Magnetic Resonance in Medicine 44:746–757 (2000)

Wavelet Transform-Based Wiener Filtering of Event-Related fMRI Data

Stephen M. LaConte, Shing-Chung Ngan, and Xiaoping Hu*

The advent of event-related functional magnetic resonance imaging (fMRI) has resulted in many exciting studies that have exploited its unique capability. However, the utility of event-related fMRI is still limited by several technical difficulties. One significant limitation in event-related fMRI is the low signal-to-noise ratio (SNR). In this work, a method based on Wiener filtering in the wavelet domain is developed and demonstrated for denoising event-related fMRI data. Application of the technique to simulated and experimental data demonstrates that the technique is effective in reducing noise while preserving neuronal activity-induced response. Magn Reson Med 44: 746–757, 2000. © 2000 Wiley-Liss, Inc.

Key words: event-related fMRI; denoising; stationary wavelet transform; Wiener filter

In the past several years functional MRI (fMRI) (1–3) based on the blood oxygen level-dependent (BOLD) (4–7) contrast has become a ubiquitously used technique for studying brain function. Analogous to PET studies, most fMRI studies have utilized "block design" experiments in which a stimulus is presented for extended periods of time. In this design, pixel intensities in images acquired during a period of stimulus (and presumably response) are compared to those from periods with no stimulus. This approach essentially provides a steady-state view of the neuronal response at an acceptable signal-to-noise ratio (SNR). Recent advances in both image acquisition and data analysis have made it possible to exploit the high temporal resolution available in MRI to perform what is termed "event-related" or "single trial" experiments (8). This new experimental design, termed ER-fMRI hereafter, closely parallels evoked response studies performed with high temporal resolution techniques such as evoked response potential (ERP) measurements. In these experiments, the subject is presented with a brief stimulus or performs a single instance of a task, and the corresponding transient response is measured. Compared to block design experiments, event-related experiments are more versatile because they can capture the temporal profile of the response, thereby potentially providing timing information. In addition, the event-related approach is also capable of capturing trial-to-trial variations that may arise from factors such as habituation, learning, arousal, and attention.

One obvious drawback of ER-fMRI is the loss of SNR due to the transient nature of the event-related response and the time needed between trials. This loss of SNR, however, has not impeded advances in ER-fMRI. Among these advances are the application of ER-fMRI to various high-order cognitive tasks, the investigation of the utility of epoch-dependent information, and the adaptation of methods developed for ERP to fMRI. Work reported by Buckner et al. (9) demonstrated that higher-order cognitive tasks such as activation of the prefrontal cortex during word-stem completion was possible despite the fact that such paradigms typically elicit much smaller signal changes than sensory stimulation paradigms. At high fields it was demonstrated that SNR is sufficient for detecting activation responses to individual trials without epoch averaging (10). Experiments with intermixed trial types such as infrequent events (11) or randomized go/no-go paradigms (12) have allowed the detection of responses to rare events, demonstrating the flexibility of ER-fMRI.

Currently, most ER-fMRI studies are performed using an approach that relies on inter-epoch averaging as first introduced by Buckner et al. (9). Despite the success of this approach, which is robust in the presence of a low SNR, the averaging ignores the information associated with each execution of the task and assumes that subject behavior and brain function do not vary during repeated trials. In fact, trial-to-trial variations are frequently what is of interest and can be correlated with task/performance variables to reveal relevant information. The possibility of exploring these time-dependent modulations is the greatest benefit ER studies have to offer. To date, however, true single trial experiments have been reported only at high fields due to SNR limitations (10). Therefore, an improvement in SNR is greatly needed to fully exploit the potential of ER-fMRI, and is thus the focus of the present work.

In developing techniques for ER-fMRI, methods introduced for ERP can be adapted. An example of such an adaptation is the work by Dale and Buckner (13) on the method of selective averaging. In the present work, which focuses on denoising, the method of time-varying Wiener filtering (14,15), along with the subsequent extension to the wavelet transform domain (16) has been adapted for processing ER-fMRI data. This technique is mainly aimed at improving the SNR in the activated time courses before they are characterized for aspects such as onset, duration, and amplitude of the BOLD response. In adapting these approaches for processing ER-fMRI data, several modifications were made. The most important modification is that in order to preserve the epoch-to-epoch variation, the filter is applied to individual epoch data. In addition, filtering is done under the framework of the translation-invariant wavelet transform (also known as the stationary wavelet transform), rather than the conventional discrete wavelet transform, in order to enhance the denoising performance.

The main purpose of this work is to develop a denoising technique for individual time courses that have already

University of Minnesota, Center for Magnetic Resonance Research, Minneapolis, Minnesota.

Grant sponsor: NIH; Grant numbers: RO1MH55346; RR08079; R03MH59245.

*Correspondence to: Xiaoping Hu, Ph.D., University of Minnesota Medical School, Center for Magnetic Resonance Research, 2021 6th Street, S.E., Minneapolis, MN 55455. E-mail: xiaoping@cmrr.umn.edu

Received 19 July 1999; revised 18 April 2000; accepted 23 June 2000.

been identified as "active" through some other method such as correlation analysis. Since the event-related design can provide the temporal information regarding the brain's response to events, enhancing the SNR of individual evoked responses will increase the accuracy of their temporal characterization, thereby improving the quantification of the commonalities and differences that may occur across repeated trials. In addition, using the filtering method as a preprocessing step to model-free algorithms, such as data clustering (17–19), may lead to improved activation detection. It should be noted that a model-free approach is more appropriate in this situation since the noise structure of the filtered data is complicated, making statistical analysis using conventional methods difficult.

THEORY AND ALGORITHM

In this section the concepts of Wiener filtering and its time-variant version are first briefly reviewed. Subsequently, a description of the discrete wavelet transform along with its stationary version is given. Finally, the Wiener filter using the stationary wavelet transform is described.

Wiener Filter

The Wiener filter is the solution to the linear minimum square error problem of estimating a signal, $x[n]$, from a measurement consisting of signal plus noise, $y[n] = x[n] + v[n]$, using a filter, $h[n]$. This estimation can be expressed as:

$$\hat{x}[n] = h[n]^* y[n], \qquad [1]$$

which is the convolution of the measured signal with a time invariant filter. Here $h[n]$, determined by minimizing the estimation error, defines the Wiener filter. It can be shown (20) that the frequency response of the noncausal Wiener filter is:

$$H(\omega) = \frac{P_x(\omega)}{P_x(\omega) + P_v(\omega)} \qquad [2]$$

where $P_x(\omega)$ and $P_v(\omega)$ represent the power spectral density of the true signal and the noise, respectively. Intuitively, the action of Eq. [2] is to keep frequency bands where the signal power is much stronger than that of the noise and to severely attenuate frequencies where noise predominates.

In deriving the time-invariant Wiener filter, both x and v are assumed to be stationary. For many applications, the assumption of stationary signal and noise characteristics is not completely satisfied. A natural extension of the Wiener filter is to allow it to be time-varying. In principle, this should only require substituting the signal and noise power spectral densities with their time-dependent counterparts. In the context of ERP analysis, de Weerd and Kap (14,15) introduced and implemented a time-varying Wiener filter. An alternative is to perform the filtering in the wavelet domain, where localization is achieved in both the frequency and time domains (16,21).

Stationary Wavelet Transform

The present work relies on the stationary wavelet transform (SWT) (22–24), which has the advantage of being shift-invariant, albeit computationally more expensive than standard implementations of the discrete wavelet transform (DWT). Detailed description of the SWT can be found elsewhere (23). For completeness, its essence is summarized in this subsection.

Since the definition of the SWT follows from that of the DWT, the definition of the DWT is first described. Given a time series, $T = \{x[0], x[1], \ldots, x[N-1]\}$ with $N = 2^J$ for some integer J, the DWT is performed based on the application of two filters: a low pass filter H and a high pass filter G, both in conjunction with a binary decimation operator D_0. The filter H is defined by a sequence $\{h_n\}_{n\epsilon Z}$, where h_n's satisfy the orthogonality condition $\Sigma_n h_n h_{n+2j} = 0$ for any integer j. The filter G is defined as $\{g_n\}_{n\epsilon Z}$ with $g_n = (-1)^n h_{1-n}$. Applying H on the time series T results in $(Hx)_k = \Sigma_n h_{n-k} x[n]$, while applying D_o on T produces $(D_0 x)_j = x[2j]$. Now, define the smooth component of the given time series at resolution level J as the original time series itself:

$$s^J = \{s^J[n]\}, \text{ with } s^J[n] = x[n] \quad \text{for } n = 0, 1, \ldots, 2^J - 1 \qquad [3]$$

and the smooth and the detail components of the time series at resolution level j as:

$$s^j = D_0 H s^{j+1} \text{ and } d^j = D_0 G s^{j+1} \qquad [4]$$

where s^j and d^j are sequences of 2^j terms. The DWT of the time series from level J to some level $L < J$ is given by 2^J terms:

$$Q = \{d^{J-1}, d^{J-2}, \ldots, d^L, s^L\} \qquad [5]$$

which provide a complete representation of the given time series.

It can be shown (23,24) that the binary decimation operator D_o causes shift-variance and hence leads to problems concerning stability. To remedy the situation, the high and low pass filters are applied to the given time series without decimation in the SWT. To describe this mathematically, a new operator Z is defined such that $(Zx)_{2j} = x[j]$ and $(Zx)_{2j+1} = 0$, and $H^{[r]}$ and $G^{[r]}$ are given by $\{Z^r h\}$ and $\{Z^r g\}$, where Z^r denotes the sequential application of the operator Z r times. To perform the SWT, set:

$$s^J[n] = x[n] \text{ for } n = 0, 1, \ldots, 2^J - 1 \qquad [6]$$

and define the smooth and detail components of the time series at level j as:

$$s^{j-1} = H^{[J-j]} s^j \text{ and } d^{j-1} = G^{[J-j]} s^j \qquad [7]$$

where s^{j-1} and d^{j-1} are now sequences of 2^J terms. The resultant $(J - L)2^J$ terms:

$$Z = \{d^{J-1}, d^{J-2}, \ldots, d^L, s^L\} \qquad [8]$$

give the stationary wavelet transform, a redundant representation of the original time series.

Wavelet-Based Wiener Filter

Roughly speaking, the denoising filter to be derived here and employed throughout the rest of this article involves a modification of Eq. [2] expressed in the wavelet domain. The concept of a wavelet-based Wiener filter is, of course, not new. As mentioned previously, similar ideas have been explored and applied to ERP (16) and electrocardiogram (ECG) data (21). The approach taken here is different from the previous approaches in two aspects. First, instead of designing a filter to denoise an averaged epoch, the filter implemented here is intended for denoising individual epochs. Second, the filter is applied in the SWT domain rather than the DWT domain. The first aspect is needed because epoch-to-epoch variations may be of interest in ER-fMRI, as discussed earlier. When the activation signal varies slightly from one epoch to another, it is advantageous to denoise each epoch separately so that the processed data still reflects this epoch-to-epoch variation. The second aspect of using the SWT is essential because the redundancy provided by this method leads to a more stable filter (24).

The derivation of Eq. [2] assumes that the power densities of the signal and noise are known. In reality, they are to be estimated from the measured data. To this end, a number of approaches have been suggested. For example, de Weerd and Kap (15) introduced a method for processing ERP data based on two types of averages, along with the assumption that each measured evoked response consists of a deterministic, time-varying component plus a random noise component. In this work, a method similar to that introduced by Bertrand et al. (16) is used.

Let us denote a time course with K epochs, each epoch consisting of N samples, as:

$$\{x_0[0], \ldots, x_0[N-1], \ldots, x_k[0], \ldots, x_k[N-1], \ldots, x_{K-1}[0], \ldots, x_{K-1}[N-1]\}. \quad [9]$$

Hence, the ensemble average reads:

$$\tilde{x}[n] = \frac{1}{K}\sum_{k=0}^{K-1} x_k[n]. \quad [10]$$

Performing the SWT on an individual epoch k yields:

$$\{d_k^{J-1}[0], \ldots, d_k^{J-1}[N-1], \ldots, d_k^{J-2}[0], \ldots, d_k^{J-2}[N-1], \ldots, d_k^{L}[0], \ldots, d_k^{L}[N-1], s_k^{L}[0], \ldots, s_k^{L}[N-1]\}. \quad [11]$$

Similarly, the same operation on ensemble average yields:

$$\{\tilde{d}^{J-1}[0], \ldots, \tilde{d}^{J-1}[N-1], \ldots, \tilde{d}^{J-2}[0], \ldots, \tilde{d}^{J-2}[N-1], \ldots, \tilde{d}^{L}[0], \ldots, \tilde{d}^{L}[N-1], \tilde{s}^{L}[0], \ldots, \tilde{s}^{L}[N-1]\} \quad [12]$$

In the wavelet domain, the desired Wiener filter takes the form:

$$H(j, n) = \frac{P_x(j, n)}{P_x(j, n) + P_v(j, n)} \quad [13]$$

where $P_x(j,n)$ is the power density corresponding to the detailed component of true signal x in location n at resolution level j. $P_v(j,n)$ is the corresponding term of the noise. It can be shown that the denominator can be approximated by:

$$P_x(j, n) + P_v(j, n) \approx \frac{1}{K}\sum_{k=0}^{K-1} (d_k^j[n])^2 \quad [14]$$

with the approximation approaching equality as K tends to ∞. To get an estimate of the numerator of Eq. [13], it is noted that:

$$(\tilde{d}^j[n])^2 \approx P_x(j, n) + \frac{1}{K} P_v(j, n). \quad [15]$$

This progression from Eq. [13] follows very closely the results in (16). Combining Eqs. [14, 15] gives:

$$P_x(j, n) \approx \frac{K}{K-1}(\tilde{d}^j[n])^2 + \left(\frac{1}{K} - \frac{1}{K-1}\right)\sum_{k=0}^{K-1} (d_k^j[n])^2. \quad [16]$$

By substituting Eqs. [14] and [16] into Eq. [13], we obtain:

$$H(j, n) = \left(\frac{\frac{K}{K-1}(\tilde{d}^j[n])^2 + \left(\frac{1}{K} - \frac{1}{K-1}\right)\sum_{k=0}^{K-1} (d_k^j[n])^2}{\frac{1}{K}\sum_{k=0}^{K-1} (d_k^j[n])^2} \right). \quad [17]$$

Since it is only possible to obtain approximate values for the power densities in Eq. [13], the values of $H(j,n)$ from Eq. [17] are not constrained to the range of [0,1]. Thus, $H(j,n)$ is truncated so that its values must fall within this desired range; negative values are replaced with zeros, and values greater than one are replaced with ones. Now, the denoised experimental data in the wavelet domain can be written as:

$$\{\hat{d}_k^{J-1}[0], \ldots, \hat{d}_k^{J-1}[N-1], \ldots, \hat{d}_k^{j}[0], \ldots, \hat{d}_k^{j}[N-1], \ldots, \hat{d}_k^{L}[0], \ldots, \hat{d}_k^{L}[N-1], s_k^{L}[0], \ldots, s_k^{L}[N-1]\} \quad [18]$$

where $\hat{d}_k^j[n] = d_k^j[n]H(j,n)$. Note that filtering is not applied to the smooth components of each epoch, $s_k^L[1], \ldots, s_k^L[N]$.

In view of Eq. [13], the action by Eq. [18] on the raw data is to shrink each of its wavelet coefficients according to the ratio between the power density due to noise and the power density due to signal plus noise in the corresponding component.

It should be noted that even though an epoch-to-epoch variation is present in the data, the estimation of the Wiener filter assumes that this variation is negligible. Therefore, the estimated filter represents an "average" Wiener filter and is not necessarily optimal for each epoch from a theoretical point of view (see Discussion).

METHODS

The above algorithm has been implemented using the WaveLab package (25). All supporting routines were written in Matlab (MathWorks, Natick, MA). This approach was investigated by applying it to simulated fMRI datasets as well as activated time courses from experimental data. In the simulation studies, the performance of the present approach was measured by the reduction of RMS error and by the improvement in detectability when it is used as a preprocessing step for a model-free detection approach. Further, these performance measures of the present approach were compared to those of an analogous Fourier domain filter. In the experimental studies, the present approach was applied to activated time courses from experimental data obtained with a visually guided motor paradigm to illustrate its practical utility. All results reported utilized a Daubechies filter of length 4, as implemented in the WaveLab package.

Denoising Activated Pixels in Simulated Data

Simulated activation was generated using a signal of the form (see Fig. 1a for an example):

$$x(t) = \left(1 - \exp\left(\frac{-t}{T_1}\right)\right)^3 \cdot \exp\left(\frac{-t}{T_2}\right) \qquad [19]$$

where T_1 and T_2 are constants that can be adjusted to obtain the desired shape, and t represents the sampling times (i.e., the image number within an epoch). In most of the simulations described below, T_1 and T_2 took the values of 5.0 and 7.5, respectively. This signal was replicated for each epoch and added to either simulated Gaussian white noise or experimentally acquired baseline data. The baseline data were collected on a healthy adult male volunteer using an EPI sequence (TE/TR: 25/500 ms, FOV: 20 $\times$ 20 cm^2, matrix: 64 $\times$ 64, slice thickness: 5 mm, 2000 images) on a Siemens 1.5 T clinical scanner. To account for possible drift in the data, the mean of each epoch was removed from the corresponding epoch before denoising.

Simulations were performed to characterize the denoising algorithm with respect to SNR, the number of epochs, and small epoch-to-epoch variations of the fMRI response. These factors were examined to see how the performance depends on the accuracy of power density estimation. In these simulations, the root-mean-square (RMS) error, calculated between the filtered time course and the true signal, is used as a performance measure and compared to the RMS error in the original data. Each simulation was repeated 10 times. The mean and the standard deviation of the results from the ten repetitions are reported. Finally, simulations were performed to compare the wavelet-based Wiener filter to the Fourier-based Wiener filter in order to demonstrate the advantage of operating under the wavelet domain (26,27).

SNR

Simulations to explore filter performance with varying levels of SNR were performed. In this case, the synthetic data were generated with both simulated Gaussian white noise and baseline data. The SNR, defined as the standard deviation of the simulated signal over the standard deviation of the added noise, ranged from 0.25–4.0. The number of epochs was eight, and the length of each epoch was 64 time samples.

Number of Epochs

The filter performance with respect to varying number of epochs was studied. The SNR of the simulated data was 1.0 and baseline data were added to the simulated activity. The number of epochs ranged from two to eight, with each epoch having 64 time samples.

Epoch-to-Epoch Variations

The robustness of the filter was tested under the two conditions of epoch-to-epoch variation in simulated time courses. Two possible scenarios were investigated. The first addresses variations in the amplitude of the response, which may arise, for example, when a task of varying difficulty is used. The second focuses on variations in the duration of the response. For the condition of variable amplitude, each epoch was generated as described above, with the signal component (Eq. [19]) multiplied by a random value ranging from 1.0–2.0. In the case of variable response duration, the signal component of each epoch was convolved with a boxcar function having a width varying from 12–25 time samples. In both cases, the SNR of the simulated data was 1.0 with baseline data as noise, the number of epochs was eight, and the length of each epoch was 64 time samples.

Fourier Domain Filtering

A comparison of the noncausal Wiener filter as described by Eq. [2] with the SWT method was performed. The filter coefficients were estimated as described by Eq. [17] with the exchange of the wavelet coefficients with their Fourier domain counterparts. The data generated to investigate the performance of the SWT-based Wiener filter at different SNR levels were also used for this comparison. Quantitative comparison of the two methods was achieved by taking the difference of the average RMS error of the two methods and dividing the result by the average RMS error of the wavelet-filtered time series.

ROC Analysis

To study the impact of this method on detection of activation, ROC analysis was performed using simulated data

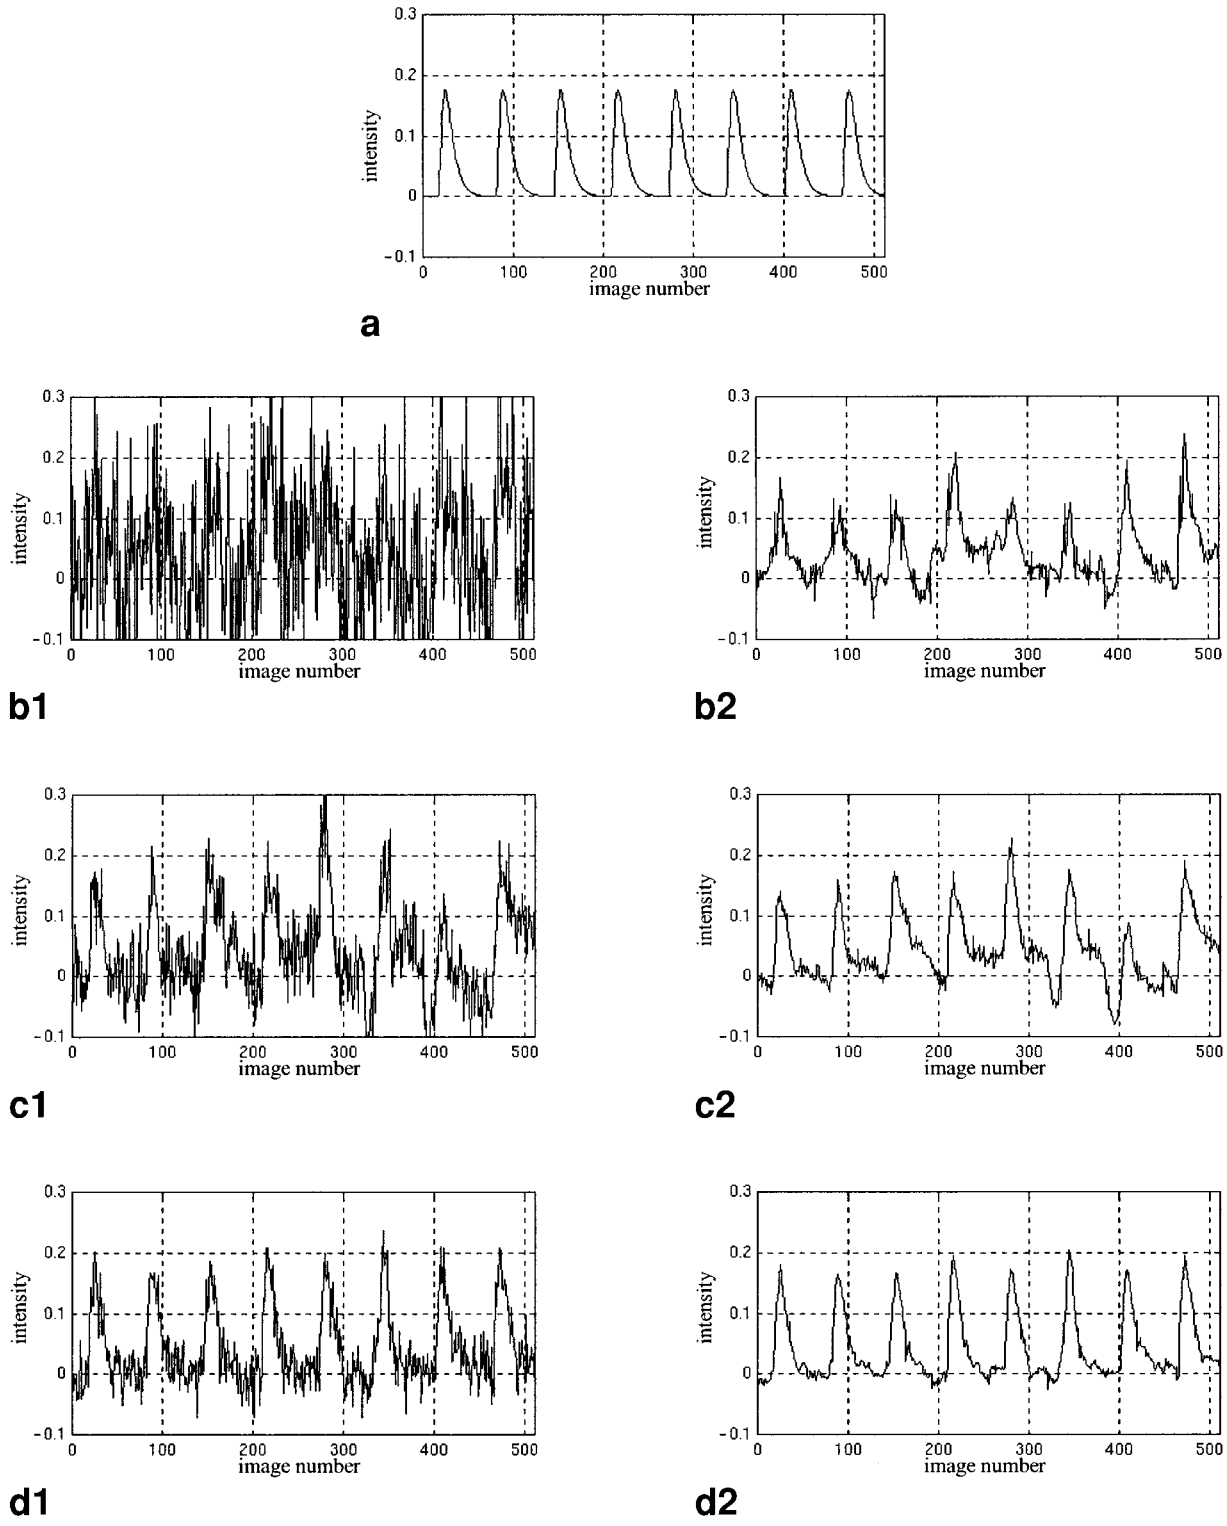

FIG. 1. Simulated time courses before and after denoising. The horizontal axis represents image number and the vertical axis corresponds to normalized signal intensity. **a:** The uncorrupted (true) signal. Baseline data were added to **a** to yield time courses with SNR of 0.5 (**b1**), 1.0 (**c1**), and 2.0 (**d1**), respectively. Their denoised counterparts are shown in **b2**, **c2**, and **d2**.

in a manner similar to that described by Xiong et al. (28). Four epochs of artificial activation (each with a duration of 64 images) with the signal pattern of a half-sinusoid and a preassigned contrast level (ranging from 1–3%) were superimposed on the experimental dataset. The spatial locations of the activated pixels correspond to a 5×5 grid of squares. Finally, an ROI encompassing the extent of the brain was delineated and the pixels within the ROI were

subjected to wavelet denoising. The nonfiltered and filtered datasets were separately submitted to Kohonen's self-organizing map (SOM), a neural network data-clustering algorithm (17,18). Upon completion of the SOM analysis, node patterns from the resulting neuron map were ordered according to their correlation with the known signal pattern. Starting from the node with the maximum correlation, the collection of activated pixels was expanded in successive steps by adding one node at a time. The true-positive fraction and the false-positive fraction were determined for each step and plotted against each other to generate the ROC curve. The area under the ROC curve was used as a measure of the detectability. In order to obtain results with statistical significance, ROC analysis was performed for a total of 10 times at each preassigned contrast level.

Denoising Activated Pixels in Experimental Data

The activation study was performed on a healthy male volunteer (with proper consent and institutional approval) using a clinical 1.5 T scanner (Siemens Medical Systems, Iselin, NJ). An oblique slice through both the motor and the visual areas was imaged using a T_2^*-weighted echo planar imaging sequence (TE = 60 ms, TR = 300 ms, FA = 55°, FOV = 22 × 22 cm^2, slice thickness = 5 mm). A T_1-weighted anatomical image of the slice was acquired with a FLASH sequence. During the acquisition of the T_2^*-weighted EPI images, the subject was asked to perform right-handed rapid finger movement when flashing LED goggles were on. In each epoch, which lasted 19.2 sec, the LED goggles were turned on for 5.4 sec. Each run consisted of 32 epochs, with epoch lengths of 64 images. This corresponded to a total of 2053 images (the first five were discarded). The data were visually inspected for subject motion by examining the center of mass and viewing the images in the cine mode using the Stimulate software package (29) and were discarded if the magnitude of the gross motion was unacceptable (e.g., >2 mm). Activated pixels were determined using standard cross-correlation analysis on raw data. For illustration, six active pixels, three from the visual cortex and three from the motor cortex, were chosen at random for filtering. To gauge the performance of the filter, it was applied to the first eight epochs of the data and the results were compared with the average of all 32 epochs, which was used as a "gold standard."

RESULTS

Results from applying the denoising algorithm on simulated data demonstrate that it is capable of substantially reducing the noise in the data under various conditions. The application of the algorithm to experimentally acquired data also exhibits a significant improvement in the fMRI time course data. Such an improvement is expected to improve the accuracy of quantitative characterization in subsequent time course analysis. In addition, the denoising of the simulated dataset demonstrated improvement in detection with the SOM method.

Denoising Activated Pixels in Simulated Data

Results from one representative simulation, where the baseline data were used as background noise, are shown in Fig. 1. The reduction of noise is evident in the filtered time course. For example, at a SNR of 0.5, little coherent structure is apparent in the original time course. After filtering, the underlying structure emerges. At SNRs of 1.0 and 2.0, substantial reduction in noise is also achieved. Quantitative measures of the filter performance on signals with various levels of SNR are summarized in Fig. 2. It is apparent that the denoising algorithm decreases the RMS error for all the cases tested. In addition, it is also interesting to note that the filter performs equally well with simulated noise and baseline data.

Figure 3 demonstrates a monotonic decrease in the RMS error of the denoised time course as the number of epochs is increased. This is expected because the increased num-

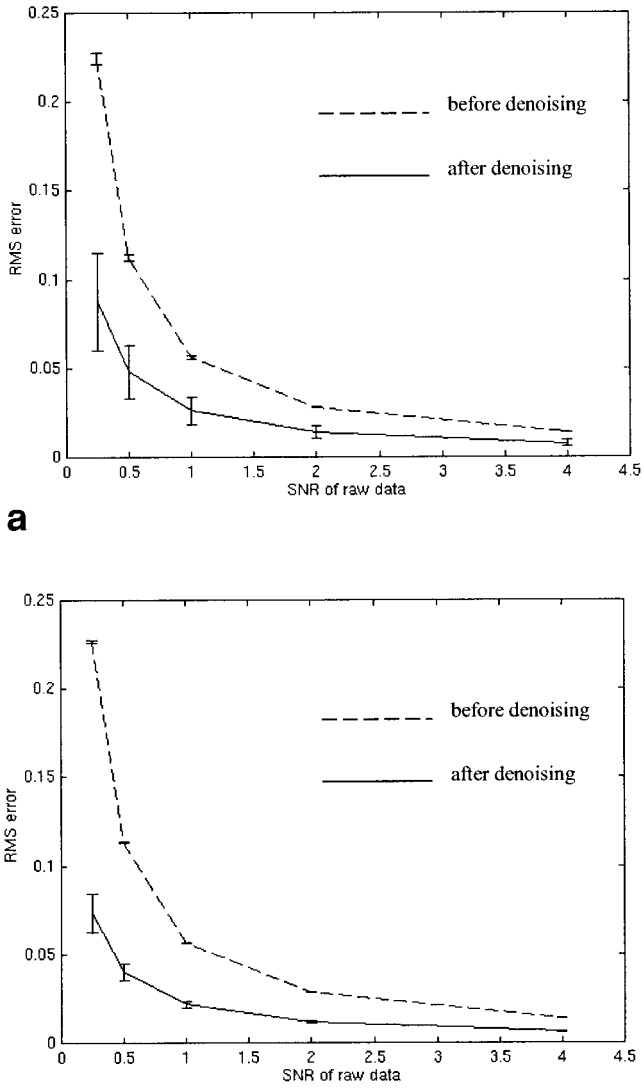

FIG. 2. Performance of filter with varying levels of SNR. The SNR of the raw data ranges from 0.25, 0.50, 1.00, 2.00, to 4.00. **a:** Results from simulation with baseline noise. **b:** Results from simulation with Gaussian white noise. The dashed line corresponds to RMS error of raw data with respect to the true signal before denoising, whereas the solid line indicates results after denoising. The error bars correspond to 1 SD from the 10 repeated simulations.

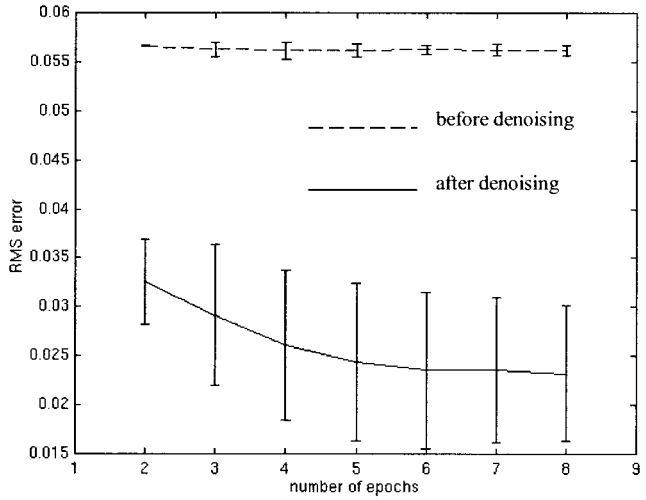

FIG. 3. Performance of filter with varying number of epochs. Baseline data were used to corrupt the true signal. The SNR of the resulting raw data was 1.0. The dashed line corresponds to RMS error of raw data with respect to the true signal before denoising, whereas the solid line indicates results after denoising. Error bars correspond to 1 SD.

ber of epochs allows us to obtain more accurate estimates for the approximation of Eq. [13]. The results are derived from simulations assuming a constant stimulus-induced response. When the response varies from epoch to epoch, the improvement in filter performance with increasing epoch number is expected to be less dramatic. Indeed, numerical experiments (data not shown due to space limitations) have revealed that this improvement is slightly less for the case of amplitude variation and more reduced for the case of width variation. The exact performance, however, depends on the SNR relative to the amount of response variation.

The results of the two numerical experiments conducted to show the effect of the wavelet filter in removing noise while preserving small epoch-to-epoch variation in the underlying signal are shown in Figs. 4 and 5. In Fig. 4, there is a small epoch-to-epoch amplitude variation in the simulated activation (top graph). The filtering operation is successful in removing a substantial amount of noise while preserving these variations. The result from a simulation in which the width of the activation varied is illustrated in Fig. 5. Again, the filter reduces a large amount of the baseline noise while preserving the variation in width in the denoised time course.

Figure 6 shows the percent difference between the average RMS error of the Fourier-filtered time course and the average RMS error of the wavelet-filtered time course. The results demonstrate that the Fourier domain implementation is consistently outperformed by the corresponding wavelet approach.

ROC Analysis

The results of the ROC analysis are shown in Fig. 7, which plots the detectability index of the prefiltered and filtered data as a function of the artificial activation contrast level. It is observed that Wiener filtering leads to statistically significant improvement in the detectability of activated regions at all contrast levels tested. Further, the wavelet domain filter showed improvement over the Fourier domain filter.

Denoising Activated Pixels in Experimental Data

The locations of the six randomly chosen pixels are shown in Fig. 8. Figure 9 contrasts the raw time courses with their denoised counterparts and clearly demonstrates a substantial improvement in the time course as a result of noise reduction. The improvements to both visual and motor time courses are verified by the numerical results listed in Table 1. These results confirm the applicability of the filter on experimental data.

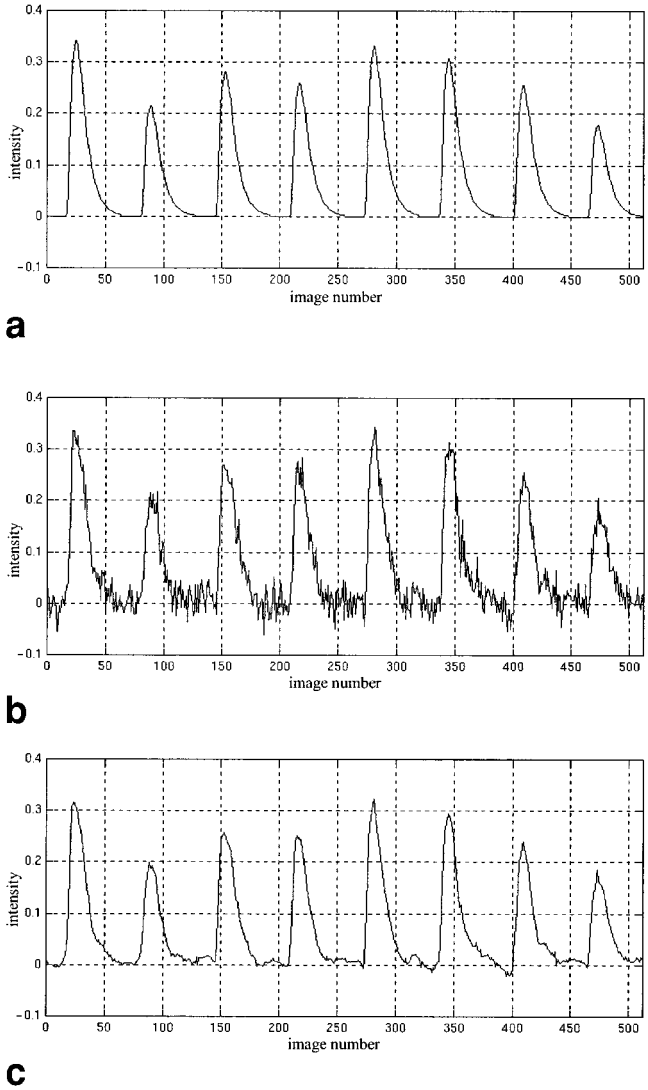

FIG. 4. Filtering of a noisy time course with intrinsic small amplitude variations in the uncorrupted signal. The horizontal axis corresponds to time in TR and the vertical axis corresponds to the signal intensity. The uncorrupted signal, which shows epoch-to-epoch variation in peak height, is displayed in **a**. The original signal is then corrupted by baseline data, resulting in a signal with SNR = 2.0, depicted in **b**. The result of denoising is shown in **c**.

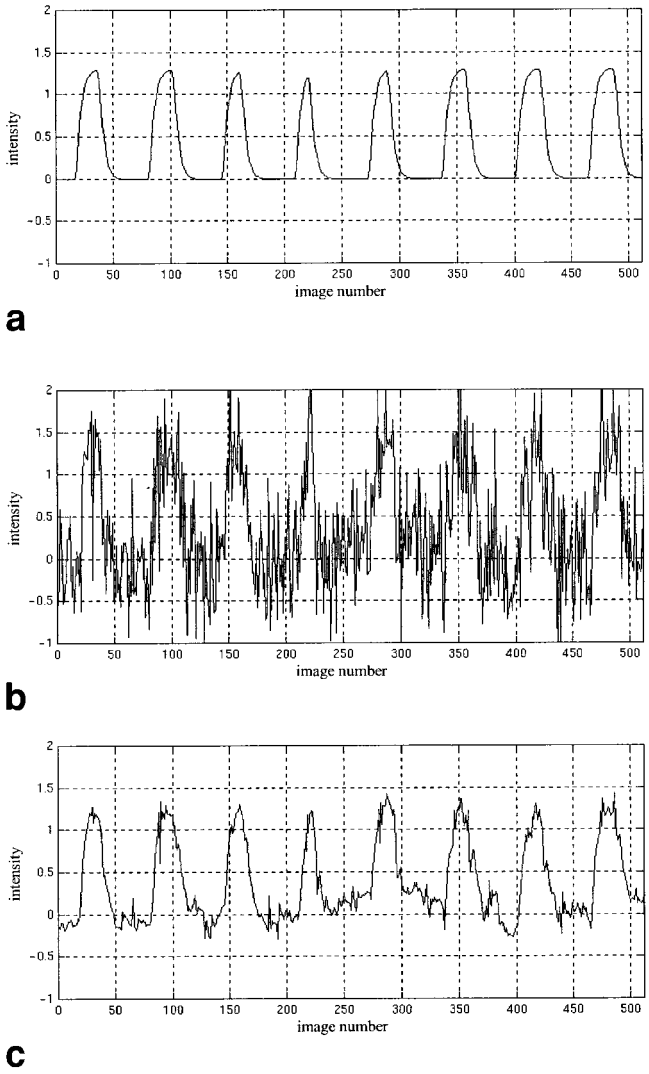

FIG. 5. Filtering of a noisy time course with intrinsic small width variations in the uncorrupted signal. The horizontal axis corresponds to time in TR and the vertical axis corresponds to the signal intensity. The uncorrupted signal is **a**, which shows the epoch-to-epoch variation in the width of the activated portion of the signal. The original signal was then corrupted by baseline data, resulting in a signal with SNR = 1.0, depicted in **b**. The result of denoising is shown in **c**.

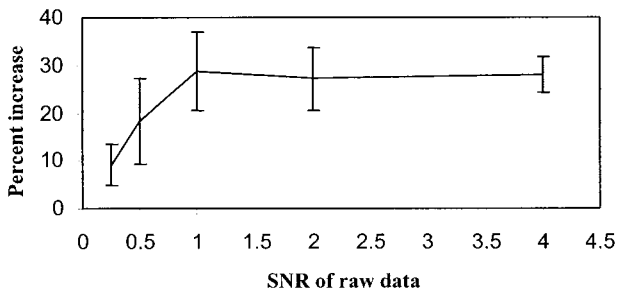

FIG. 6. Percent increase in the average RMS error of the Fourier-filtered time series with respect to the average RMS error of the wavelet-filtered time series for varying levels of SNR. Error bars correspond to 1 SD.

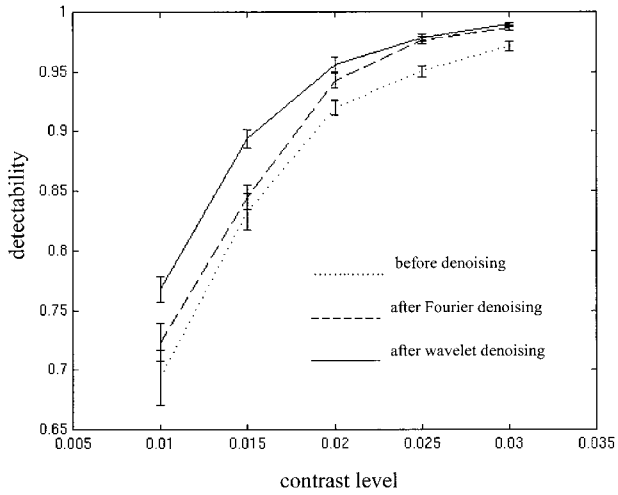

FIG. 7. Detectability index as a function of the artificial activation contrast level imposed upon experimental baseline data. The dotted line corresponds to results before denoising, the dashed line indicates results after Fourier denoising, and the solid line indicates results after wavelet denoising. At increased contrast level, the three curves tend toward 100% detectability asymptotically. Error bars correspond to 1 SD.

DISCUSSION

The wavelet transform-based Wiener filter, as developed in this work, has been shown to improve ER-fMRI data from repeated trials. A substantial reduction in noise (around a factor of three) was observed. This reduction is expected to improve subsequent quantitative analysis for extracting characteristics, such as amplitude, width, and onset time, of the individual trial responses. While the filter was developed for processing event-related fMRI studies, it is also applicable to block design experiments where stimulus/control blocks of identical duration are repeated. Indeed, its use could be extended to any data consisting of a measurement of such repeated blocks.

Other than the initial screening of the data for gross subject motion, the effect of head motion was not specifically addressed. The epoch-specific mean adjustment ac-

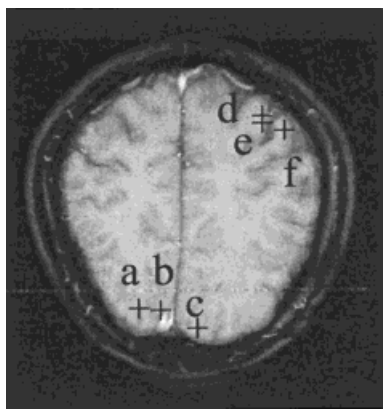

FIG. 8. Sample pixels selected from the visual and motor region for subsequent denoising. Pixels **a–c** are taken from the visual cortex, while pixels **d–f** are taken from the motor cortex.

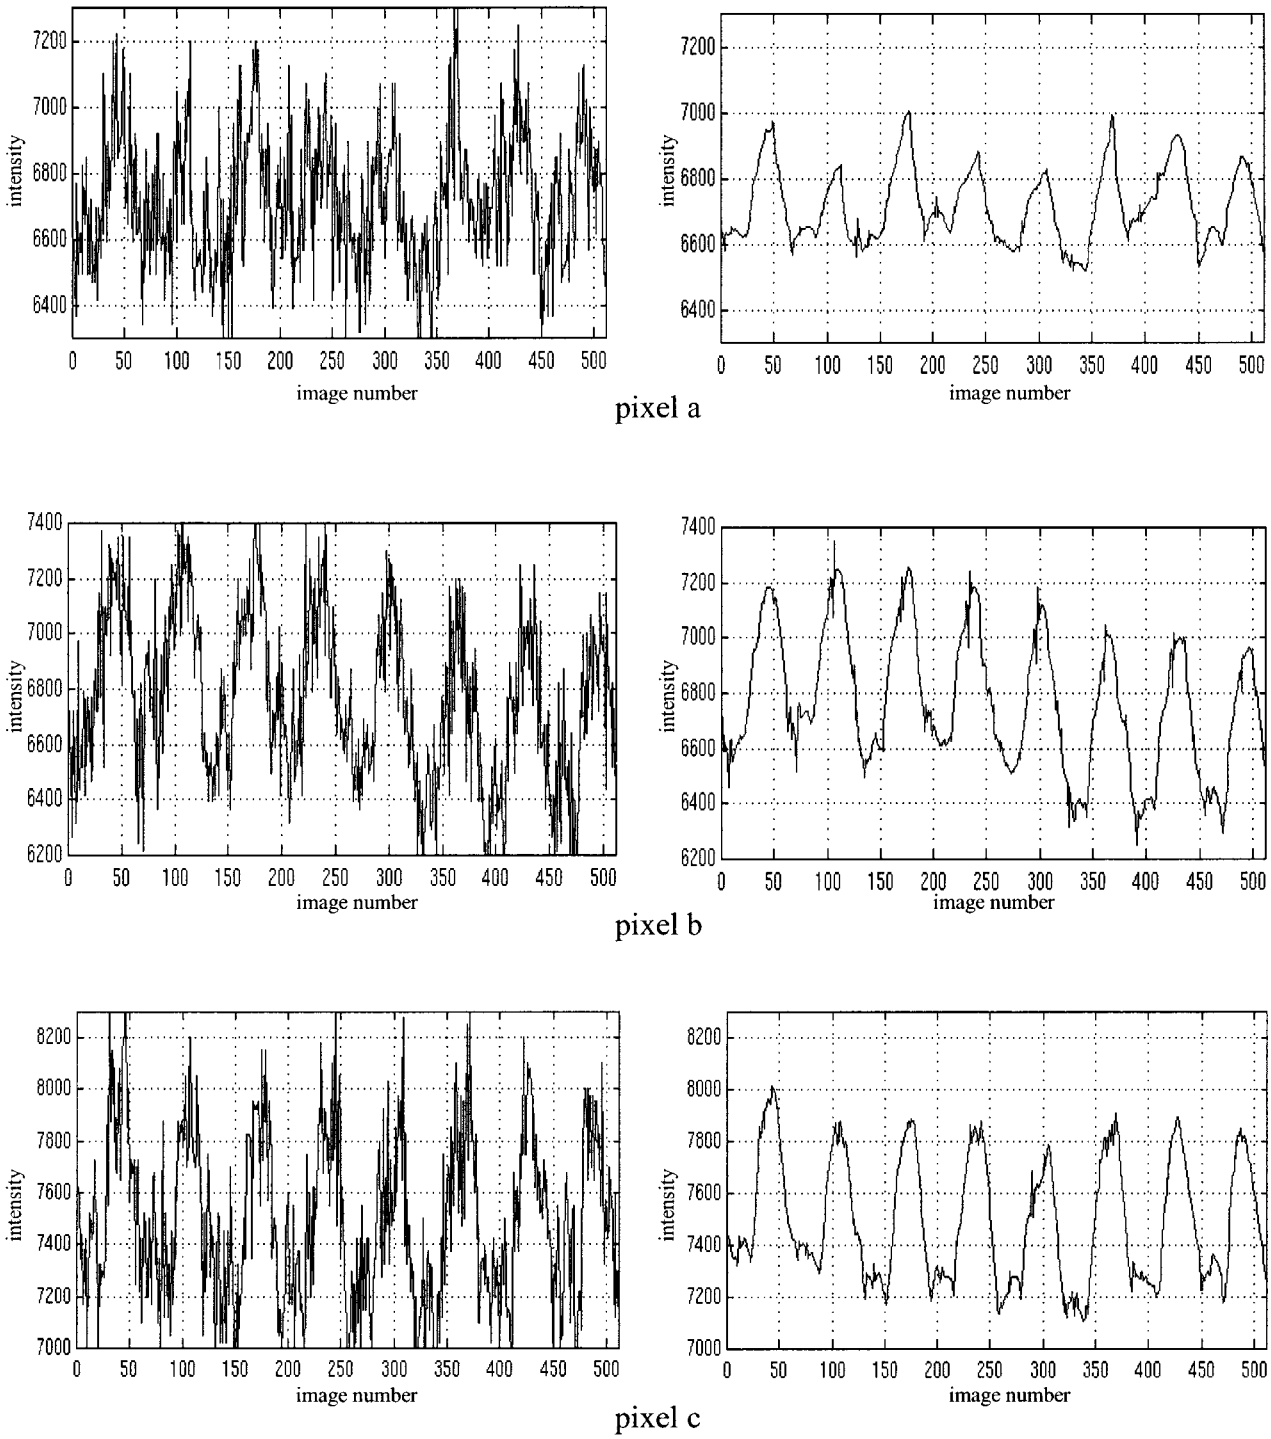

FIG. 9. Denoising of pixel time courses taken from the visual and motor cortexes. The horizontal axis corresponds to time in TR and the vertical axis corresponds to the signal intensity. The graphs on the left-hand side are the raw pixel time courses obtained from three sample pixels from the visual cortex (points **a**, **b**, and **c**) and from three sample pixels from the motor cortex (points **d**, **e**, and **f**) as indicated in Fig. 8. Time courses on the right are the result of denoising.

counts for some epoch-dependent movement effects in the data. If motion-related effects are large, the performance of the Wiener filter will be degraded due to an error in the estimation of the power densities. If the subject motion is synchronized with the stimulus, the movement effect will be preserved as a component of the filtered signal. On the other hand, for random motion, the filter is expected to reduce the motion-related fluctuation, as the fluctuation will be regarded as noise. Of course, standard motion correction algorithms can be applied before the application of the present method to avoid the complication due to motion if needed.

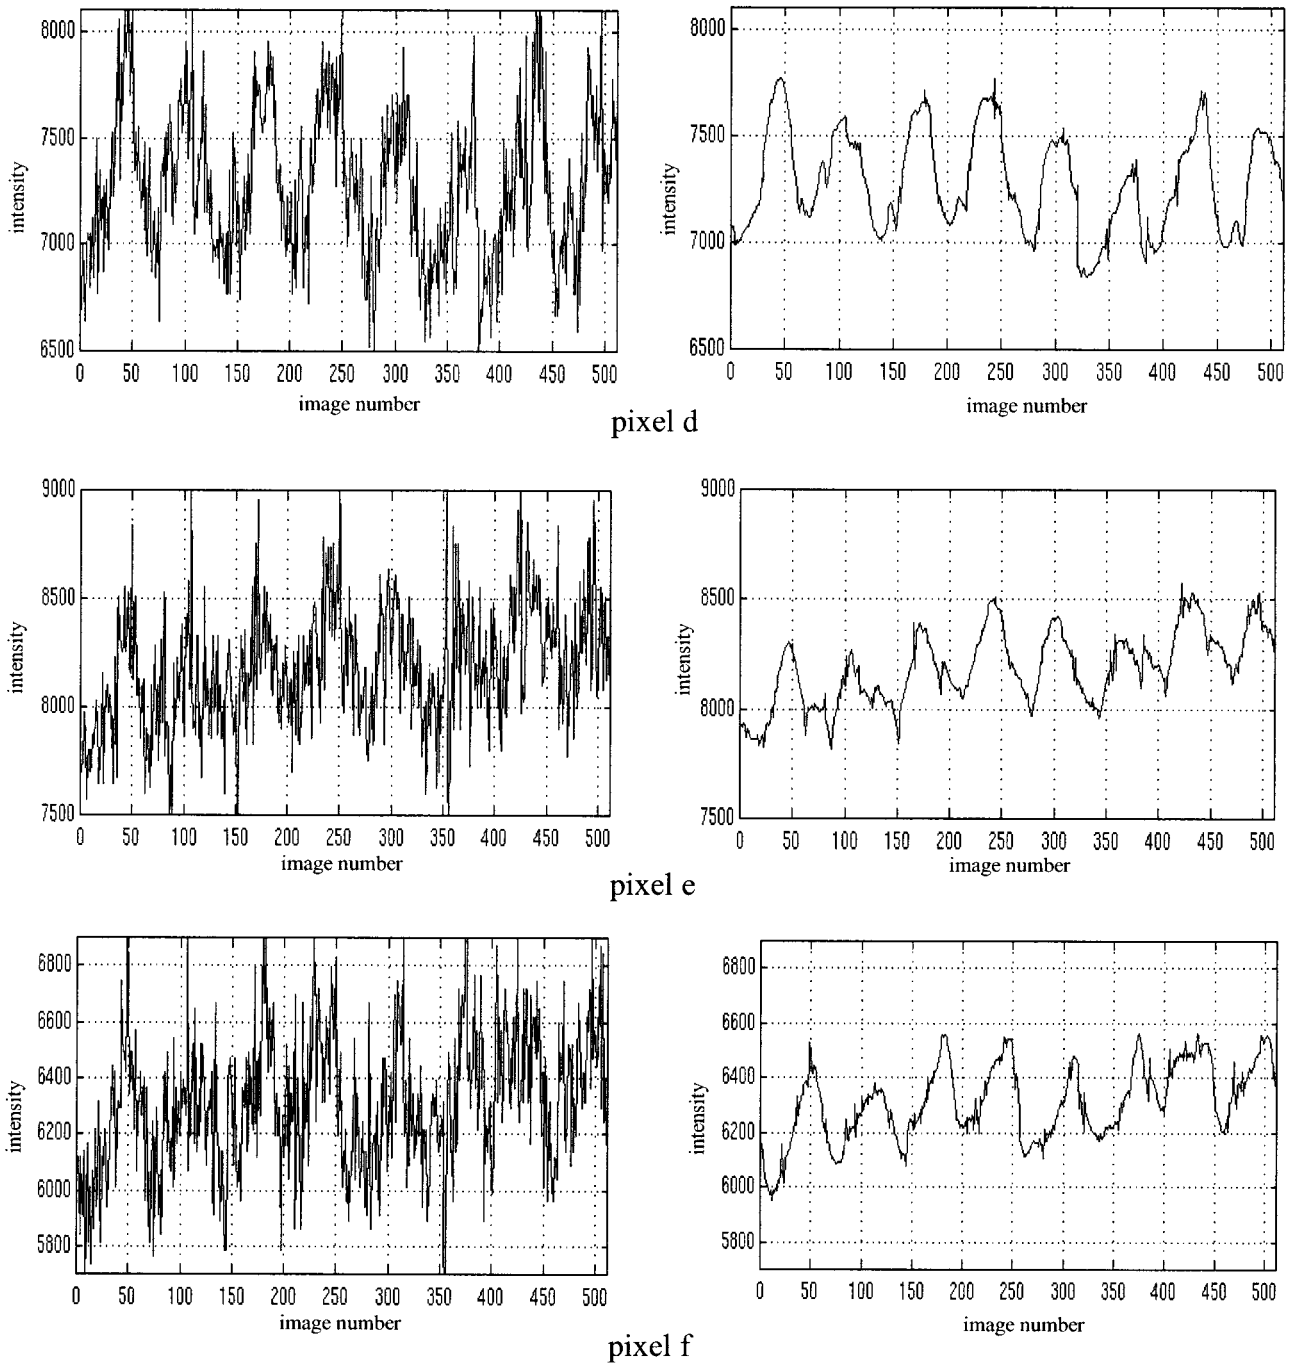

FIG. 9. Continued.

For the results reported above, a four-sample length Daubechies wavelet was used. It is likely that the filter performance depends on the choice of the wavelet function. To explore this possibility, six different wavelet functions (referred to here as filter name-parameter value) available in the WaveLab package (25), Symmlet-4, Symmlet-8, Coiflet-1, Coiflet-3, Daubechies-4, and Daubechies-8, were investigated using simulation. The RMS

Table 1
RMS Error of Pixel Time Courses Before and After Denoising, With Respect to the "Gold Standard"

	Pixel a	Pixel b	Pixel c	Pixel d	Pixel e	Pixel f
RMS error before denoising ($\times 10^2$)	2.42 ± 0.22	2.51 ± 0.27	2.78 ± 0.20	3.37 ± 0.57	2.84 ± 0.45	3.15 ± 0.36
RMS error after denoising ($\times 10^2$)	0.75 ± 0.14	0.90 ± 0.09	1.02 ± 0.16	1.40 ± 0.39	0.68 ± 0.14	1.26 ± 0.29

Pixels a, b, and c are selected from the visual cortex. Pixels d, e, and f are selected from the motor cortex.

errors of the filtered time course were found to be insensitive to the specific choice of the wavelet function (data not shown). It is not entirely surprising that very little change in performance was observed even when drastically different mother wavelets were used. This result arises from the utilization of the SWT, which is redundant. Because of this redundancy, the mismatch between the shape of the signal and that of the wavelet is not as critical as it would be in a nonredundant context. For this reason, it has been shown in other denoising algorithms that the use of SWT improves robustness (24).

The simulations resulted in similar improvements for both Gaussian white noise and baseline noise. Admittedly, the actual noise structure of the baseline data was not explicitly studied, nor was the algorithm tested under other noise models. However, others have shown that there is substantial autocorrelation in fMRI time series (30,31). In retrospect, the data used in the present study was found to follow an autoregressive model with a coefficient between 0.05 and 0.3, indicating significant deviation from Gaussian white noise. Despite this deviation, the results in this study demonstrated that the filter is applicable to experimental fMRI data.

The simulations also demonstrated that small epoch-to-epoch variations in the activation profile do not affect the performance drastically. In the presence of such variations, the final filter represents a tradeoff between accounting for individual differences and obtaining a good estimate of the power densities. While this variation represents a violation of the assumptions made in estimating the power densities, substantial improvement in the trial-specific responses was still achieved. This result indicates that the error in the power density estimates due to this epoch-to-epoch variation does not adversely affect the performance of the filter. While reducing the noise in the data, the filter also preserved the variations in the simulated signal. Hence, it will permit the extraction of epoch-specific information, which is likely to be of interest in single trial studies.

Because filtering alters the noise property of the data, making the application of conventional statistical detection methods difficult, these methods were not applied here. Instead, a model-free approach based on SOM was applied in conjunction with ROC analysis. This analysis showed that the wavelet-based Wiener filter improved the detectability of activated pixels over the Fourier-based Wiener filter.

A close comparison between the filtered and the original time courses, both simulated and experimental, reveals that there is a slight reduction in peak height associated with the filtering. Nonetheless, it should be noted that the filters are derived adaptively and the tradeoff between the bias introduced by the filter and the reduction in noise is, to a large extent, optimized. Therefore, a small reduction in the high-energy signal is acceptable.

Undoubtedly, the performance of the filter depends on the accuracy of the power density estimates. There are a number of possible approximations available for the type of data dealt with in this work. It is not clear which approach provides the best estimates. However, experience in other fields shows that the approach used here is sufficiently robust, and should perform at least as well as the potential alternatives.

CONCLUSION

A wavelet transform-based Wiener filter was developed and implemented for denoising single-trial fMRI data. By applying this filter to synthetic and experimental data, it has been demonstrated that the filter is effective in reducing noise in ER-fMRI time courses while preserving features of the activation response. In addition, application to an experimental visually guided motor study demonstrated the practical utility of the filter for enhancing the epoch-specific features of an experimentally acquired time course. Therefore, this filter is expected to improve quantitative analysis in ER-fMRI, although this utility remains to be fully demonstrated.

ACKNOWLEDGMENTS

The authors thank S. Sarkar, E. Yacoub, and J.C. Zhuang for help in experimental data acquisition, W. Auffermann for improving the exposition of the article, and the anonymous reviewers for useful comments and suggestions.

REFERENCES

1. Ogawa S, Tank DW, Menon R, Ellermann JM, Kim S-G, Merkle H, Ugurbil K. Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. Proc Natl Acad Sci USA 1992;89:5951–5955.
2. Kwong KK, Belliveau JW, Chesler DA, Goldberg IE, Weisskoff RM, Poncelet BP, Kennedy DN, Hoppel BE, Cohen MS, Turner R, Cheng H-M, Brady TJ, Rosen BR. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. Proc Natl Acad Sci USA 1992;89:5675–5679.
3. Bandettini PA, Wong EC, Hinks RS, Tikofsky RS, Hyde JS. Time course EPI of human brain function during task activation. Magn Reson Med 1992;25:390–398.
4. Ogawa S, Lee T-M. Magnetic resonance imaging of blood vessels at high fields: in vivo and in vitro measurements and image simulation. Magn Reson Med 1990;16:9–18.
5. Ogawa S, Lee T-M, Kay AR, Tank DW. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. Proc Natl Acad Sci USA 1990;87:9868–9872.
6. Ogawa S, Lee T-M, Nayak AS, Glynn P. Oxygenation-sensitive contrast in magnetic resonance image of rodent brain at high magnetic fields. Magn Reson Med 1990;14:68–78.
7. Thulborn KR, Waterton JC, Mattews PM, Radda GK. Oxygenation dependence of the transverse relaxation time of water protons in whole blood at high field. Biochem Biophys Acta 1982;714:265–270.
8. Rosen BR, Buckner RL, Dale AM. Event-related functional MRI: past, present, and future. Proc Natl Acad Sci USA 1998;95:773–780.
9. Buckner RL, Bandettini PA, O'Craven KM, Savoy RL, Petersen SE, Raichle ME, Rosen BR. Detection of cortical activation during averaged single trials of a cognitive task using functional magnetic resonance imaging. Proc Natl Acad Sci USA 1996;93:14878–14883.
10. Richter W, Georgopoulos AP, Ugurbil K, Kim S-G. Detection of brain activity during mental rotation in a single trial by fMRI. NeuroImage 1997;5:S49.
11. McCarthy G, Luby M, Gore J, Goldman-Rakic P. Infrequent events transiently activate human prefrontal and parietal cortex as measured by functional MRI. J Neurophysiol 1997;77:1630–1634.
12. Konishi S, Nakajima K, Uchida I, Sekihara K, Miyashita Y. Temporally resolved no-go dominant brain activity in the prefrontal cortex revealed by functional magnetic resonance imaging. NeuroImage 1997;5:S120.
13. Dale AM, Buckner RL. Selective averaging of rapidly presented individual trials using fMRI. Hum Brain Map 1997;5:329–340.

14. de Weerd JPC. A posteriori time-varying filtering of averaged evoked potentials. I. Introduction and conceptual basis. Biol Cybern 1981;41: 211–222.
15. de Weerd JPC, Kap JI. A posteriori time-varying filtering of averaged evoked potentials. II. Mathematical and computational aspects. Biol Cybern 1981;41:223–234.
16. Bertrand O, Bohorquez J, Pernier J. Time-frequency digital filtering based on an invertible wavelet transform: an application to evoked potentials. IEEE Trans Biomed Eng 1994;41:77–88.
17. Kohonen T. Self-organizing maps. New York: John Wiley & Sons; 1998.
18. Ngan S-C, Hu X. Analysis of functional magnetic resonance imaging data using self-organizing mapping with spatial connectivity. Magn Reson Med 1999;41:939–946.
19. Scarth G, McIntyre M, Wowk B, Somorjai R. Detection of novelty in fMRI using fuzzy clustering. In: Proc 3rd Annual Meeting ISMRM, Nice, France, 1995. p 238.
20. Lim JS. Two-dimensional signal and image processing. Englewood Cliffs, NJ: Prentice Hall; 1990.
21. Woolfson MS, Huang XB, Crow JA. Time varying Wiener filtering of the fetal ECG using the wavelet transform. In: IEE Colloquium on Signal Processing in Cardiography, 1995. p 11.1–11.6.
22. Pesquet JC, Krim H, Carfantan H. Time invariant orthonormal wavelet representations. IEEE Trans Signal Process 1996;44:1964–1970.
23. Nason GP, Silverman BW. The stationary wavelet transform and some statistical applications. In: Antoniadis A, Oppenheim G, editors. Wavelets and statistics. New York: Springer-Verlag; 1995. p 281–299.
24. Coifman RR, Donoho DL. Time-invariant denoising. In: Antoniadis A, Oppenheim G, editors. Wavelets and statistics. New York: Springer-Verlag; 1995. p 125–150.
25. Buckheit J, Chen S, Donoho D, Johnstone I, Scargle J. WaveLab .701, available at http://www-stat.stanford.edu/~wavelab
26. Le TH, Hu X. Methods for assessing accuracy and reliability in functional MRI. NMR Biomed 1997;10:160–164.
27. Metz CE. ROC methodology in radiologic imaging. Invest Radiol 1986; 21:720–733.
28. Xiong J, Gao JH, Lancaster J, Fox PT. Assessment and optimization of functional MRI analyses. Hum Brain Map 1996;4:153–167.
29. Strupp JP. Stimulate: a GUI based fMRI analysis software package. NeuroImage 1996;3:S607.
30. Purdon PL, Weisskoff RM. Effect of temporal autocorrelation due to physiological noise and stimulus paradigm on voxel-level false-positive rates in fMRI. Hum Brain Map 1998;6:239–249.
31. Tagaris GA, Richter W, Kim S-G, Georgopoulos AP. Box-Jenkins intervention analysis of functional magnetic resonance imaging data. Neurosci Res 1997;27:289–294.