As evident from Eq. (11), the relative magnitude of the BOLD spurious component M_b/M_q can be reduced either by using a shorter echo time TE or by using background suppression to reduce M_0. In addition, for a given TE, the relative magnitude depends on spatial and temporal variations of $|\Delta R_2^*[n]|$. Relative magnitudes versus percent BOLD changes ranging from 0.1% to 10% are shown in Fig. 3 using the experimental parameters described above and assuming either no background suppression (solid lines) or 90% background suppression (dashed lines). Also shown are the term M_{qm}/M_q and the sum of the spurious magnitudes for both background suppression conditions. The magnitude M_{qm}/M_q of the perfusion spurious component does not vary with either M_0 or the percent BOLD change, and therefore serves as a lower bound on the sum of spurious components. As a result, when M_b/M_q becomes less than M_{qm}/M_q, the use of shorter echo times and background suppression has a smaller effect on the overall sum of spurious components. For example, the curves in Fig. 3 show that for BOLD percent changes less than 0.6%, the use of 90% background suppression reduces the overall unfiltered spurious components by less than a factor of 2.

Perfusion weighting of BOLD time series

In examining the perfusion weighting of the BOLD time series, we assume the use of a presaturation pulse since this has been shown to reduce the degree of perfusion weighting (Wong et al.,

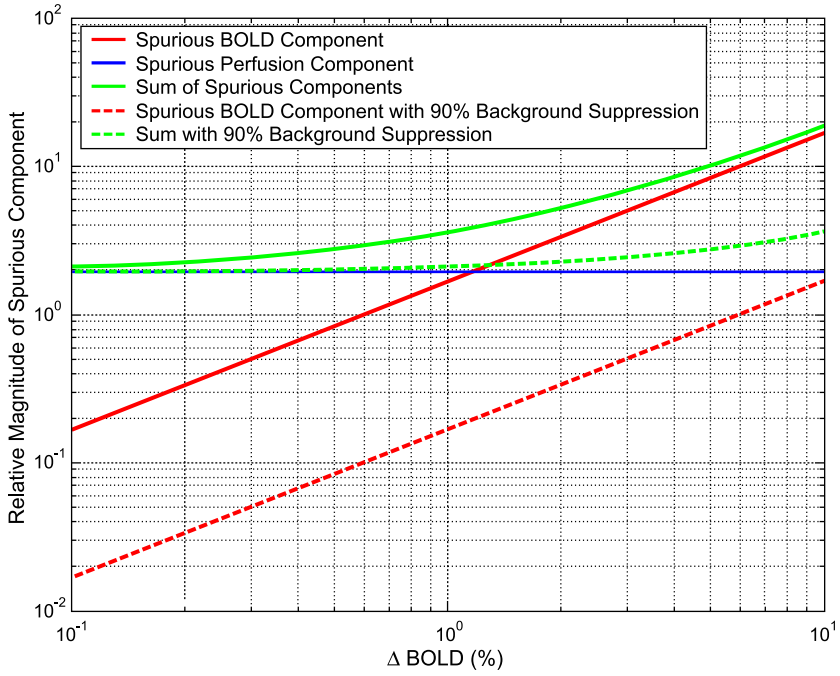

Fig. 3. Relative magnitudes of unfiltered spurious components in the perfusion estimate versus percent change in BOLD. Assumed experimental parameters are described in the text.

1997). The expression for the BOLD estimate in Eq. (5) may be rewritten as the sum $\hat{b}[n] = b_b[n] + b_q[n] + b_e[n]$ where

$$b_b[n] = \left[b[n]\left(s_M M[n] + s_q q[n]\right)\right] * g[n] \tag{13}$$

is the sum of BOLD-weighted static tissue and perfusion terms that are low pass filtered,

$$b_q[n] = \left((-1)^{n+1}\alpha b[n] q[n] e^{-TI/T_{1B}}\right) * g[n] \tag{14}$$

is the modulated and low-pass-filtered perfusion component, and

$$b_e[n] = e[n] * g[n] \tag{15}$$

is the output noise term. Here, $b_q[n]$ is the spurious perfusion-weighted term that is attenuated by the low pass filter. Because the terms $b_b[n]$ and $b_q[n]$ differ from $q_b[n]$ and $q_q[n]$ only in the placement of the modulation term $(-1)^{n+1}$, the selection criteria for choosing a filter to reduce the spurious term $b_q[n]$ are similar to those used for reducing BOLD contamination of the perfusion signal. If the BOLD and perfusion signals have the same functional form, then the same filter minimizes both perfusion weighting and BOLD contamination.

In addition to the perfusion weighting due to $b_q[n]$, there is perfusion weighting of the $b_b[n]$ term that is not affected by the filtering process. It is useful to rewrite this term as

$$b_b[n] = \left[s_M b[n]\left((M[n] + q[n]) + \left(s_q - s_M\right)q[n]/s_M\right)\right] * g[n]. \tag{16}$$

The term $M[n] + q[n]$ is the sum of the static magnetization and the magnetization due to inflowing spins, and is constant if the magnetization densities of static tissue and flowing blood are the same. However, if blood volume increases with functional activation and the density of blood is higher than that of static tissue, this term will tend to increase with activation. The term $(s_q - s_M)q[n]/s_M = (\beta e^{-TI_p/T_1} - \alpha e^{-TI/T_{1B}})q[n]/(1 - \beta e^{-TI_p/T_1})$ reflects perfusion weighting due to the difference in the longitudinal time constants of blood and static tissue. With typical parameters of $\alpha = \beta = 1$, $TI_p = TI = 1400$ ms, $T_{1B} = 1300$ ms, and $T_1 = 1000$ ms, this term is $-0.19\ q[n]$, which is similar to the effect calculated in Wong et al. (1997). If we assume a resting $q[n] \approx 0.01\ M_0$ and a 100% increase in CBF, this translates into a 0.19% increase in the BOLD signal without any true BOLD change. The residual perfusion weighting term can be reduced by decreasing TI_p to compensate for the shorter time constant of static tissue. However, the effect of the delayed presaturation pulses on the bolus of flowing spins that have entered the gap between the tagging region and imaging slice can introduce errors into the quantitation of CBF. In addition, the degree to which the perfusion weighting is reduced requires knowledge of the time at which tagged blood water exchanges with tissue water and begins to relax with the time constant of tissue. Note that if presaturation is not used ($\beta \neq 1$), the perfusion weighting of the BOLD signal can be much larger and can be either positive or negative depending on β. Another approach to the reduction of perfusion weighting is the use of a multi-echo sequence to directly measure changes in the transverse relaxation rate.

Selection of an optimal filter

As discussed in the previous section, the criteria for selecting an optimal filter for minimizing either BOLD contamination or perfusion weighting are similar. In this section, we focus on the selection of an optimal filter for the reduction of the spurious components in the perfusion estimate. The analysis is most easily

presented in the frequency domain. The discrete-time Fourier transforms of the perfusion and spurious components are

$$Q_q(f) = \alpha e^{-TI/T_{1B}} G(f)[B(f) * Q(f)] \quad (17)$$

and

$$Q_b(f) \approx -G(f)\left(s_M M_0 \Delta B(f+0.5) + s_q B(f) * Q(f+0.5)\right) \quad (18)$$

where f denotes normalized frequency, that is, f is the frequency divided by the sampling frequency (Oppenheim and Schafer, 1989). An optimal filter minimally attenuates the BOLD-weighted perfusion spectrum $B(f) * Q(f)$ while maximally attenuating the modulated spurious components. To simplify the analysis, we assume that the percent increase in perfusion is much greater than the percent BOLD increase, which leads to the approximations $b[n]q[n] \approx b_0 q[n]$, $B(f) * Q(f) \approx b_0 Q(f)$, and $B(f) * Q(f+0.5) \approx b_0 Q(f+0.5)$. This is a reasonable assumption for voxels with functional perfusion increases, where typical percent increases range from 50% to 100% as compared to BOLD percent changes of 1–5% (Wong et al., 1997). With this assumption and the expansion of the perfusion signal as the sum $q[n] = q_0 + \Delta q[n]$ of a baseline perfusion term q_0 and a dynamic perfusion term $\Delta q[n]$, we may rewrite Eq. (18) as

$$Q_b(f) \approx -G(f)\left(s_M M_0 \Delta B(f+0.5) + s_q b_0 \Delta Q(f+0.5)\right), \quad (19)$$

where the q_0 term is eliminated by the zeros of $G(f)$ at $f = \pm 0.5$. With the approximation that the dynamic BOLD and perfusion spectra are described by the same function ($\Delta B(f) = \Delta Q(f)$), the spurious term defined in Eq. (19) is proportional to $G(f)\Delta B(f + 0.5)$. The dynamic portion of the perfusion term of Eq. (17) is proportional to $G(f)\Delta Q(f)$. Comparison of the peak amplitude of $G(f)\Delta B(f + 0.5)$ with that of $G(f)\Delta Q(f)$ provides the relative attenuation of the spurious components by the filter. Combining the filter attenuation with the estimates of the relative magnitudes of the unfiltered spurious components then yields estimates of the magnitudes of the spurious components.

As examples of filter selection, we consider the spectra associated with two widely used experimental designs for fMRI. These are a block design and a randomized event-related design. The response to the block design is obtained by convolving a stimulus consisting of 4 cycles of 30 s on/off with a canonical hemodynamic response (HDR) modeled as a gamma density function of the form $h(t) = (\tau n!)^{-1}(t/\tau)^n \exp(-t/\tau)$ with $\tau = 1.2$, $n = 3$, and $\Delta t = 1$. Fig. 4a shows the normalized spectra $\Delta Q(f)$ and $\Delta B(f + 0.5)$ for the block design assuming a TR of 2 s. Also shown are the spectra $G(f)$ for $g[n] = [1\ 1]$, $g[n] = [1\ 2\ 1]/2$, and $g[n] = \mathrm{sinc}[n/2]$. Fig. 4b shows the spectra after filtering with $g[n] = [1\ 1]$ and $g[n] = [1\ 2\ 1]/2$. Consistent with the findings of Aguirre et al. (2002), it is clear that the sinc filter provides minimal attenuation of $\Delta Q(f)$ and maximal attenuation of $\Delta B(f + 0.5)$. The relative performance of the other filters can be measured by comparing the amplitude of the perfusion component at the normalized fundamental frequency $f = 1/30$ to the amplitude of the spurious component at the modulated fundamental frequency $f = 0.5 - 1/30$. This is given by the ratio $|G(0.5-1/30)|/|G(1/30)|$. The relative gains of the filters at the spurious frequency are 0.10 and 0.01 for $g[n] = [1\ 1]$ and $g[n] = [1\ 2\ 1]/2$, respectively. Combining these relative gains with previous estimates for the overall relative magnitudes of spurious components assuming a 1% BOLD change yields filtered spurious component relative magni-

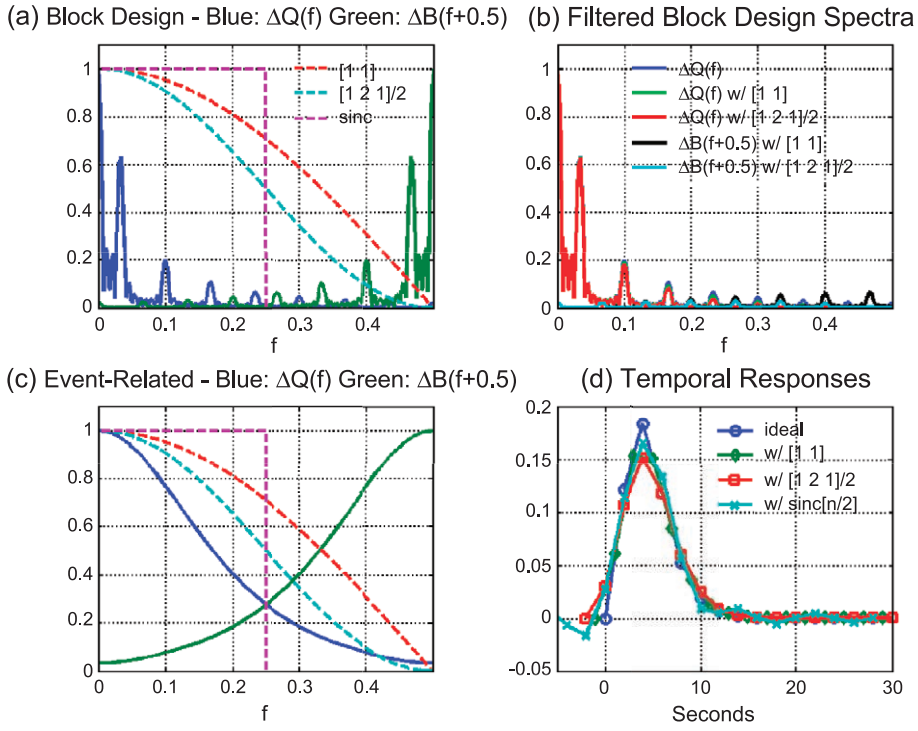

Fig. 4. (a) Perfusion spectrum $\Delta Q(f)$ and modulated BOLD spectrum $\Delta B(f + 0.5)$ for a block design stimulus convolved with a canonical hemodynamic response (HDR) function. Filter responses are also shown. (b) Perfusion and modulated BOLD spectra after filtering. (c) Perfusion and modulated BOLD spectra for a randomized event-related design. (d) Ideal HDR and filtered HDR functions.

tudes of 40% and 4%, respectively. The performance of the filters is consistent with the findings of Aguirre et al. (2002) who found that BOLD contamination with surround subtraction was much less than obtained with pairwise subtraction and slightly higher than the contamination observed with sinc subtraction.

In randomized event-related designs, the stimulus spectrum is nearly white, so that $\Delta Q(f)$ and $\Delta B(f)$ can be approximated by spectrum of the HDR. Fig. 4c shows the spectra $\Delta Q(f)$ and $\Delta B(f+0.5)$ and the filter responses. Because of the broad bandwidth of the HDR spectra, there is a trade-off between attenuation of the spurious component and preservation of the desired perfusion spectrum. The sinc filter not only maximally attenuates the spurious component for normalized frequencies above 0.25, but also significantly reduces the bandwidth of the desired perfusion spectrum. In contrast, the filter $g[n] = [1\ 1]$ has the largest bandwidth and thus best preserves the bandwidth of the perfusion spectrum, but provides the least attenuation of the spurious component. To gauge the relative performance of the various filters, it is useful to return to the temporal domain and consider the perfusion estimate when the perfusion signal $q[n] = h[n]$ is equal to the sampled HDR function $h[n] = h(2n)$. When the tag image coincides with the first sample of the HDR, the estimate has the form

$$\hat{q}_1[n] \approx \left(h[n] + (-1)^{n+1} h_b[n]\right) * g[n] \tag{20}$$

where $h_b[n]$ is the sampled BOLD HDR. If the perfusion signal is delayed by one time sample, then the control image coincides with the first HDR sample, and the estimate has the form

$$\hat{q}_2[n] \approx \left(h[n-1] + (-1)^{n+1} h_b[n-1]\right) * g[n]. \tag{21}$$

Because of the randomized stimulus presentation, these two estimates occur with roughly equal frequency, so that an overall estimate may be formed by advancing $\hat{q}_2[n]$ by one time sample and averaging with $\hat{q}_1[n]$ to obtain

$$\hat{q}[n] \approx (\hat{q}_1[n] + \hat{q}_2[n] * \delta[n+1])/2 = h[n] * g[n]. \tag{22}$$

Thus, the randomization tends to eliminate the spurious component, yielding an estimate that is the HDR filtered by $g[n]$. A similar argument would apply to a periodic single trial experiment in which the start times of half of the trials coincide with a tag image, with the remaining start times coinciding with control images. Fig. 4d shows the canonical HDR after convolution with the various filters. Temporal broadening is minimized for the $g[n] = [1\ 1]$ filter, which has the largest bandwidth. This is consistent with previous findings reported in Liu et al. (2002).

Perfusion noise characteristics

As stated in the Eq. (9), the perfusion noise term is $q_e[n] = [(-1)^{n+1} e[n]] * g[n]$. To derive the autocorrelation of the noise, we first note that the autocorrelation function of $(-1)^{n+1} e[n]$ is $(-1)^n \rho[n]$ where $\rho[n]$ is the autocorrelation function for $e[n]$. The autocorrelation of the perfusion noise is then given by

$$\rho_q[n] = (-1)^n \rho[n] * g[n] * g[-n] \tag{23}$$

and the power spectrum is

$$\hat{S}_q(f) = S(f+0.5)|G(f)|^2 \tag{24}$$

where $S(f)$ is the Fourier transform of $\rho[n]$ (Oppenheim and Schafer, 1989). To examine the effect of filtering on the perfusion noise, we use the autocorrelation model presented in Burock and Dale (2000) as an example of $1/f$ fMRI noise. In this model, the noise is the weighted sum of a white noise process and a first order autoregressive process and has autocorrelation $\rho[n] = \sigma^2 (\lambda\delta[n] + (1-\lambda)a^{|n|})$ with $a = 0.88$ and $\lambda = 0.75$. Fig. 5a shows $S(f)$ and $S(f+0.5)$ and $|G(f)|^2$ for the various low pass filters, and Fig. 5b shows the filtered spectra. Note that the modulated spectrum $S(f+0.5)$ is fairly flat out to a normalized frequency of 0.35. Consistent with the experimental findings of Wang et al. (2003), sinc subtraction provides the flattest noise spectrum for low frequencies ($f < 0.25$, corresponding to frequencies below 0.125 Hz), with the pair-wise and surround subtraction providing progressively less flat spectra. However, the sinc subtraction filter has a sharp cut-off at $f = 0.25$, so that the output noise spectrum is not flat over the entire frequency range. Although maximizing the flatness of the spectrum for low frequencies may be a reasonable criteria for block designs in which most of the energy is concentrated around a low frequency fundamental, a more generally valid criterion is to maximize the flatness over the entire frequency range, or equivalently minimize the width of the autocorrelation function. The normalized autocorrelation functions are shown in Fig. 5c. The width is minimized for the pair-wise subtraction filter, which is also the filter that appears to provide the flattest spectra over the entire frequency range. To gain insight into the form of the autocorrelation functions, we note that for the noise model in our example the modulation and filtering process greatly attenuate the contribution of the autoregressive term, so that the autocorrelation can be approximated as $\rho_q[n] \approx \sigma^2 \lambda g[n] * g[-n]$. The convolution products $g[n] * g[-n]$ are equal to [1 2 1], [1 4 6 4 1]/4, and sinc $[n/2]$ for pairwise, surround, and sinc subtraction, respectively. These products are shown in normalized form in Fig. 5d and approximate the functions in Fig. 5c very well.

Discussion

We have shown that the various perfusion estimates based upon subtraction processing can be represented by the general form presented in Eq. (3), consisting of a modulation operation followed by a low pass filtering operation. The subtraction methods differ only in the selection of the low pass filter. The performance of the methods can be evaluated using the signal processing model presented in Fig. 1. The perfusion estimate is the sum of a low pass filtered, BOLD-weighted perfusion component and low pass filtered spurious components (containing modulated BOLD and perfusion terms). For block design experiments, sinc subtraction provides the greatest attenuation of the spurious components, although the attenuation provided by surround subtraction is probably sufficient for most experiments. For randomized event-related experiments, the random occurrence of stimuli tends to cancel out the spurious components, so that the filter with the largest bandwidth, corresponding to pair-wise subtraction, should be used to preserve the spectrum of the perfusion component. The same conclusion holds for a periodic single trial design with appropriate stimulus shifting such that half of the trial start times coincide with tag images. BOLD weighting of the desired perfusion term and the BOLD spurious component can be minimized with shorter echo times or background suppression. However, these methods do not attenuate the spurious perfusion component, which is roughly the same size as the desired perfusion term prior to filtering. Thus, even in experiments with minimal BOLD weighting, either sinc subtraction or surround subtraction

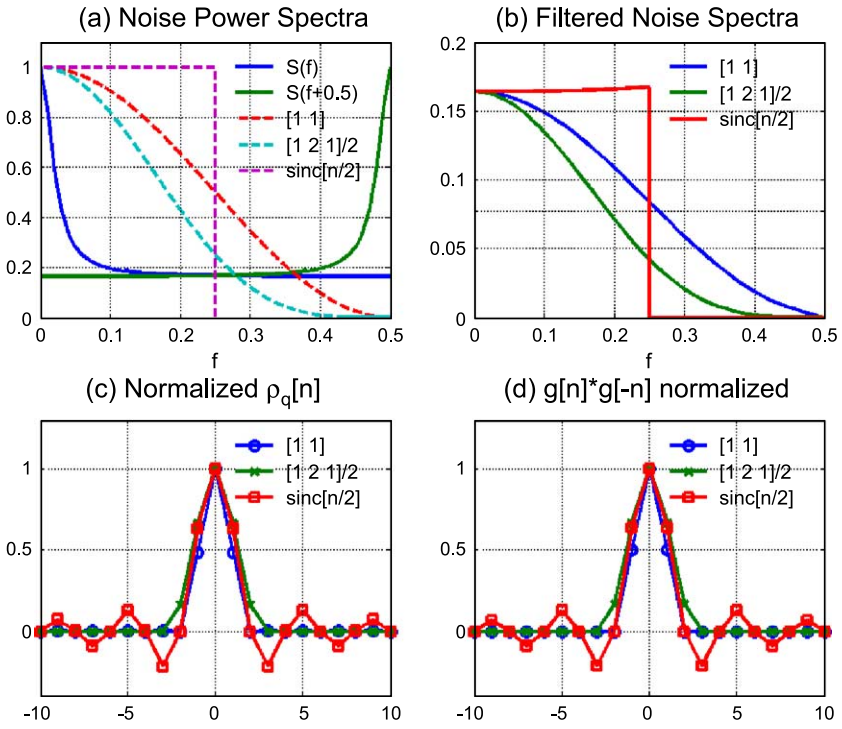

Fig. 5. (a) Noise power spectrum $S(f)$, modulated noise power spectrum $S(f + 0.5)$, and low pass filter responses. (b) Filtered noise power spectra. (c) Normalized perfusion noise autocorrelation functions. (d) Convolution of low pass filters with their mirror images.

should still be used for block designs, unless stimulus shifting methods are applicable (e.g., for estimating an average block response).

The ASL subtraction process has the advantage of whitening the $1/f$ noise typically observed in fMRI experiments. Sinc subtraction provides the flattest noise spectrum for low frequencies. However, when the entire experimental frequency range is considered, both pair-wise and surround subtraction yield narrower autocorrelation functions.

Our analysis has not explicitly modeled the low frequency drift terms that are often observed in fMRI scans (Friston et al., 1995) and contribute to the apparent $1/f$ type noise. These terms are often treated as additional regressors in the statistical analysis of fMRI experiments (Liu et al., 2002). For the perfusion estimates, these terms will be modulated and filtered in the same manner as the noise (e.g., Eq. (9)). The selection of an optimal filter for reducing the low frequency confounds depends on assumptions about the bandwidth of the low frequency confounds. In most block design experiments, however, the bandwidth of drift terms will be less than the fundamental frequency of the stimulus, so that a filter optimized to reduce BOLD contamination also minimizes the confounds.

When taken together, the BOLD contamination and noise whitening results suggest that either sinc or surround subtraction is preferable for block design experiments, while pair-wise subtraction is preferable for randomized event-related experiments and periodic single trial designs with appropriate stimulus shifting. These conclusions are consistent with the results of previous experimental and simulation studies (Aguirre et al., 2002; Liu et al., 2002; Wang et al., 2003).

In general, the performance of the estimate is optimized by appropriate selection of the low pass filter, and the optimal filter may not correspond to any of the three methods considered in this paper. The selection of an optimal filter depends on the specifics of the experiment and the performance criteria established by the investigator. As an example, if the experiment corresponds to the block design described in this paper, and the design criterion is to achieve the flattest frequency response while nulling out the spurious component at $f = 0.5 - 1/30$, then the filter $g[n] = \mathrm{sinc}[0.9n]$ would outperform both the sinc and surround subtraction filters. Other filters, such as windowed sinc filters, may be designed using standard digital filter design methods (Oppenheim and Schafer, 1989).

It is important to note that all of the subtraction methods result in a filtered estimate of the actual perfusion. In event-related experiments, this can lead to significant broadening in estimates of the perfusion HDR. In these cases, a direct matrix deconvolution may be used to obtain an unfiltered HDR estimate (Liu et al., 2002).

In conclusion, we have presented a theoretical model of the ASL signal processing chain and have shown how the model provides a general framework for analyzing the performance of various processing methods. In showing the applications of the model we have made a number of assumptions regarding the ASL sequence, such as the use of a presaturation pulse. The modification of these assumptions for the analysis of other ASL sequences is readily accomplished within the framework of the model.

Appendix A

Here we show that estimate given in Eq. (4) is equivalent to that in Eq. (3). The downsampling operator is defined as $y[n] \downarrow 2 = y[2n]$ (Oppenheim and Schafer, 1989). The tag and control time series are then $y_{con}[n] = y[n + 1] \downarrow 2 = y[2n + 1]$ and $y_{tag}[n] = y[n] \downarrow 2 = y[2n]$. Invoking linearity, the filter in Fig. 2b may

be placed after the subtraction operation yielding $\hat{q}[n] = g[n] * ((y_{con}[n] \uparrow 2) * \delta[n-1] - y_{tag}[n] \uparrow 2)$, so it is sufficient to show that the term in parentheses is equal to $(-1)^{n+1}y[n]$. Substituting for the tag and control terms yields $((y[2n+1] \uparrow 2) * \delta[n-1] - y[2n] \uparrow 2)$. The upsampling operator is defined as $y[n] \uparrow 2 = y[n/2]$ for $n/2$ equals to an integer and zero otherwise (Oppenheim and Schafer, 1989). Thus, the term $y[2n] \uparrow 2$ is equal to $y[n]$ for n even and zero otherwise, and the term $(y[2n+1] \uparrow 2) * \delta[n-1]$ is equal to $y[n]$ for n odd and zero otherwise. The difference of these two terms is equal to $(-1)^{n+1}y[n]$.

References

Aguirre, G.K., Detre, J.A., Zarahn, E., Alsop, D.C., 2002. Experimental design and the relative sensitivity of BOLD and perfusion fMRI. NeuroImage 15, 488–500.

Alsop, D.C., Maldjian, J.A., Detre, J.A., 2000. In-vivo MR perfusion imaging of the human retina. Proceedings of the ISMRM 8th Scientific Meeting, Denver, pp. 162.

Blamire, A.M., Styles, P., 2000. Spin echo entrapped perfusion image (SEEPAGE). A nonsubtraction method for direct imaging of perfusion. Magn. Reson. Med. 43 (5), 701–704.

Bracewell, R.N., 1965. The Fourier Transform and its Applications. McGraw-Hill, New York.

Burock, M.A., Dale, A.M., 2000. Estimation and detection of event-related fMRI signals with temporally correlated noise: a statistically efficient and unbiased approach. Hum. Brain Mapp. 11, 249–260.

Buxton, R.B., Wong, E.C., Frank, L.R., 1998. Dynamics of blood flow and oxygenation changes during brain activation: the balloon model. Magn. Reson. Med. 39, 855–864.

Duhamel, G., Bazelaire, C.d., Alsop, D.C., 2003. Evaluation of systematic quantification errors in velocity-selective arterial spin labeling of the brain. Magn. Reson. Med. 50 (1), 145–153.

Duong, T.Q., Kim, D.S., Ugurbil, K., Kim, S.G., 2001. Localized cerebral blood flow response at submillimeter columnar resolution. Proc. Natl. Acad. Sci. U. S. A. 98, 10904–10909.

Duyn, J.H., Tan, C.X., van Gelderen, P., Yongbi, M.N., 2001. High-sensitivity single-shot perfusion-weighted fMRI. Magn. Reson. Med. 46 (1), 88–94.

Friston, K.J., Frith, C.D., Turner, R., Frackowiak, R.S.J., 1995. Characterizing evoked hemodynamics with fMRI. NeuroImage 2, 157–165.

Golay, X., Hendrikse, J., Lim, T.C., 2004. Perfusion imaging using arterial spin labeling. Top. Magn. Reson. Imaging 15 (1), 10–27.

Hoge, R.D., Atkinson, J., Gill, B., Crelier, G.R., marrett, S., Pike, G.B., 1999. Stimulus-dependent BOLD and perfusion dynamics in human V1. NeuroImage 9, 573–585.

Liu, T.T., Wong, E.C., 2004. A signal processing model for arterial spin labeling perfusion fMRI. Proceedings of the ISMRM 12th Scientific Meeting, Kyoto, pp. 368.

Liu, T.T., Wong, E.C., Frank, L.R., Buxton, R.B., 2002. Analysis and design of perfusion-based event-related fMRI experiments. NeuroImage 16 (1), 269–282.

Luh, W.M., Wong, E.C., Bandettini, P.A., Ward, B.D., Hyde, J.S., 2000. Comparison of simultaneously measured perfusion and BOLD signal increases during brain activation with T(1)-based tissue identification. Magn. Reson. Med. 44 (1), 137–143.

Mani, S., Pauly, J., Conolly, S., Meyer, C., Nishimura, D., 1997. Background suppression with multiple inversion recovery nulling: applications to projective angiography. Magn. Reson. Med. 37 (6), 898–905.

Oppenheim, A.V., Schafer, R.W., 1989. Discrete-Time Signal Processing. Prentice-Hall, Englewood Cliffs, NJ.

Wang, J., Aguirre, G.K., Kimberg, D.Y., Detre, J.A., 2003. Empirical analyses of null-hypothesis perfusion fMRI data at 1.5 and 4T. NeuroImage 19, 1449–1462.

Wong, E.C., Buxton, R.B., Frank, L.R., 1997. Implementation of quantitative perfusion imaging techniques for functional brain mapping using pulsed arterial spin labeling. NMR Biomed. 10, 237–249.

Wong, E.C., Luh, W.-M., Liu, T.T., 2000. Turbo ASL: arterial spin labeling with higher SNR and temporal resolution. Magn. Reson. Med. 44, 511–515.

Ye, F.Q., Frank, J.A., Weinberger, D.R., McLaughlin, A.C., 2000. Noise reduction in 3D perfusion imaging by attenuating the static signal in arterial spin tagging (ASSIST). Magn. Reson. Med. 44 (1), 92–100.

Yongbi, M.N., Fera, F., Mattay, V.S., Frank, J.A., Duyn, J.H., 2001. Simultaneous BOLD/perfusion measurement using dual-echo FAIR and UNFAIR: sequence comparison at 1.5T and 3.0T. Magn. Reson. Imaging 19 (9), 1159–1165.