NeuroImage **17,** 1018–1026 (2002)
doi:10.1006/nimg.2001.1017

Temporal Sensitivity of Event-Related fMRI

Luis Hernandez, David Badre, Douglas Noll, and John Jonides

FMRI Laboratory, University of Michigan, Campus Box 2108, Ann Arbor, Michigan 48109-2108

Received June 15, 2001

The temporal resolution of event-related fMRI is limited by the low sampling rate of typical MR whole brain protocols and by the slow rate of the BOLD response. Within the assumptions of the General Linear Model, we explore the tolerance of regression analyses of fMRI data to errors in the timing of the model relative to the experimental data under a number of circumstances. Given the sensitivity of the analysis to temporal shifts of the model relative to the data, one can search for the time a neuronal event occurs with temporal resolution on the order of a few hundred milliseconds with 95% confidence. This confidence level is strongly dependent on the signal-to-noise ratio of the observed BOLD responses (approximately ±200 ms in our example of visual stimulation data collected at 1.5 T).

INTRODUCTION

In order to study the flow of control and information between brain structures using fMRI, fine temporal resolution is necessary. In such studies, one would like to know what the sequence of neuronal events is in a particular task. The limitation of fMRI in resolving neuronal events arises mainly from the slow hemodynamic response observed by fMRI. The sampling rate of the MRI scanner is also limited (it takes less than 3 s to sample the whole brain), but luckily, since the response function is so slow and smooth (up to 20 s in duration), the sampling rate is still adequate under the Nyquist sampling theorem.

Thus, two related questions often arise in the design of an event-related functional MRI experiment. The first is whether we can distinguish between closely spaced events (200 ms) that happen in different brain structures (Badre *et al.,* 2000; Dale and Buckner, 1997; Wagner *et al.,* 1998; Buckner *et al.,* 1998). In that case, the limitations and the optimal experimental conditions are also of interest. If the neuronal events were shifted slightly in time, one would expect that a shift in the BOLD response should also occur. The issue is whether such a shift is observable, given the very low amount of detectable signal of the BOLD response (commonly 5% or less of the MR signal intensity), which makes the detection of the signal difficult in the first place.

The second question concerns the *tolerance* to time shifts in the model with respect to the data. This issue arises when there is an uncertainty in the onset time of the neuronal events or when there is lack of tight control over the synchronization between scanner and stimulus/response hardware. In this case, the experimenter expects a simpler activation pattern in the brain and doesn't need to distinguish between cognitive subprocesses that may occur at different times, so it would be beneficial for the investigator to have a somewhat higher tolerance of timing errors.

The temporal characteristics of the BOLD response have been explored previously by others. For example, Lee and Glover (1995) observed delays of 4 to 8 s in the vascular response to photic stimulation between voxels in gray matter and in larger vessels. In their study, they examined the phase of the correlation between the data and the model. Menon *et al.* (1998) used a "time-locked" approach to discriminate the difference in the onset time of the BOLD response of the right and left visual cortices to independent photic stimulation as short as 120 ms apart. This first study involved very fast imaging (TR = 100 ms) at 4 T at the expense of spatial resolution (single slice). Acquiring five slices (thickness 10 mm) every 480 ms, they were also able to discriminate response onset delays between SMA and visual and motor cortices in a visual–motor task (as low as 28 ± 11 ms). Their analysis was based on fitting straight lines to the on-ramp BOLD response. This approach served them quite well with the very high signal-to-noise ratio (SNR) responses they measured. Although perhaps not suitable for work at lower magnetic field and finer spatial resolution, that work demonstrates that fine temporal resolution of neuronal events can be resolved by measurements of the BOLD response. The drawback of such "time-locked" approaches is that they limit the experimental design to sparsely separated trials that allow for the BOLD response to decay completely (>20 s). The current trend, however, is to "deconvolve" the BOLD response out of an experimental data set and a stimulus function.

1053-8119/02 $35.00

These approaches include the use of a Wiener filter (Glover, 1999), linear regression (Hinrichs, 2000), or subtraction of the prior and future responses (Dale, 1997; Miezin, 1999). The main drawback to the deconvolution approach is that it is quite susceptible to noise. In the context of linear regression, it requires fitting 20 or 30 parameters at a time, in addition to the estimates of the confounds, which uses a lot of degrees of freedom. In the context of a deconvolution filter, one has to assume knowledge of the noise structure in order to design the filter appropriately.

The present study will try to answer the questions of sensitivity and tolerance within the framework of the General Linear Model approach to functional imaging data analysis and assuming knowledge of the hemodynamic response function from previous measurements. We'll take a three-step approach to the question. First we will derive an analytical equation for the correlation coefficient between the model and the data as a function of noise and time shift discrepancy, in order to optimize the experimental design. Second, we will generate simulated data and compute the correlation coefficients as we introduce noise into the data and as we introduce a time shift into the model. We will then examine what happens to the temporal resolution when there are errors in the model of the BOLD impulse response function, the measurement of the onset times. Third, we will test the theory on experimental data in order to validate it.

METHODS

Analytical Approach

We derived an analytical expression for the correlation between simulated real data and a model to be tested, including terms for the dependence on the timing shift and on the noise level of the data. The underlying assumptions for the derivation were that the model accounted for all sources of variance, in other words, that the model was an exact representation of the data, except for noise and a time shift. The second assumption was that noise in the time series was Gaussian and independent of the BOLD signal, and the third assumption was that the BOLD response was linear and time-invariant.

The derivation of the expression is as follows. The correlation coefficient between any two functions is usually computed by

$$\rho_{xy} = \frac{E[(X - \bar{X})(Y - \bar{Y})]}{(\sigma_x\, \sigma_y)}. \tag{1}$$

In our ideal case, the two functions are the BOLD response plus noise and the model of the response, which is assumed to be perfect except for a time offset (the other variable we wish to study). Thus,

$$Y(t) = B(t) + \varepsilon(t), \tag{2}$$

$$X(t) = B(t - T), \tag{3}$$

where $B(t)$ is the BOLD response to a set of stimuli, T is the time shift between the data and the model, and $\varepsilon(t)$ is the zero-mean noise in the measurement. Since the BOLD response is assumed to be linear and time-invariant, the response to a set of stimuli can then calculated by convolving the hemodynamic response function (HRF) of choice with a train of events (delta functions) occurring at the stimulus times, such that

$$B(t) = HRF(t) \times \sum_{trial} \delta(t - T_{trial}). \tag{4}$$

By substituting the BOLD response into the correlation coefficient equation above, and performing a few algebraic manipulations, one can show that

$$\rho_{xy} = \frac{E[B(t)B(t - T)] - \bar{B}^2}{\sigma_B \sqrt{\sigma_\varepsilon^2 + \sigma_B^2}} \tag{5}$$

($\bar{B}$ is the average BOLD signal over the interval and σ_B is its standard deviation). Note that this equation requires a sufficiently large data sample such that $E[B(t)] = E[B(t - T)]$. Additionally, we can compute the corresponding t score of the correlation coefficient as a function of the noise and the time shift between the model. The t score is usually computed from the correlation coefficient by

$$t = \frac{\rho_{xy}\sqrt{n - 2}}{\sqrt{1 - \rho_{xy}^2}}, \tag{6}$$

where n is the degrees of freedom (Shott, 1990). Note that the term $E[B(t)B(t - T)]$ in Eq. (5) is the definition of the autocorrelation function of $B(t)$, so the correlation coefficient equation is dominated by the shape of the autocorrelation function of the response. Thus, the sensitivity will be dependent on the choice of hemodynamic response model.

Simulations

All computations were carried out using MATLAB (Mathworks, South Natick, MA), on a Pentium III PC.

A model of the BOLD response to a set of stimuli was created by convolution of a set of stimuli with a canonical HRF. The stimuli were created by a spike train in which the intertrial intervals (ITI) of the stimuli were normally distributed about a mean of 16 s with a standard deviation of 4 s (the distribution was truncated to avoid negative ITIs) over a period of 600 s. The ITI

distribution was chosen to be Gaussian as an arbitrary example to simply illustrate the technique, although the 16-s mean ITI was chosen to preserve the linearity of the BOLD response. The data were created using a 0.01-s sampling period and subsampled to match the TR of the simulated experiment. These data were intended to serve as *B(t)* in Eq. (5). The correlation's *t* score was computed using Eqs. (5) and (6) for several levels of noise in the data (σ_ε ranging from 0.3 to 3 times the amplitude of a single BOLD response) and temporal shifts of the model (T ranging from −1 to 1 s at 0.1-s increments).

The canonical HRF was based on two gamma variate functions to characterize the rise and the ensuing undershoot observed in event-related fMRI experiments. This function is based on the shape of empirical data (Friston, 1994; Lange, 1997) and given by the following equations in the implementation of the SPM99 analysis software package (Wellcome Department of Cognitive Neurobiology, Institute of Neurology, University College of London, London, UK; Evans, 1993; Abramowitz, 1964):

$$HRF(t) = \frac{l_1^{\tau_1} t^{(\tau_1 - 1)} e^{(-l_1 \cdot t)}}{\Gamma(\tau_1)} - \frac{l_2^{\tau_2} t^{(\tau_2 - 1)} e^{(-l_2 \cdot t)}}{\Gamma(\tau_2)}, \quad (7)$$

where

$$\Gamma(\tau) = \int_0^\infty t^{(\tau - 1)} e^{-t} dt. \quad (8)$$

In order to verify the results from Eqs. (5) and (6), a model was created by convolving the stimulus function and the HRF, as before, and a data set was simulated by adding Gaussian noise to the model. Unless otherwise indicated, we used the default parameters of the software, which follow: $\tau_1 = 6$, $\tau_2 = 16$, $l_1 = 1$, $l_2 = 1$. The model was shifted in time between −1 and 1 s at 0.1-s increments to simulate the time shift errors in question. Thus, the correlation between the model with different time shifts and the data sets was then calculated. The correlations were computed using a multivariate linear regression approach, as different noise levels were added to the simulated data (zero mean Gaussian noise with variance 0.3 to 3 times the amplitude of a single BOLD response). Note that throughout this paper, we will use the variance of the noise as a measure of the noise level relative to the peak amplitude of the BOLD response. All the simulations were repeated 1000 times, and the average *t* scores were computed for each noise level.

In an additional set of simulations (noise levels ranging from 0 to 10 times the amplitude of a single BOLD response), we computed the temporal shift that produced the maximum *t* score, in order to identify the temporal shift of the model. After 1000 repetitions at each noise level, we calculated the mean and standard deviation of the results, in order to model the error of measurements of the temporal shift in a given model.

Similar simulations of time shifts were also carried out introducing additional discrepancies between the model's and the data's parameters. The errors were examined for two fixed levels of noise (variance 0.0001 and 3 times BOLD response), as well as in a visual stimulation experiment's data (described below). We examined errors in the HRF model, the ITI, and the number of events present in the data.

The effect of errors in the hemodynamic response function used to model the response was studied by varying the parameter τ_1 (τ_1 = 5.5, 6, and 6.5 s) in the HRF used to generate the model, while keeping it constant in the simulated data (τ_1 = 6 s). The parameter τ_1 changes both the peak time and the width of the rising portion of HRF, according to Eqs. (7) and (8). We examined the effect of errors in the onset times of the events (or ITI) by adding an extra amount of Gaussian jitter (standard deviation ranged from 0 to 1 s, mean = 0) to the timing of the events in the model relative to the simulated data and recomputing the correlations as before. We also examined the effect of discrepancies between the number of events present in the model and in the data, by adding and removing events from the simulated data (the discrepancy in number of events ranged from 10 to −10). Such errors can occur if noncompliance of the subject with the task goes unnoticed.

Experimental Data

In order to assess the validity of the simulations, a data set from a visual stimulation experiment was reanalyzed with different temporal shifts of the model, as well as including the mentioned discrepancies of the model's parameters. Six hundred volumes taken at an oblique angle of a subject's visual cortex were acquired on a 1.5-T Signa LX MR scanner (General Electric, Milwaukee, WI). The scans were acquired using a single-shot spiral imaging sequence (TR/TE/flip = 1000/35/35, matrix 64 × 64 × 12 slices, FOV 200 mm, slice thickness 3.4 mm) in a 10-min period, during which an alternating checkerboard pattern was flashed (flicker rate 8 Hz) in front of the subject's eyes, using a MR-compatible stimulus presentation system (IFIS; Psychology Software Tools, Pittsburgh, PA), for 0.5 s at a time. The intervals between stimuli were normally distributed about a mean of 16 s with a standard deviation of 4 s. After reconstruction of the *k* space data, the time series of images was interpolated to correct for differences in the acquisition times of each slice within

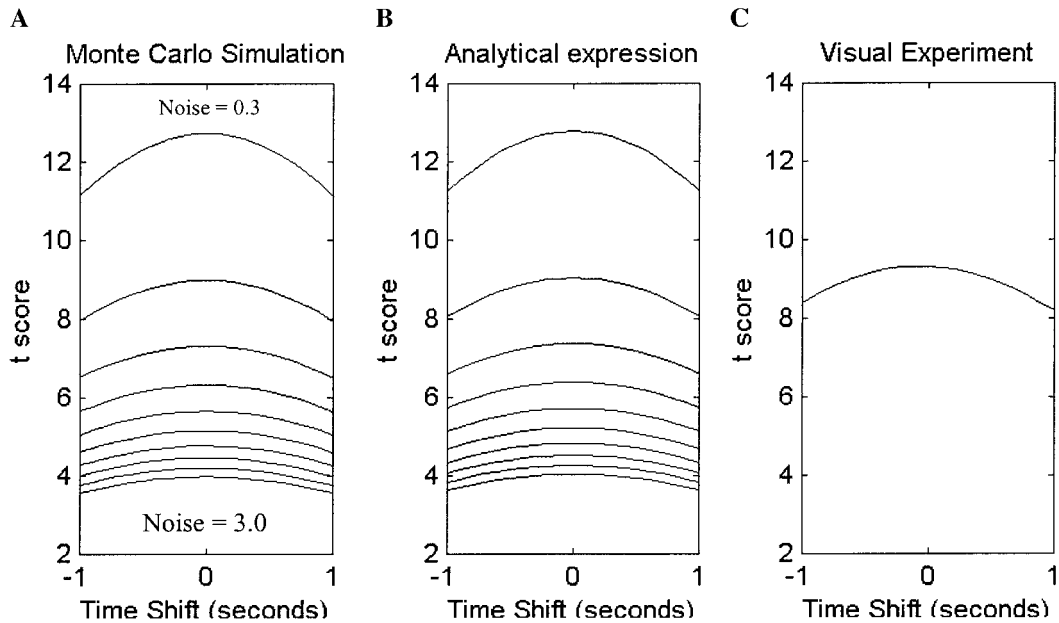

FIG. 1. Simulated correlations between model and data. The *t* score of the correlations is plotted as a function of the temporal shift between the model and the data, at different levels of noise present in the BOLD signal, ranging from 0.3 to 3 times the amplitude of a single BOLD response. The *t* score of the correlation is reduced as the noise level is increased. Note that the rate of change of the *t* score as a function of the time shift is considerably faster when the SNR is high. (A) The *t* scores obtained from analyzing data using a multivariate linear regression. In (B), the plot is generated by evaluating Eqs. (5) and (6). The same statistical analyses were performed on experimental data from the visual stimulation experiment, and the resulting scores are shown in (C). As expected, the three approaches are in excellent agreement.

each image. The images were then realigned to compensate for head motion and, subsequently, spatially smoothed with a Gaussian kernel (FWHM 6 mm in all directions) in order to increase the SNR and to better fit the random fields theory (Peterssen, 1999; Worsley, 1996). A multivariate regression analysis was carried out using SPM99 (Wellcome Department of Cognitive Neurobiology, University College of London, UK), in order to identify the voxels whose signals fit a model of the hemodynamic response to the experimental paradigm. The analysis included no additional linear confounds with the exception of baseline, in order to mimic the simulations as close as possible. The movement parameters were examined to verify that no gross movements (>2 mm, or 2°) were present. Once an active region was identified in the visual cortex, a time series was extracted from a region of interest within that active region, by averaging the voxels inside a 5-mm radius sphere centered on the maximum of the activation cluster. This time series was then subjected to the same tests as the simulated data above (temporal shifts, discrepancies in the HRF, ITI, and number of trials).

RESULTS AND DISCUSSION

As we discussed, we examined the effects of noise and time shifts of the model, in two ways: first, by plotting the correlation obtained from the derived Eqs. (5) and (6) as a function of time shift and noise level and second, by examining the correlation between model and simulated data. Figure 1a shows a family of plots (corresponding to different noise levels) of the correlation between the model and the response data as a function of the temporal shift of the model, as predicted by the equations. Figure 1b shows a similar family of plots of the equations as a function of the time offset between the model and the simulated data as we added noise to the simulated data. Figure 1c shows the result of analyzing the data in a voxel of interest from the visual cortex in the visual experiment (described under Methods) with different temporal shifts. The two simulation plots (a and b) are in excellent agreement, suggesting that the equations hold and can be used to predict the sensitivity of the model to temporal shifts. In addition, the experimental data in Fig. 1c also agree with the predicted outcome of the analysis. By comparison with the simulated data, the plot of Fig. 1c suggests that the variance of the noise is in the neighborhood of 50% of the amplitude of the BOLD response. There are two points to note about these plots: First, as the noise is incremented, the *t* scores are reduced and one can see that the slope of the curve gets flatter, indicating that there is less temporal sensitivity as the noise level is increased (Eq. (5) predicts that the noise acts as a scaling factor of the correlation coefficient).

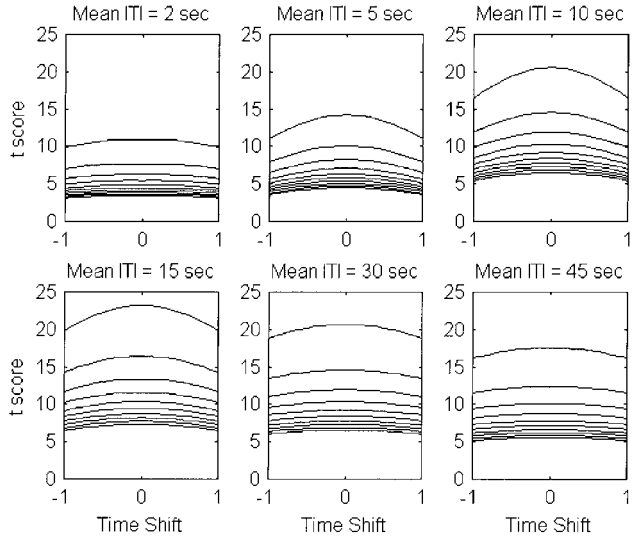

FIG. 2. Plots of the correlation's *t* score as a function of temporal shift of the model, showing the dependence on the mean intertrial intervals.

This implies that we are more likely to cross the chosen *t*-score threshold as we shift the model in time. At the same time, the *t* score is not reduced more than 10% of its value by the temporal shift in the model over the simulated ranges, which means that event-related fMRI experiments are quite tolerant of such timing errors. Second, the reduction of the *t* score as a function of the temporal shift of the model is symmetric within the 2-s window explored.

As we have shown, the derived equations can be used to examine the sensitivity of the model to temporal shifts, given an experimental design. For example, Fig. 2 shows the effect of the choice of ITI on the temporal resolution. The choice of ITI will have a direct impact on the variance of the response, σ_B, and on the autocorrelation function, which we showed to determine the shape of the sensitivity curve. By inspection, one can see that there is an optimum ITI for which the test is more sensitive to temporal shifts of the model around 15 s. This is not surprising, since Cox and Bandettini (1998; Bandettini, 2000) reported an ITI of 15 s as optimum to achieve maximum statistical power in event-related experiments, given the characteristics of the canonical HRF, and some of the earlier plots suggest that higher statistical significance increases the sensitivity of the analysis to temporal shifts of the model. Note that a shorter ITI allows for more events to be packed in the paradigm, which increases the energy of the model, but below a certain ITI, one starts to approximate a blocked design with only one condition, i.e., the model becomes closer to a flat line, so the estimate of the response amplitude becomes confounded with the estimate of the intercept (baseline signal) in the general linear model. One should keep in mind that this optimum ITI applies only to this particular trial randomization scheme and that the optimum ITI would be different for other randomization schemes, such as the semirandom one employed by Dale and Buckner (1997) and Miezin *et al.* (2000), which reportedly increases the statistical power of the test. We arbitrarily chose a normal distribution of ITIs for our studies, even if this is not optimum for statistical power, but the behavior of the statistical scores relative to the temporal shift shown here should apply for other distributions as well.

Although one usually does not have any control over the HRF of the subject, it is interesting to note that the time shift sensitivity of the analysis is also dependent on the HRF's shape. A sharper HRF is expected to be more sensitive than one that is slow and smooth. Equations (5) and (6) predict that the shape of the *t* score will be closely linked to the shape of the autocorrelation of the BOLD response function of the whole experiment, which is determined by the response to a single stimulus. This is illustrated in Fig. 3, in which the simulations are repeated, varying the parameter τ_1 in the HRF (which determines the width of the response and its time to peak). It should be noted that this equation is based on the assumption that there is an accurate measurement of the HRF. We will soon examine the effects of errors in the model of the HRF on the temporal sensitivity of the analysis more closely.

There are a number of mathematical models for the HRF (Cohen, 1997; Friston, 1994; Aguirre *et al.*, 1998; Buxton *et al.*, 1998; Franssin *et al.*, 1998; Hoge *et al.*, 1999; Boynton *et al.*, 1996; Lange and Zeger, 1997). Here, we have used a simple linear model based on a

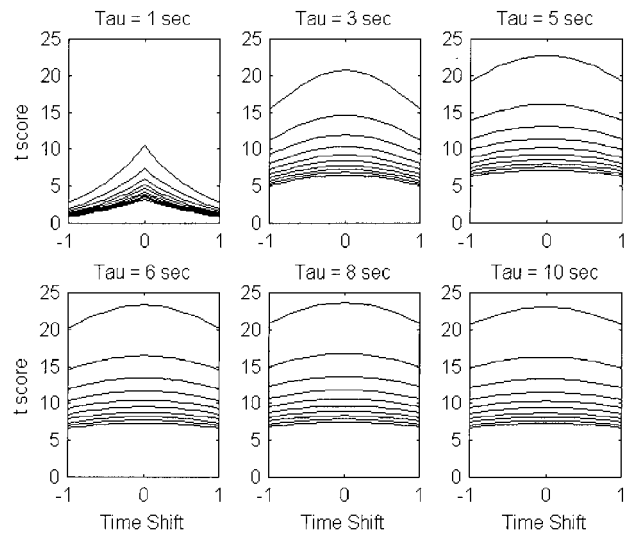

FIG. 3. Plots of the correlation's *t* score as a function of temporal shift of the model, showing the dependence on the shape of the HRF, as determined by the parameter τ [τ_1 in Eq. (7)]. Longer τ lengthens the duration of the HRF, whereas shorter ones make it sharper, as illustrated by the more extreme values of τ. The usual "physiological" value used in the vast majority of fMRI analyses is 6 s.

gamma variate function, because it resembles the physiological HRF and it has the flexibility to alter the width and the delay time. However, one should note that the same principles about time sensitivity should still apply to other models of the HRF. Since the HRF has been shown to exhibit nonlinear properties, one must take these into consideration as well, but that is beyond the scope of this study.

The simulations presented so far were performed under the assumption that the only error in the model was a temporal shift between the model and the data. The following results give an indication of what happens in the presence of additional experimental errors. Figure 4 shows a "matrix" of plots showing the dependence of the regression's *t* score on temporal shifts of the model. Each column shows the results of simulating a type of error that can be potentially found in the model in addition to the temporal shift, and each row represents the different kinds of conditions in the simulations as well as experimental data.

The first column of plots was obtained from simulations in which the HRF of the model and the HRF of the simulated data were in disagreement. The second was generated from data in which an extra jitter was added to the event times in the model relative to the data, in order to investigate the effects of inaccuracies in the measurement of the actual ITI. The third column explores what happens in the case in which some trials are missed due to subject noncompliance or in which extra events are present in the data, but not accounted for in the model.

The first row of plots was obtained from simulations with minimal noise levels (1–3% of the amplitude of the HRF); the noise was increased to more realistic levels (150% of the amplitude of the HRF) for the plots in the second row. The third row of plots was obtained from the visual stimulation experiment data. The fourth row shows the effects of the errors alone in the absence of a temporal shift in the model for both the experimental and the simulated data. Note that the jagged appearance of the simulation plots is due to the fact that a new random noise vector was used in each simulation, introducing variability into the correlation between simulations, and this variability remains even after 1000 averages of the simulations.

The errors in the model's canonical HRF are the most significant ones. Not only do they reduce the statistical score of the analysis, but they also cause the time shift sensitivity curves to peak earlier or later than one would expect. The shift of the curve is not very surprising since a faster HRF would peak and decay earlier. This sort of error would be very damaging if we were trying to identify the exact timing of neuronal events by this sort of "time-shifting" approach, but it means a reduction of statistical significance only in analyses that exclusively try to identify activation. It is a little more surprising to note that one also gets higher *t* scores (although the curve is shifted) when the HRF model is a little slower than the true HRF than when the model and the true HRF are in agreement. We hypothesize that, since a slower HRF function has more energy associated with it, it is less sensitive to noise. Thus, as the level of noise increases in the simulations, the *t* score is reduced faster in the case of a "sharper" HRF (i.e., Tau is shorter) than in the case in which it is a slower HRF. This is apparent comparing the plots in Fig. 4 that show the simulation of HRF errors with and without noise (plots A and D). In our initial simulations, we observed that all three curves got smaller as we gradually increased the noise. However, the curve on the left (i.e., with the erroneously slower HRF) did so slower than the other two. The puzzling result is that the *t* score of a regression analysis can be higher if the model HRF is too slow and there is a temporal shift than if the model HRF is correct and there is no temporal shift. We attribute this to the greater amount of power in the slower HRF models. Note that this ceases to be true as the noise in the simulations is reduced.

Discrepancies in the ITI and in the number of events between the model and the data result in a loss of statistical power, but the temporal shift sensitivity curves still peak at zero shift. In the case of extra events (i.e., ones not included in the model) being present in the data, the reduction of the *t* score is less than in the opposite case because these events simply contribute to the "noise," or unaccounted variance, affecting only the denominator of the *t* score (recall that a *t* score is essentially a signal to standard deviation ratio). In the opposite case, in which the subject misses events, the *t* score is more greatly reduced, since the variance of the data is increased *and* there is a reduction of the parameter estimate of the data, which affects the numerator of the *t*-score calculation.

The results from the experimental data are in excellent agreement with the simulated data, except for the experiments regarding the number of trials present. The reason is that the number of events was varied in the model, rather than the data, thus affecting the power of the regression, rather than its *t* score. This results in a much faster decrease of the *t* score as we introduce discrepancies in the number of events. One must also note the slight asymmetry in the shift sensitivity plots of the experimental data; these can be blamed on erroneous parameters in the HRF model, as discussed.

The above simulations suggest that it should be possible to determine the temporal shift between stimulus and response with great accuracy by searching for the temporal shift in the model that produces the maximum correlation. One can think of this approach as fitting an additional parameter [the temporal shift, *T*, in Eq. (3)] of the linear model whose slope parameters have already been estimated. Of course, such a mea-

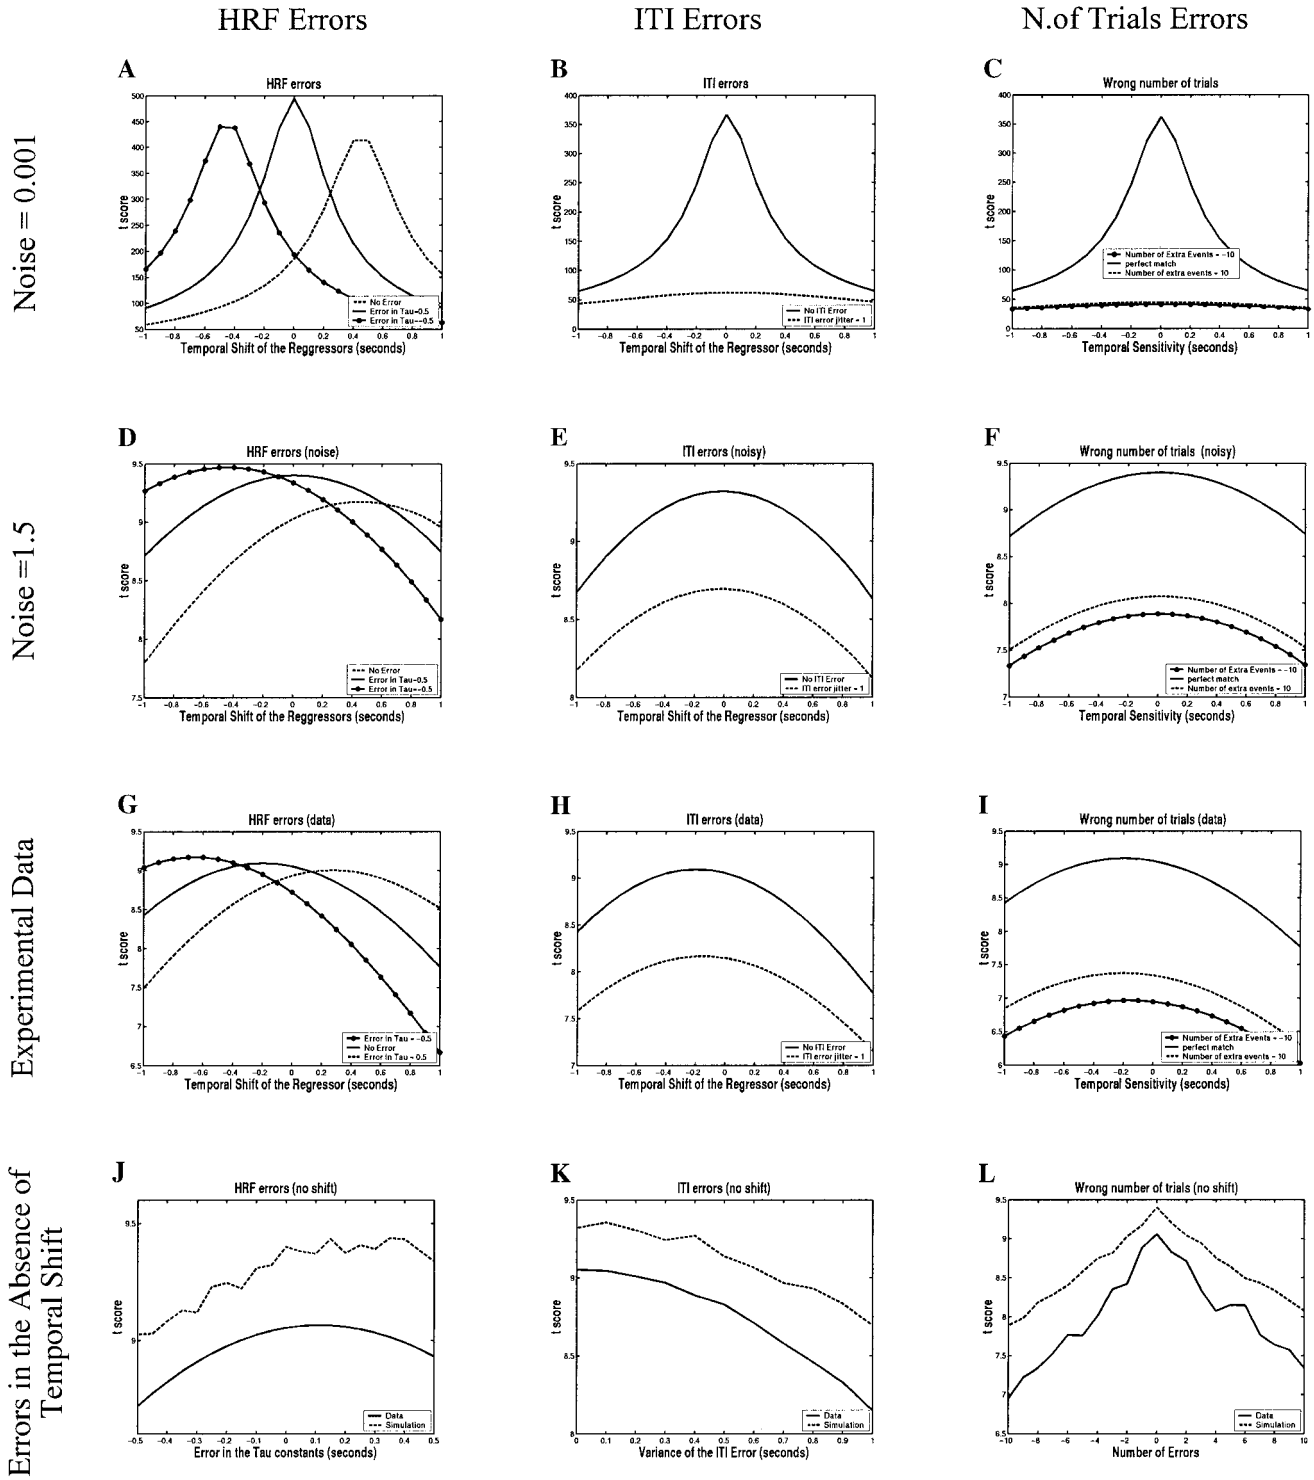

FIG. 4. Each column represents a type of error that can be potentially found in the model in addition to the temporal shift, and each row represents the different kinds of conditions in the simulations as well as experimental data. From this family of plots, we infer that all these errors lead to a reduction in the *t* score. However, discrepancies between the model and the data's HRF parameters lead to a shift of the temporal sensitivity curves. If the HRF is modeled to be faster than it really is, the *t* score will peak when the model is delayed relative to the real experimental paradigm. In other words, one will mistakenly identify the onset of the event at a later time. In the opposite case, in which the HRF is modeled too wide, the onset of the event will appear at an earlier time and, interestingly, the reduction of the *t* score will be a lot less than in the opposite case.

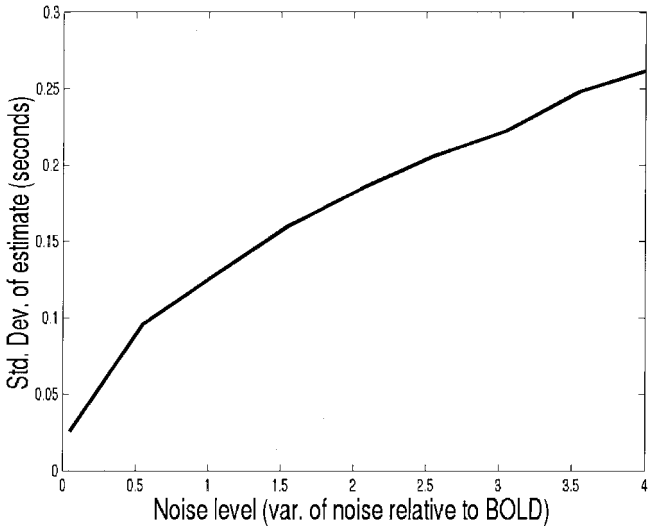

FIG. 5. Uncertainty of the estimate of the temporal shift in a model relative to the data as a function of the amount of noise in the data.

surement would have a variance associated with it, which would depend on the amount of noise present in the data. Figure 5 shows a plot of the variance of the model's temporal shift as a function of the noise in the data. The standard deviation was computed from a set of 10,000 simulations with noise levels ranging from 0.05 to 4.05 times the size of the BOLD response, in which the shift (T) was varied from -2 to 2 s at 0.2-s intervals. As the noise increases, so does the standard deviation of the measurement of the temporal shift (a histogram of the estimates for noise variance of 3.05 can be seen in Fig. 6). In a typical fMRI experiment, one might expect the noise to be up to three times the size of a single BOLD response. In such a case, the standard deviation of the temporal shift estimate would be approximately 200 to 400 ms. In the example of visual stimulation presented here, the noise is approximately 50% of the BOLD signal amplitude, which suggests a standard deviation of 100 ms according to our simulations. This in turn translates into an interval of $+/-200$ ms certainty at the 95% confidence level. There is a great variability of SNR depending on the field strength, subject's response, scanner hardware, etc. In our experience at 3 T, it is not uncommon to see noise levels under 10% of the BOLD response for similar visual stimulations. On the other hand, it is well known that other areas of the brain do not give such clear BOLD responses, especially at lower field strengths.

One must note that this estimate is possible only if the hemodynamic response of the tissue is well characterized. Otherwise, as indicated by the simulations involving errors in the HRF (Fig. 4), the estimate would be biased to an incorrect temporal shift and thus an incorrect timing of neuronal events. Thus, if timing of events is desired, it will be important to measure the HRF directly using a separate paradigm designed for that purpose.

The data presented here suggest also that regression analyses of fMRI data are quite tolerant, although observable, to time shifts in the experimental model. The effect is a small reduction of the significance of the analysis within the 2-s window explored due to the slow hemodynamic response of the human brain. However, in the case of a small response or high amount of noise in the data, this reduction of the significance can cause active voxels to fail the significance test because of a temporal shift.

The assumptions made in our model were that the HRF was linear and time invariant, which also underlie the theory of the General Linear Model. These assumptions have been examined by others, and they are reasonably met under normal experimental conditions (Boynton *et al.*, 1996), although they do break down under certain conditions (Buxton *et al.*, 1998; Vazquez and Noll, 1998). For our derivation approach, the noise in the BOLD signal was assumed to be independent of the signal, with a mean of zero, regardless of its distribution. The BOLD noise is made up of a number of components, among them respiratory and cardiac motion, as well as patient motion. Depending on the paradigm and the subject's compliance, this noise can have a certain degree of correlation with the paradigm and, consequently, the experimental model of the BOLD signal. We believe that our assumption is reasonable, since it is again one of the underlying assumptions for GLM analysis.

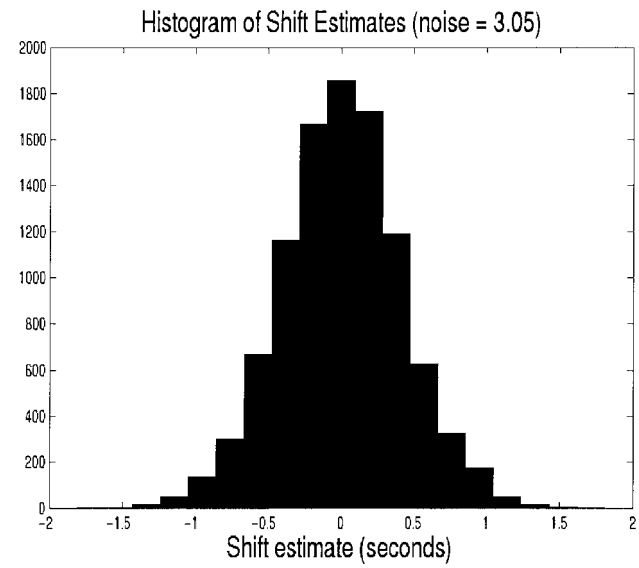

FIG. 6. Histogram of the estimates of the temporal shift, when the noise's variance is 3.05 times the amplitude of the BOLD response.

SUMMARY

Event-related fMRI can resolve closely spaced (on the order of hundreds of milliseconds) neuronal events when the data are analyzed through an iterative process of temporally shifting the model and computing a linear regression statistic (General Linear Model) until a maximum is found. This sort of technique has the potential to resolve questions regarding the timing of neuronal events and thus help us resolve issues of control between brain structures. However, accurate knowledge of the BOLD impulse response function, or hemodynamic response function, is crucial to determine the correct event time. The HRF could be measured using time-locked averaging or deconvolution of a simpler paradigm, as discussed in the Introduction. Some factors that affect the temporal sensitivity of the measurement are the intertrial interval and errors in its measurement, the variance of the noise present in the data, and the shape of the HRF as well as having an accurate measurement of it.

REFERENCES

Aguirre, G. K., Zarahn, E., and D'Esposito, M. 1998. The variability of human, BOLD hemodynamic responses. *NeuroImage* **8:** 360–369.

Badre, D., Jonides, J., Hernandez, L., Noll, D. C., Smith, E. E., and Welsh, R. 2000. Behavioral and neuroimaging evidence of dissociable switching mechanisms in executive function. Poster presented at the Cognitive Neuroscience Society Annual Meeting, San Francisco, CA.

Bandettini, P. A., and Cox, R. W. 2000. Event-related fMRI contrast when using constant interstimulus interval: Theory and experiment. *Magn. Reson. Med.* **43:** 540–548.

Boynton, G. M., Engel, S. A., Glover, G. H., and Heeger, D. J. 1996. Linear systems analysis of functional magnetic resonance imaging in human V1. *J. Neurosci.* **16:** 4207–4221.

Buckner, R. L., Goodman, J., Burock, M., Rotte, M., Koutsaal, W., Schacter, D., and Rosen, B. 1998. Functional–anatomic correlates of object priming in humans revealed by rapid presentation event-related fMRI. *Neuron* **20:** 285–296.

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

Clark, V. P., Maisog, J. V., and Haxby, J. V. 1998. fMRI study of face perception and memory using random stimulus sequences. *J. Neurophysiol.* **79:** 3257–3265.

Cohen, M. S. 1997. Parametric analysis of {fMRI} data using linear systems methods. *NeuroImage* **6:** 93–103.

Cox, R. W., and Bandettini, P. A. 1998. Single Trial FMRI: The Optimal Inter-Stimulus Interval. In *Proceedings of the Sixth ISMRM Conference,* Philadelphia, PA, p. 244.

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

Friston, K. J. 1994. Analysis of functional MRI time series. *Hum. Brain Mapping* **1:** 153–171.

Fransson, P., Kruger, G., Merboldt, K. D., and Frahm, J. 1998. Temporal characteristics of oxygenation-sensitive MRI responses to visual activation in humans. *Magn. Reson. Med.* **39:** 912–919.

Glover, G. H. 1999. Deconvolution of impulse response in event-related BOLD fMRI. *NeuroImage* **9:** 410–429.

Hinrichs, H., Scholz, M., Tempelmann, C., Woldorff, M. G., Dale, A. M., and Heinze, H. J. 2000. Deconvolution of event-related fMRI responses in fast-rate experimental designs: Tracking amplitude variations. *J. Cognit. Neurosci.* **12**(Suppl. 2): 76–89.

Hoge, R. D., Atkinson, J., Gill, B., Crelier, G., Marrett, S., and Pike, B. 1999. Investigation of BOLD signal dependence on cerebral blood flow and oxygen consumption: The deoxyhemoglobin dilution model. *Magn. Reson. Med.* **42:** 849–863.

Lange, N., and Zeger, S. 1997. Non-linear Fourier time series analysis for human brain mapping by functional magnetic resonance imaging (with discussion). *Appl. Stat. J. R. Stat. Soc. Ser.* C **46:** 1–29.

Lee, A. T., Glover, G. H., and Meyer, C. H. 1995. Discrimination of large venous vessels in time-course spiral blood-oxygen-level-dependent magnetic resonance functional neuroimaging. *Magn. Reson. Med.* **33:** 745–754.

Menon, R. S., Luknowski, D. C., and Gati, J. S. 1998. Mental chronometry using latency-resolved functional MRI. *Proc. Natl. Acad. Sci. USA* **95:** 10902–10907.

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

Petersson, K. M., Nichols, T. E., Poline, J. B., and Holmes, A. P. 1999. Statistical limitations in functional neuroimaging (I and II). *Philos. Trans. R. Soc. London Ser. B Biol. Sci.* **354:** 1239–1281.

Shott, S. 1990. *Statistics for Health Professionals.* Saunders, Philadelphia.

Vazquez, A. L., and Noll, D. C. 1998. Non-linear aspects of the blood oxygenation response in functional MRI. *NeuroImage* **7:** 108–118.

Wagner, A. D., Schacter, D. L., Rotte, M., Koutstaal, W., Maril, A., Dale, A. M., Rosen, B. R., and Buckner, R. L. 1998. Building memories: Remembering and forgetting of verbal experiences as predicted by brain activity. *Science* **281:** 1188–1191.

Worsley, K. J., Marret, S., Neelin, P., Vandal, A. C., Friston, K. J., and Evans, A. C. 1996. A unified statistical approach for determining significant signals in images of cerebral activation. *Hum. Brain Mapping* **4:** 58–73.