11. Effective Paradigm Design

David I. Donaldson and Randy L. Buckner

David I. Donaldson, Ph.D., Post-doctoral Research Fellow, Department of Psychology, Washington University, St. Louis, MO.

Randy L. Buckner, Ph.D., Assistant Professor, Departments of Psychology, Radiology, and Anatomy and Neurobiology, Washington University, St. Louis, MO.

In *Functional Magnetic Resonance Imaging of the Brain: Methods for Neuroscience.* (PM Matthews, P Jezzard, AC Evans, Eds.). Oxford: Oxford University Press, in press. 2000.

CONTENTS

11.1 Introduction

Advances in functional Magnetic Resonance Imaging (fMRI) are affording numerous new ways to design and analyze brain-imaging studies. A major driving force for these advances has been the widespread use of MRI scanners capable of ultra-fast imaging. For example, MRI scanners in many laboratories and clinical settings can acquire whole-brain images once every two seconds or even faster (see Chapter 6). However, the temporal resolution of fMRI is limited by the underlying physiological signals through which neuronal activity is indirectly measured (see Chapter 3). The optimal choice of an MRI experimental paradigm is determined by an interaction between methodological limitations and the properties of the physiology being measured. Extending the range of questions that can be addressed using the fMRI technique depends at least in part upon continued methodological advancement. This chapter aims first to review key aspects of MRI technology and the underlying physiology that constrain paradigm design for functional brain imaging.

Perhaps one of the most influential recent developments in terms of paradigm design has been the evolution and utilization of 'event-related' procedures. These procedures allow images to be formed of the transient neuronal changes associated with individual cognitive and sensory events, and even of the sub-processing stages within the events themselves. Event-related fMRI (ER-fMRI) allows experiments to depart solely from the use of 'blocked' procedures, introducing paradigms that isolate individual trial events separated by as little as a few seconds. A second aim of this chapter is to review issues related to the design of event-related experimental paradigms and how event-related procedures have extended the range of questions that can be asked. Finally, we will introduce the concept of mixed 'blocked and event-related' procedures. The development of such 'mixed' procedures widens the spectrum of task designs and analytical techniques that can be used, further extending the range of questions that can be addressed with fMRI.

Although aspects are discussed elsewhere in this book, we first present a brief review of the underlying characteristics of the hemodynamic response, to provide a context for discussing event-related methods and their derivatives.

11.2 Origins of the signal

Hemodynamic techniques such as blood oxygenation level dependent (BOLD) fMRI and PET (Positron Emission Tomography) index changes in blood properties (associated with relative oxygenation levels or blood flow respectively) that follow from changes in neural activity (as discussed in previous chapters). Hemodynamic techniques depend upon there being a close relation between neural activity and cerebral blood flow regulation within the brain – a coupling that is slow and that is affected by differences in regional vascular anatomy (cf. Lee et al. 1995; Robson et al. 1998; discussed in more detail below). In fact, temporal evolution of the fMRI signal is an order of magnitude poorer than the evolution of the neural activity itself, at least when brief sensory events are considered. Thus, it is worthwhile reviewing the temporal properties of the signal, their limitations, and how certain forms of paradigm design can best exploit the temporal information provided by the signal.

11.2.1 The BOLD-contrast hemodynamic response

BOLD is the most commonly measured contrast mechanism used for fMRI (Ogawa et al. 1990; Kwong et al. 1992; Ogawa et al. 1992; see Chapters 1 and 2). At present the exact origin of the BOLD signal remains unclear. Nonetheless, even our present understanding of how the BOLD-contrast signal evolves in relation to neuronal activity (henceforth referred to as the 'hemodynamic response') provides important guides to effective fMRI paradigm design. Thus, we briefly consider these findings here.

Perhaps the most relevant constraint of the hemodynamic response for issues related to paradigm design is that, as already noted, the response exhibits considerable temporal blurring in relation to the underlying neuronal activity. Although neural activity can occur very rapidly (in the order of milliseconds) in responses to a sensory event (Robinson and Rugg, 1988), changes in the hemodynamic response occur much more slowly (on the order of seconds). For example, even a brief period of sensory stimulation (as little as half a second) that presumably results in a proportionately brief period of neural activity, produces hemodynamic changes that do not begin for one to two seconds, and take place over a 10 to 12 second period (Bandettini 1993; Blamire et al. 1992; Boynton et al. 1996; Konishi et al. 1996). Moreover, recent evidence suggests that, although the robust positive deflection of the hemodynamic response evolves over a 10-12 second time interval, there may be physiological effects that last considerably longer (e.g., the 'post-stimulus undershoot', see Fransson et al. 1998; 1999; Buxton et al. 1998). Thus, inherent to fMRI data is the fact that the changes that are measured are lagged in time relative to the neural activity itself. As is discussed below, this is one of many factors that must be accounted for in the analysis procedures employed in event-related fMRI.

11.2.2 Reliability of the signal

The reliability of the BOLD signal as an indirect measure of neuronal activity has received support from a number of empirical observations. Most directly relevant to paradigm design is the finding that the response within a given subject and within a given region of cortex is extremely consistent from one set of measurements to the next (Dale and Buckner 1997; Menon et al. 1998; Miezin et al. 2000). Thus, although the signal is temporally blurred, the lag of onset and time course of signal evolution are highly reproducible.

Reliable maps of brain activity further suggest consistent properties of the BOLD-contrast exist across the cerebral cortex. For example, Ojemann et al. (1998; see also Casey et al. 1998; Clark et al. 1996) demonstrated reproducible patterns of activity associated with a speech production task in a comparison across studies produced in multiple laboratories employing both fMRI and PET. Several studies employing fMRI have demonstrated patterns of activity within striate and extrastriate cortex that could be predicted on the basis of the well-understood retinotopic organization of primary visual cortex (e.g., see Sereno et al. 1995; DeYoe et al. 1996; Engel et al. 1997). Such evidence builds confidence that the BOLD response is tightly coupled to neuronal activity.

The studies discussed above demonstrate the stability and reliability of the signal provided by the hemodynamic response. Nonetheless, evidence suggests that there is some variability in the exact form of the response from one subject to the next and from one brain region to another. For example, Miezin et al. (2000) revealed differences in the timing and amplitude of the hemodynamic response from the primary visual cortex across subjects. More significant however, is evidence that the hemodynamic response varies across brain regions within a given subject. This variability is exhibited in terms of the onset and shape of the response. For example, data from Buckner et al. (1998) revealed the onset of the response in extrastriate cortex to be about 1 second earlier than the response in prefrontal cortex during a word generation task. A related finding was presented by Schacter et al. (1997) in the context of a memory study, revealing a difference of several seconds between activity in anterior and dorsal prefrontal cortex. Even within visual cortex itself, Bandettini (1999) has shown variance in the timing of the response of one to two seconds across voxels. As these findings make clear, paradigm design and analysis must take into account the possibility of variance in the timing and shape of the hemodynamic response across regions.

11.3 Blocked paradigms

By far the most commonly used fMRI experimental paradigm is the 'blocked' task paradigm, whereby a series of trials in one condition are presented during a discrete epoch of time. The signal acquired during one blocked condition is then compared to other blocks involving different task conditions. In a typical study, task blocks will range in duration from 16 seconds to a minute and, in a single fMRI run (continuous period of data acquisition), multiple task blocks will be presented to allow the contrast of fMRI signals between task blocks. For example, to determine brain areas active during a language production task, task blocks in which subjects see words and produce a response to them could be contrasted with alternating task blocks in which subjects passively view words.

Reminiscent of earlier PET paradigms, blocked paradigms employ time-integrated averaging procedures. Not only were blocked paradigms the first approach to be employed in fMRI studies (Ogawa et al. 1992; Kwong et al. 1992; Bandettini et al. 1992), they were also the first targeted for the development of serious statistical analysis (e.g., Bandettini et al. 1993; Friston et al. 1994). Regions of activity change between one condition and another can be identified with considerable statistical power. Blocked procedures allow considerable experimental flexibility, for example allowing parametric designs and multi-factorial designs to be employed (Frackowiak and Friston 1995). The kinds of design possible with blocked paradigms have previously been reviewed extensively (for excellent reviews see Binder and Rao 1994; D'Esposito et al. 1999) and we will therefore only touch upon one important conceptual consideration.

Because fMRI methods are limited to finding relative changes between task comparisons, a crucial question is how to best design task comparisons in a given study. One important question is how closely matched the different tasks should be. A useful distinction that can be made in this regard is between 'tight' and 'loose' task comparisons (cf. Buckner 1996; Buckner and Logan, in press). Tight comparisons aim at holding as many extraneous variables as possible constant across tasks. For example, consider a blocked design by Demb et al. (1995) that examined meaning based decisions on visually presented words. They compared blocks involving a meaning based semantic decision (deciding whether each item was abstract or concrete) relative to equivalent blocks of a nonsemantic decision (deciding whether each item was presented in upper or lower case letters). To isolate differences related to processing word meaning, the two tasks were designed to differ in terms of the 'depth' of meaning based processing afforded to each item, whilst sharing other task demands, such as similar presentation of words, identical stimuli (when counterbalanced across runs), and grossly similar response demands.

An alternative approach is possible however. That is, 'loose' comparisons between tasks that are not closely matched, employing a much broader comparison across task variables. For example, visual fixation could be used as a low-level task. Buckner and Koutstaal (1998) used just such a low-level visual fixation task in a study that also employed both the meaning-based semantic and nonsemantic tasks used by Demb et al. (1995). Buckner and Koutstaal were thus able to compare the two meaning-based task conditions to each other (tight comparison) and either task condition to the 'fixation' task (loose comparisons). The resulting contrasts yielded brain areas active in common between the two word processing tasks and those that differed.

The use of 'loose' task comparisons described above may appear unnecessary. However, there are several clear merits in employing both loose and tight task comparisons simultaneously. First, the comparison between two closely matched tasks will only reveal regions that differ – any regions that are shared across the tasks will be subtracted out. It is often useful to see the entire set of regions that are activated during performance of a task, and not simply those regions that differ between two tasks, as a means of checking data quality and also as a means of identifying the entire network of brain regions active in a given

task. For example, the two word processing tasks described above would, a priori, be expected to give rise to activity associated with viewing the stimuli and making a motor responses – excellent markers of whether the experiment has succeeded and of the overall quality of the data. These regions could only be examined using a 'loose' task comparison between each task and the low-level fixation condition. Moreover, several component processes may be shared between conditions being compared but nonetheless be important components of the high-level demands of a task. Shared regions of activation often only become apparent when multiple comparisons are made to low-level reference tasks (e.g., Buckner 1996). A final reason to employ reference to a low-level control is that the same reference task can serve for multiple conditions and even studies. As fMRI data provide information about relative signal change, interpretations across conditions and studies can be made easier if contrasts are made with respect to the same reference task.

This is not to imply that 'loose' comparisons (which come with their own interpretive problems, e.g., see Binder et al. 1999; Shulman et al. 1997) are an alternative to the more closely matched 'tight' task comparisons. Nonetheless, it is often possible to include a variety of task comparisons within a study, building confidence in the quality of the data and answering theoretically driven questions – both loose and tight task comparisons should be considered.

11.4 Event-related paradigms

Event-related paradigms differ from blocked paradigms in that individual trial events (or even sub-components of trial events) are measured rather than a temporally-integrated signal. For example, in a typical event-related paradigm two separate trial types (e.g., words versus pictures) might be randomly intermixed in rapid succession. Then, the separate signal contributions of the two kinds of trial type are compared directly. Such a design may seem counterintuitive given the lag and temporally blurring inherent in the hemodynamic response. However, informative analysis of the complex signal is possible and affords considerable statistical power.

Several features of fMRI data proved to be critical in allowing event-related procedures to be developed. First, technological advances in the speed with which fMRI data could be acquired. Second, the fact that even very brief periods of neural activity give rise to measurable signal changes – despite the delayed and prolonged nature of the time-course of the hemodynamic response. Third, the hemodynamic response has been shown to provide a highly consistent response that summates over sequential events in a roughly linear fashion. Each of these issues is considered in turn below.

11.4.1 Rapid data acquisition

Event-related procedures necessarily require that the signal of interest can be sampled frequently and repeatedly, allowing the data to be acquired over the time course of an individual event. Thus, it is an even more important feature of ER-fMRI than of block designs that signals can be acquired extremely rapidly (e.g., using echo planar or spiral imaging methods). If only a few slices through the brain are required, measurements can be repeated in less than 1 second. Generally, however, functional imaging studies require whole brain coverage, thus many slices must be acquired, increasing the total acquisition time (sampling rate or repetition time, TR). Even so, a scan that provides whole brain coverage can be acquired with a total TR as little as 2 seconds on current systems.

Whilst the actual sampling rate of fMRI measures is ultimately limited (due to factors such as saturation of spins in the brain and the speed with which the imaging gradients can be switched), it is possible to further reduce the effective sampling rate by appropriate paradigm design. By staggering the timing of stimulus presentation relative to the timing of image acquisition, reconstructed signals can be produced with a shorter effective sampling rate than

is implied by the actual TR (cf. Josephs et al. 1997; Miezin et al. 2000). For example, consider the case whereby stimuli are presented every 5 seconds, and the signal is sampled (TR) every 3 seconds. The first sample is at zero seconds relative to the onset of the first stimuli, the second sample is at three seconds, the third sample is at one second relative to the onset of the second stimuli, etc. Thus, over the course of an experiment signals will, in effect, be acquired every second relative to the onset of the stimuli (Joesephs et al. 1997). An even simpler example is shown in figure 11.1, whereby the effective sampling rate is doubled using an "interleaved" procedure (Miezin et al. 2000). For a fixed TR of 2 seconds two data sets are acquired, providing odd and even sample points that contribute to a composite waveform with a 1 second sampling resolution.

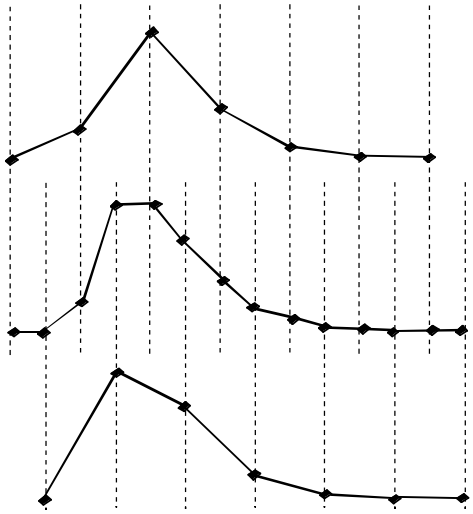

Fig. 11.1: Sampling the BOLD signal. The temporal resolution of the sampling of fMRI data can be increased beyond that of the basic sampling rate (TR) by simply varying the relationship between the onset of stimuli and the image acquisition. The figure demonstrates how staggering the point at which the signal is sampled can increase the temporal resolution. The averaged waveform (**middle**) is a composite of data from two separate sets of trials that sample sparsely, providing either odd (**top**) or even (**bottom**) data points. This method can be implemented by simply staggering the presentation of stimuli, presenting half the stimuli time-locked to the onset of the image acquisition (i.e., immediately) and the remaining stimuli half way between successive image acquisitions (i.e., delayed). Note that this method of improving the temporal resolution beyond the basic TR relies upon the averaging of data from different sets of stimuli and could not therefore be employed with single trial data.

Of course, for a given number of sampling points, such procedures result in less data being acquired, and therefore a poorer signal-to-noise ratio at each time point in the sample. Importantly, this does not imply a loss of power in the ability to detect the hemodynamic response signal over successive time points. In fact, in circumstances where fixed versus intermittent sampling procedures have been directly compared, the statistical power has been found to be either equivalent (Miezin et al. 2000) or better for intermittent sampling (Price et al. 1999).

In summary, it is clearly possible to increase the effective sampling rate beyond that of the TR, providing an increase in the temporal resolution of the data. Nonetheless, without such procedures fairly rapid sampling procedures (~ 2 sec per brain volume) are still achievable using currently available imaging systems, yielding a powerful means to estimate the hemodynamic response in event-related paradigms.

11.4.2 Sensitivity of the signal

The second feature of fMRI that was important in the development of event-related procedures is that the signal is highly sensitive, such that even a very brief period of neural activity elicits a measurable signal change. Transient increases in BOLD signal were initially demonstrated in sensory and motor cortex. For example, Blamire et al. (1992) presented subjects with visual stimuli lasting 2 seconds, intended to elicit correspondingly brief periods of neuronal activity in visual cortex. As expected, transient increases in the BOLD signal were observed in response to each stimulus. Likewise, Bandettini (1993, 1999) observed the effects of brief versus prolonged motor activity, measuring signal changes in motor cortex

while subjects made finger tapping movements lasting from 0.5 - 5 seconds in duration. Signal changes began around 2 - 3 seconds after finger tapping began, and continued for around 3 - 5 seconds, before returning to baseline. Increases in signal of around 2% in magnitude were seen for even the briefest movements. Figure 11.2 shows an extreme example of the sensitivity of the BOLD signal, demonstrating that visual stimulation as brief as 34 msec in duration will elicit small, but clearly detectable, signal changes (cf. Savoy et al. 1995; and see Konishi et al. 1996).

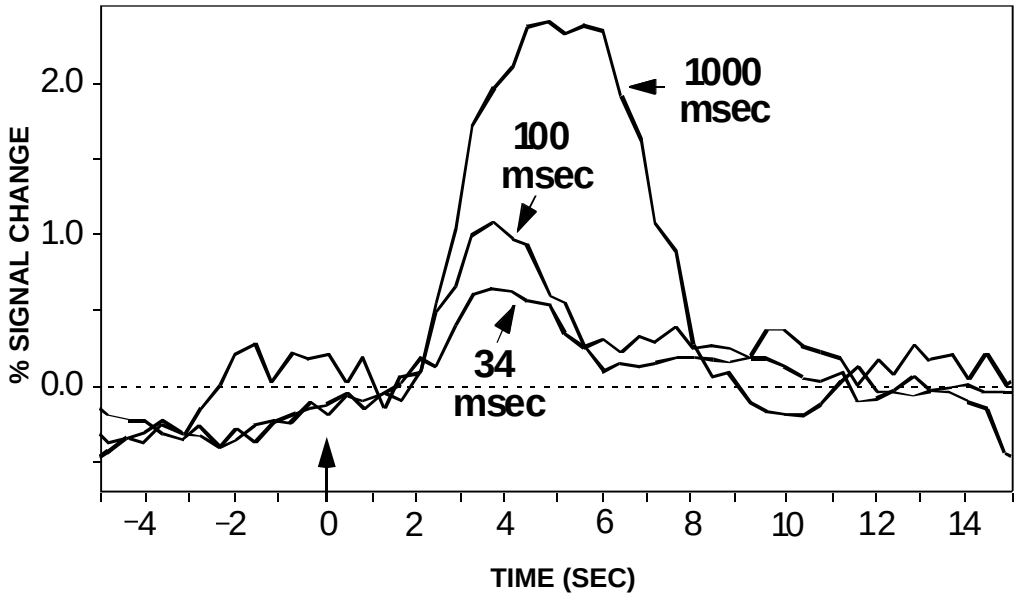

Fig. 11.2: Sensitivity of the BOLD signal. The fMRI BOLD signal change to visual stimulation. The stimuli varied in duration, lasting either 34, 100 or 1000 msec. Even the briefest period of stimulation gives rise to clear signal change. The signal change clearly increases in both amplitude and duration as the period of visual stimulation increases. Moreover, the signal is both delayed in onset and temporally blurred relative to the visual stimulation. Adapted from Savoy et al. (1995).

Importantly, similar transient signal increases have also been demonstrated in cognitive task paradigms, even though the signal changes can be considerably smaller than in studies using sensory and motor tasks (i.e., with observed signal changes of less than 1%). For example, Buckner et al. (1996) demonstrated that signal changes in visual and prefrontal brain areas can be detected during isolated trials of a word generation task. Similarly, Kim et al. (1997) defined hemodynamic responses in motor and visual cortex with tasks involving subject-initiated motor preparation. McCarthy et al. (1997) were able to measure transient responses in parietal cortex during infrequent presentation of target letter strings. Notably, all of these initial studies employed the constraint that trials were widely separated in time, allowing the hemodynamic response from one event to fully restore before the beginning of the next event.

Dale and Buckner (1997) extended the use of event-related procedures to circumstances in which different classes of stimuli were intermixed and presented close together in time. Dale and Buckner employed simple visual stimuli (flickering checkerboards), ensuring that relatively large and robust signal changes could be observed. In two experiments they showed that reliable signals could be detected in visual cortex when a single class of stimuli was presented rapidly – brief full-field flickering checkerboards, with as little as 2 second separation between successive stimulus presentations. Importantly, the response to these rapidly presented stimuli summed nearly linearly, so the individual responses to sequential events could be estimated. In two further studies Dale and Buckner demonstrated that reliable signals could be detected when two classes of stimuli were

intermixed (stimuli were randomly presented to either left or right hemifield). Even though stimuli were randomly presented to left and right hemi-fields in too rapid a succession for the hemodynamic response to return to baseline, it was possible to extract the same lateralised patterns of visual cortex activity seen with much longer inter-stimulus intervals. Collectively, these observations made clear that fMRI can detect hemodynamic responses to extremely brief, rapidly presented and randomly intermixed neuronal events, making it possible to employ event-related signal-averaging procedures (for review see Rosen et al. 1998). The principals and specific procedures by which the hemodynamic response can be estimated in trains of rapidly presented stimuli will be discussed below.

11.4.3 Linearity of the signal

The study by Dale and Buckner (1997) discussed above and a seminal study by Boynton and colleagues (1996) highlight the third notable feature of fMRI data that proved crucial for the introduction of rapid presentation event-related procedures: namely that the shape of the BOLD hemodynamic response to a given period of stimulation is predictable and relatively stable across events, even when there is an overlap in the responses to successive events.

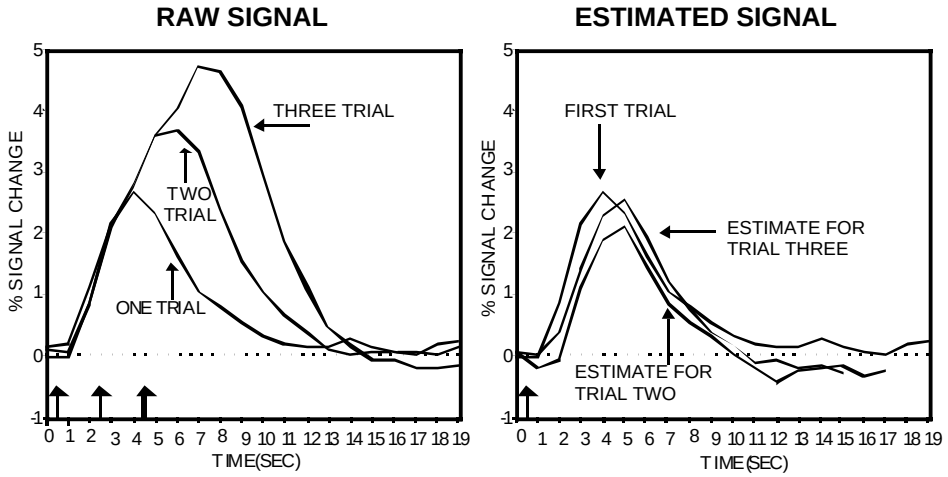

Fig. 11.3: Linear summation of the BOLD signal. Event-related fMRI data show approximately linear summation of the hemodynamic response in early visual cortex for closely spaced trials (2 sec apart) of a 1 sec visual checkerboard stimulus. **Left.** The raw fMRI signal intensity elicited by the presentation of one, two, or three trials. Note that the response is increased and prolonged by the addition of multiple trials, but does not saturate as successive trials overlap. **Right.** Estimates of the individual response to each of the three trials. Estimates were obtained by subtraction between the trial conditions (i.e., estimates for the second trial are the one trial condition subtracted from the two-trial condition, estimates for the third trial are the two-trial condition subtracted from the three-trial condition). Although subtle departures from linearity can be seen, the estimates are clearly very similar, supporting the assumption that the response to each stimulus summates linearly, at least under these specific conditions. Adapted from Dale and Buckner (1997).

Of particular interest is the finding that roughly the same response will be evoked even if the response to an initial event has not decayed, i.e., that the hemodynamic response summates in a roughly linear fashion (cf. Dale and Buckner 1997; Boynton et al. 1996; but see Vasquez and Noll 1998; Friston et al. 1997; Miezin et al. 2000). As can be seen in figure 11.3, fMRI signals evoked by a single event (1 second flickering checkerboard) exhibit the characteristic impulse-response function, with the hemodynamic responses to subsequently presented trials additively superimposing. Moreover, figure 11.3 shows that the estimated response to each successive trial closely matches that of the first trial with an onset around 2 seconds post-stimulus, a peak of between 4 and 6 seconds, and a duration of around 12 seconds. This finding is important, because linear models of the hemodynamic response

function underlie both the proper analysis of fMRI data, and the logic of certain forms of experimental designs such as rapid stimulus presentation and parametric manipulations (e.g., see Braver et al. 1997).

The even stronger claim of 'time-intensity separability' has been made in support of event-related procedures. This claim is that variations in the magnitude and duration of activity are essentially independent (Boynton, et al. 1996). This means that changes in the intensity (strength) of neural activity result in equivalent changes in the intensity (magnitude) of the hemodynamic response, without any concomitant change in the duration of the response. In addition, a given change in the duration of neural activity results in an equivalent change in the duration of the hemodynamic response, without concomitant change in the magnitude of the activity.

Whilst the linearity and time-intensity separability assumptions may be valid working assumptions in many cases, it should be remembered that evidence suggests that there are limits to these assumptions. In particular, it seems possible that certain changes in the intensity of experimental stimulation are likely to break the linearity assumption – indeed even some of the earliest investigations of the BOLD signal showed some non-linear behaviours. In fact, all studies showing roughly linear summation properties have also revealed evidence for subtle (or not-so-subtle) nonlinearities. For example, Bandettini (1993) varied the rate at which subjects made finger-tapping movements, demonstrating that the magnitude of the signal change in motor cortex did not vary in a monotonic fashion, i.e., it was not linearly related to the frequency of finger tapping.

One difficulty in resolving questions about how the hemodynamic signal summates over time is the fact that in many situations it is not known whether the underlying neuronal activity is itself linearly additive across time and trials. Thus, it is unclear whether departures from linearity reflect an intrinsic non-linear property of the hemodynamic response or of the underlying neuronal activity. For example, auditory word stimuli have been shown to exhibit roughly linear responses when stimuli were presented as frequently as one per 2 sec or slower, but robust non-linearities in the response are observed at higher stimulus presentation rates (Friston et al. 1997). It may be the case that the neuronal response to auditory words at such rapid rates is different to that of more widely spaced words; alternatively, the neuronal response to words may be constant across rates but the hemodynamic response may saturate.

A practical question to ask – which to some degree is agnostic as to the underlying mechanism by which nonlinearities might arise – is whether the same, or different estimates of the hemodynamic response would be obtained at different trial presentation rates. Would the same response estimate be obtained when trials are spaced widely apart as when trials are so closely spaced as to yield overlapping responses? This was one of the questions asked in a recent analysis presented by Miezin et al. (2000). The rate of trial presentation was varied from one trial every 5 seconds (on average) to one trial every 20 seconds. Similar response function estimates were obtained at all rates. However, an amplitude reduction of about 20% was noted at the fastest trial presentation rates, again suggesting either modest nonlinear summation in the hemodynamic response itself, or interactions between trials. Neither the time-to-onset of the response, nor time-to-peak of the response was affected by the rate of trial presentation. Moreover, because of the increased numbers of trial events at the fastest rate, the greatest statistical power was present for detecting a response. Thus, in practical terms, these findings suggest that the hemodynamic response adds in a sufficiently linear fashion to provide a statistically powerful means of estimating responses – the necessary foundation for rapid event-related paradigm procedures.

11.5 Paradigm design and analysis of event-related studies

Hemodynamic imaging was originally limited to the use of 'blocked' recording procedures. As Raichle (in press) points out in an historical review of neuroimaging, the use of blocked

procedures was a compromise, and the introduction of event-related procedures represents a major advance. Why? Fundamentally, because it allows the same procedures to be employed that have been so successfully employed elsewhere in the cognitive neurosciences. Enthusiasm for the 'event-related' approach stems from the fact that cognitive neuroscientists have traditionally employed procedures that allow for measurements of individual trials, or even sub-components of trials. Consider reaction time data, for example. Typically, reaction times are measured for individual trials, and then averaged according to different response categories (based on the type of stimuli presented and subject responses), rather than being averaged across the entire experimental session. Equivalent procedures can now be easily employed in fMRI studies.

11.5.1 Extracting the signal of interest

From a methodological viewpoint, event-related and blocked procedures differ in the means by which the signal of interest is extracted. Blocked paradigms integrate activity over extended periods of up to a minute during which a whole series of experimental trials (or a continuous extended task) are presented. By contrast, event-related procedures achieve similar signal-to-noise improvements by averaging data according to stimulus trial types, revealing activity that is time-locked to the onset of individual trial events. In most common settings these two forms of analysis largely proceed independently and are used to resolve different forms of signal. As will be discussed below however, in certain situations a given experimental data set can be analyzed both in a blocked fashion and as part of an event-related analysis.

In blocked paradigms the time-averaged signal from one group of sequential trials is contrast with the time-averaged signal from another group. The signal reflects an average response from all trials presented during the task block – the signal changes associated with individual trials or trial types are not estimated separately. Moreover, averaged signal may also reflect multiple sources of activity. For example, imagine a visual brain region that increases in activity when a subject is asked to make judgements on visual stimuli, and further increases in activity each time an individual visual stimulus occurs during the task. Two separate components of the signal would likely be present differing in their temporal evolution: (1) a sustained signal change across the entire task block and (2) a transient signal change with onset in a fixed relation to each trial event. In a blocked paradigm in which activity is averaged over an extended period the sustained and transient signal changes would not be resolved.

Event-related analyses explore signal changes in relation to the onsets of individual trial events, taking advantage of the finding that the BOLD hemodynamic response produces a transient (if temporally blurred) response to isolated trial events. The strength of the approach is two-fold. First, transient signal changes that are not sustained across multiple trials can be detected. Second, because the transient signals are determined by individual trials, different kinds of trials can be randomly intermixed (or determined post-hoc) for comparison. This allows the use of the kinds of paradigms that have been so well developed in other disciplines of cognitive neuroscience, such as for evoked potential studies. This allows the response to a specific event of interest to be defined. Note however, that a consequence of event-related procedures is that (as typically applied) they will likely miss activity that spans across successive trials – an issue that we will return to below in the section "Introducing mixed 'event-related and blocked' designs". First however, we provide a more extensive discussion of event-related procedures and their use in identifying transient signal events.

11.5.2 Time-locked averaging

As described above, event-related signal averaging most often requires the repetition of experimental trials, such that repeated time-locked epochs of data can be recorded and subsequently averaged together (cf. Buckner et al. 1996; Clark et al. 1998; Dale and Buckner 1997; Zarahn et al. 1997; Friston et al. 1998a; Josephs et al. 1997). This, by definition, requires that an event is repeatable – or more commonly, that multiple instances of a class of event can be presented (such as old or new items in a memory test). The averaging procedure also requires that the data can be acquired in such a fashion that they can be aligned with the event of interest (i.e., averaged together based on a consistent reference point). As is illustrated in figure 11.4, in a typical event-related paradigm the presentation of each experimental stimuli is used as the temporal event to which the data is time-locked. However, this need not be the case. It is possible to time-lock to other events, such as the behavioral response made to each stimuli, and (at least in principle) to physiological measures such as heart rate. Moreover, in some circumstances it may be that one is interested in the pattern of neural activity leading up to an event, defining the event of interest as the end rather than the beginning of a trial.

In the most straightforward case the analysis of event-related reduces to simple selective averaging, i.e., calculating the mean and variance of the fMRI signal at each time point for each kind of trial event (e.g., Buckner et al. 1996; Dale and Buckner 1997). Inferential statistics can then be employed to ask questions about the presence or absence of differences between hemodynamic responses for each type of event (equivalent to the analysis of other physiological and behavioral data). Importantly, because of the temporal span of the hemodynamic response, it is necessary to consider the data across a range of time points. Whilst averaging procedures are typically performed off-line, alternative on-line procedures are currently being developed that allow real-time data processing and even statistical analyses (cf. Voyvodic 1999; Posse et al. 1998; Cox et al. 1995). Either way, given that the signal of interest is systematically associated with the time-locked event and invariant across trials, while background noise is random, the averaging procedure increases the signal-to-noise for signal changes that are time-locked to the experimental event. Clearly, the greater the number of trials that are averaged together the higher the signal-to-noise ratio becomes.

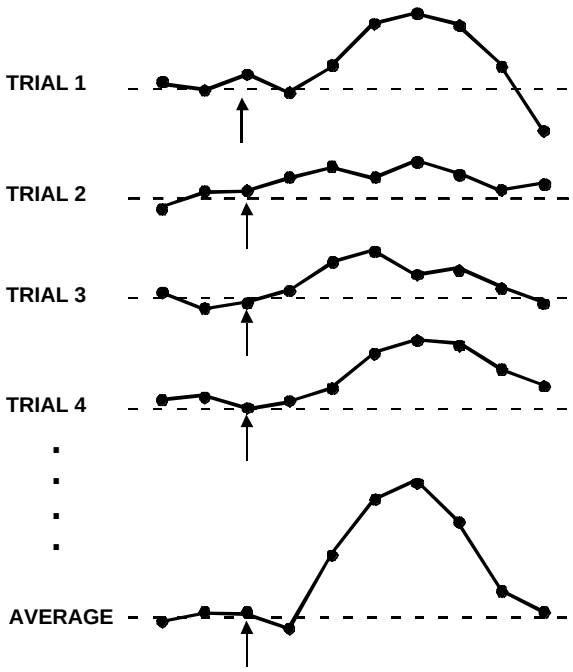

Fig. 11.4: Event-related averaging. The figure shows how averaging can be used to reveal hemodynamic response that is consistently related to an event of interest. A series of trials are presented, involving the presentation of a stimulus on each trial, and hemodynamic changes are measured time-locked to the presentation of each stimulus. Individual trials are then averaged together to provide a representative average signal. Note that the signal on any given trial contains both the signal of interest and noise, and thus need not appear highly similar to the average signal. Nonetheless, assuming that noise is random across trials, the average signal reveals that part of the hemodynamic response that is systematically related to the event of interest.

More recent approaches to the analysis of event-related fMRI data now involve a full implementation of the general linear model (GLM; cf. Friston et al. 1994; Worsley and Friston 1995; Josephs et al. 1997; Zarahn et al. 1997; Miezin et al. 2000). Analysis within the GLM is rooted in the simple assumption that the variance in the evolving fMRI BOLD-contrast signal that is systematically time-locked to the event is a direct measure of the hemodynamic response. In practice, analysis of event-related data within the GLM requires that an explicit model is generated of the factors (i.e., effects) that are thought to contribute to variability in a data set. Effects can be modeled to account for the different kinds of trial events, as well as confounding effects such as a mean run intensity or slope that are theoretically uninteresting, but nonetheless present in the data.

For each voxel in a data set, estimates of the response to each effect are calculated by representing each time point in the data set by a linear equation. For each time point in the data set the linear equation represents the measured BOLD signal as the sum of the hemodynamic responses occurring at that point plus variance from noise (as illustrated in figure 11.5).

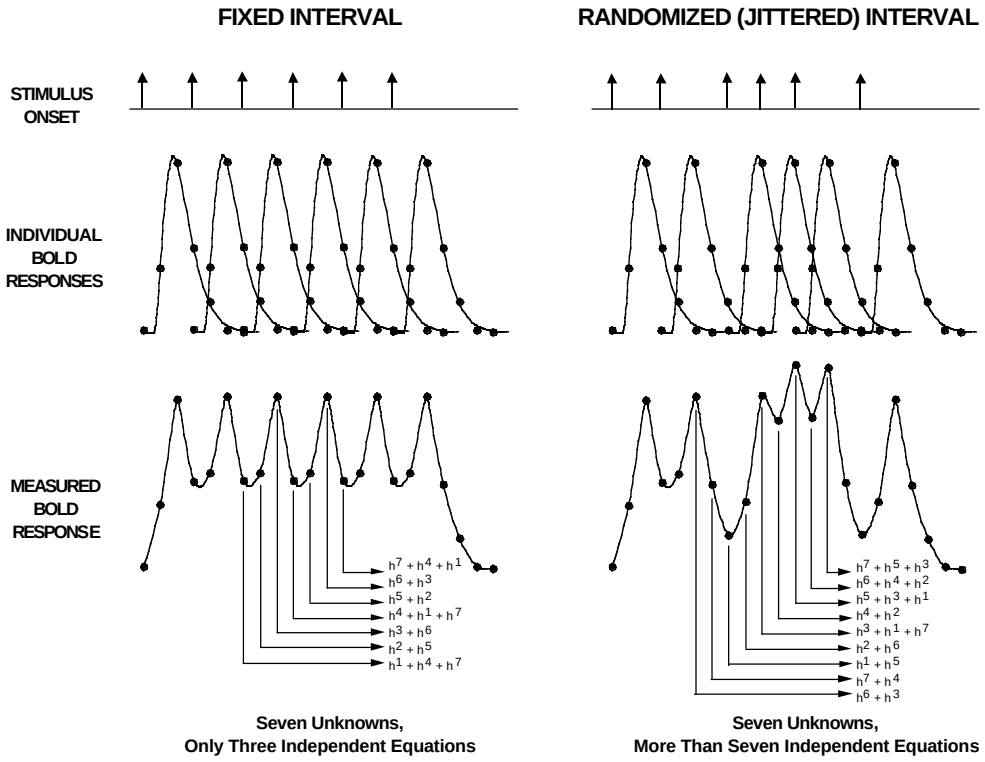

Fig. 11.5: The general linear model. The general linear model can only effectively estimate the separate contributions of the different trial types to the measured waveform when there are more equations than unknowns. The use of linear equations to estimate the hemodynamic response of overlapping events is consequently dependent upon the introduction of jitter between successive trial types. Each stimulus is associated with a transient signal change, evolving over several time points. The measured fMRI signal will be the linear summation of the signals to each stimulus – a complex waveform. The independent contributions of the different trials are unclear in the summated signal, however the contributions can be estimated using general linear equations. **Left.** When a fixed order of stimuli is employed in rapid event-related designs, with a fixed interval between successive trials, the resulting fMRI signal is repetitive. Analytically, the waveform is represented by fewer equations than unknowns, rendering the model unsolvable. **Right.** By introducing temporal jitter (variation) into the presentation rate of the trials the measured waveform becomes more complex. Consequently additional equations are required to represent the waveform, producing more equations than unknowns, making the general linear model tractable. Adapted from Miezin et al. (2000).

Because of overlap of successive hemodynamic responses to closely spaced trials, the ability to estimate the BOLD response is highly dependent upon the time between trials (cf. Burock et al. 1998). An important feature of the GLM approach is that the linear equations representing a data set can only be solved (i.e., estimates can only be derived) if there are as many equations as there are unknowns. When there is a fixed spacing between successive trials there is not enough information in the time-course data to arrive at a unique solution; there are more unknowns than equations, and the model is unsolvable. However, the introduction of "jitter" (i.e., a variable temporal delay between successive trials) significantly increases the variance within the data set, providing more information from which to derive estimates of the BOLD response. The effect of introducing jitter is illustrated in figure 11.5.

The introduction of the GLM allows sophisticated approaches to data analysis to be employed, such as interactions between response categories, time and performance measures (e.g., reaction time). Data can be selectively analyzed for multiple independent factors (e.g., different response categories), extracting out that part of the on-going signal that is associated with each factor. Moreover, as will be discussed below, different factors that contribute to the ongoing waveform can be separated out using the GLM, allowing transient stimulus-related activity to be dissociated from more sustained state-related activity.

11.5.3 Trial ordering and jitter

When trials are spaced widely in time, the ordering of trial types has a marginal affect on the power of a paradigm. However, when trials are rapidly presented in succession, and must be linearly estimated, the trial ordering and/or temporal jittering between trials becomes a critical factor in designing a study. For example, Dale and Buckner (1997) provided a simple selective averaging procedure that allowed the separate contributions of different trials types to be estimated when they overlapped in time. It is important to stress that the power of the selective averaging procedures used by Dale and Buckner (1997) derived in part from the counterbalancing procedures employed. In their design, all trial types were made to follow each other equally often. That is, if trial type "A" was followed by a "B" (yielding A-B), it was also followed by equally often an "A" at another point in time (yielding A-A). When differences between event types were considered, the history of overlap simply therefore subtracted out.

The more general form of this counterbalancing technique is straightforward: trial orders are organized such that each type of trial is followed and preceded by every other trial type equally often. Because of the potential complexity of this, a computer program can be used to generate such trial sequences. The consequence of such counterbalancing is that when differences between trial types are considered the overlap between adjacent hemodynamic responses cancels out by simple subtraction. Within the context of the simple signal averaging procedure employed by Dale and Buckner (1997), without the use of this form of counterbalancing the overlapping hemodynamic responses to each successive stimuli are mixed together, making it difficult to assess the response for each type of stimuli.

It is not always possible to counterbalance stimuli in this way, however, when interest lies in unpredictable responses to stimuli that may exhibit temporal dependencies or, by chance, yield sequences that are not counterbalanced. More generally, whenever trials are sorted post-hoc according to subject performance it is unlikely that different trial types will be perfectly counterbalanced when only small numbers of trial events are considered. In such circumstances there are two simple solutions. First, if a signal averaging procedure is used, then one solution is to change the experimental paradigm by spacing successive trials sufficiently far apart (i.e., around 10-12 seconds) for the hemodynamic response to return to baseline between trials (see Bandettini et al. 1998). This solution is probably not ideal as it will be of low power because relatively few trials can be presented in an acceptable time period. A second option is more appealing because it involves changing the analysis, rather

than the paradigm. With this approach a more sophisticated analysis procedure is employed to obtain an overlap-corrected estimate of the hemodynamic response. Estimation within the GLM provides a powerful strategy. Although the GLM approach may provide a better estimate when stimuli are fully counterbalanced, it does not require the stringent counterbalancing procedures upon which the signal averaging technique is based. Power will increase to the degree that a variable amount of jitter exists between reoccurrences of the same trial type, which is near its maximum in the design of Dale and Buckner (1997). However, a considerable amount of power will exist so long as there is sufficient jitter between reoccurrence of the same event type.

The significance of introducing jitter between the representations of the same trial type is clear when one considers the difference between procedures that have a single trial type presented at fixed positions in time (i.e., without jitter) and those in which the trial is represented at varying intervals. When the interval between trials is fixed, very long intervals allow the hemodynamic response to return to baseline, producing some variance in the signal. By contrast, at very short intervals the response reaches a plateau because successive responses systematically overlap and never returning to baseline, which reduces any variance in the signal. Put another way, the overlap cancels any modulation in signal; as the response increases for the next event, it will obscure the return to baseline of the earlier events. However, when the interval between trials fluctuates on any given trial there may be either a large or small amount of overlap between the hemodynamic response elicited on successive trials, depending on whether or not the trials were close together. Thus, as the presentation rate increases so does variability in the amount of overlap between trials, and hence variance in the signal will be larger (Burock et al. 1998). The greater this variance, the greater the information for deconvolution of individual time courses.

The foregoing discussion should make clear that with rapidly presented stimuli it is necessary to introduce fluctuations in the interval between successive trials. Jitter increases the variance in the signal and thereby increases the amount of information available, thus allowing the underlying hemodynamic response to be estimated for each class of trial. In practical terms this feature of the event-related procedure is easy to incorporate into any given experimental design by simply introducing randomly inter-mixed ‘null trials’ – such as the presentation of a blank screen or fixation point during a visual experiment. From the subject’s perspective the gap between trials simply appears to vary slightly.

11.5.4 Detecting the response and constructing activation maps

A critical consideration when examining the signal using event-related procedures lies in estimating the time course of the signal, so that the averaged data includes the full evolution of a potential hemodynamic response, and then using that estimate to detect any hemodynamic response. The estimation procedures discussed above can be used without incorporating an explicit mathematical representation of the properties of the hemodynamic response. Linear estimation based on selective averaging or estimation within the GLM does not assume a specific response shape; only when statistical maps or quantitative variables must be derived are response shapes assumed.

Statistical map generation depends on the ability to detect those voxels in which a hemodynamic response is present. The most straightforward approach is to make no assumptions about the form (such as shape and temporal properties) of the hemodynamic response by examining for an effect over time (cf. Cohen et al. 1997). The disadvantage with this however, is that it is potentially less powerful than approaches that make use of the well-known characteristics of the hemodynamic impulse-response function. However, in using a procedure that assumes a specific response shape and extracts activity that conforms to the particular properties of the hemodynamic response (e.g., using cross-correlation methods),

the sensitivity will depend on what the shape of the hemodynamic response actually is, relative to what it is modeled as.

While general characteristics of the response can be defined, as described above the specific shape of the response may vary in timing considerably between different regions of the cortex (Lee et al., 1995; Schacter et al. 1997; Buckner et al. 1996; Robson et al. 1998). Thus, when an idealized hemodynamic response is employed as the basis for analysis, it is important to employ a range (or set) of functions, to account for differences in the onset of the hemodynamic response across brain regions. Analysis that is restricted to a single function may systematically ignore activity that does not conform to the behavior of the chosen hemodynamic response (cf. Buckner et al. 1996; Schacter et al. 1997, discussed above). Fortunately, such variability can be relatively easily accounted for in practice. Multiple hemodynamic responses can be employed, using multiple delays, providing some robustness in detecting experimentally induced activity when a basic shape is assumed.

11.5.5 Some consequences of being event-related

Clearly, event-related averaging is a powerful technique for extracting the signal-of-interest. As with any method however, in addition to the inherent strengths, it has weaknesses. One clear caveat is that the averaged signal is not a direct measure of the response occurring on individual trials, and thus the averaged waveform could potentially show little relation to those of individual trials. The possibility of such distortion through averaging is important, because it leads to uncertainties in interpretation. For example, consider two averaged responses from different experimental conditions for a single brain region. If the magnitude of activity is greater in one case than the other this could be taken to reflect the presence of an underlying neural process that is graded, i.e., with greater processing in one case than the other. This need not necessarily be the case, however. The averaged responses also could reflect activity associated with an all-or-none neural process, with the difference between the two conditions lying in the proportion of trials in which the process is active.

Of course, this is not a problem specific to neuroimaging data – similar concerns arise in the interpretation of any dependent measures when averaged events are considered, whether behavioral or physiological. Nonetheless, such concerns reinforce the importance of studies examining single trial data (cf. Kim et al. 1997; Richter et al. 1997) or the distributions from many individual trials. Rapid event-related designs, in particular, are not easily modified to estimate the distribution of responses underlying the average. Estimation using parametric statistical methods such as the GLM implicitly ignores the possibility that responses may not come from a single uni-modal distribution.

Notwithstanding concerns over the interpretation of differences between averaged event-related signals, it is clear that the introduction of event-related procedures has a major effect on the way fMRI data can be analyzed. Perhaps the most significant effect of the event-related method is that it becomes possible to analyze data from different experimental trials according to either the type of stimuli presented, or subjects' responses to stimuli. Such approaches allow the neural correlates of performance on each experimental trial to be sorted post-hoc. The advantages of this should be clear. For example, in studies of memory where the focus of interest lies in memory errors that are likely to occur relatively infrequently, the neural correlates of interest may only be visible using a trial based approach, where infrequently occurring memory errors can be cleanly isolated.

Studies of 'subsequent memory' effects highlight the utility of post-hoc sorting. Wagner et al. (1998; also see Brewer et al. 1998) used event-related methods to study the neural correlates of memory encoding, by post-hoc sorting hemodynamic response estimates according to whether subjects remembered each item during a later memory test. It was found that differences in the activation response associated with individual items at presentation were predictive of subjects' subsequent memory for those items. It is worth noting that the

ability to sort data extends to other aspects of performance, such as reaction time (and in principal, to physiological measures such as heart rate). Thus, the use of event-related procedures allows one to ask questions concerning the covariation between neural activity and measurable aspects of behavior that are also likely to differ on a trial by trial basis.

Finally, the ability to sort trials post-hoc also introduces the possibility of rejecting artifact-laden trials. It is clear that there are circumstances in which it is desirable to exclude responses to certain types of trial – for example based on performance measures such as outlier responses (e.g., as measured by response times), incorrect responses, or when large movement artifacts are present. ER-fMRI procedures make this possible by allowing investigators to determine post-hoc which trials contain artifacts and exclude them from further analysis. This kind of artifact rejection seems likely to be of particular use in studies involving difficult population groups, such as in experiments involving children (e.g., see Thomas et al. 1999; Logan 1999) or clinical patients (e.g., see Jezzard and Song 1996). Note that with blocked procedures this form of analysis is necessarily precluded because neural activity is averaged over a series of successive trials.

11.5.6 Issues related to experimental design

Not all of the effect of employing event-related procedures relates to data collection and analysis issues; the procedure also directly affects experimental design. The introduction of event-related procedures makes it more meaningful to inter-mix different types of stimuli. That is, to employ randomised rather than blocked experimental designs. A study by McCarthy et al. (1997) highlights the way in which being able to employ randomized designs extends the range of questions that can be addressed using fMRI. The aim of the study was to examine the neural correlates of relatively infrequently occurring (rare) events, occurred in amongst a series of frequently occurring (common) events. By averaging data in a time-locked manner dependent on each class of event, they were able to show that the infrequently occurring events elicited transient increases in activity in prefrontal and parietal cortex. This could not have been shown using blocked procedures however, because activity associated with each type of event could only have been separated had they been presented in separate blocks – making it difficult to have infrequent (rare) events.

Another clear example of the potential problems associated with this comes from studies of memory retrieval where old (studied) and new (unstudied) test items are presented in separate test blocks. After considering the first few trials subjects will be easily able to respond correctly to others in that block, without needing to consider whether the items are actually old or new. Employing a randomized experimental design can not only prevent subjects from suffering from fatigue and boredom, but also stop them from employing such strategies. However, it is worth stressing that the use of blocked designs in neuroimaging does not necessitate that trials of a certain type are grouped together. Indeed, researchers studying memory retrieval have purposely mixed old and new items at test, parametrically varying the proportions of each trial type across blocks (e.g., see Rugg et al. 1996). The problem with this is that when the neural activity is measured across the block it is not possible to distinguish specifically the differences that occur between the different types of trials that are presented. This is only possible when event-related procedures are employed.

While the preceding discussion illustrates the potential problems of employing blocked data analysis procedures, it should not be taken to suggest that the displaying stimuli in blocks is undesirable per se. Indeed, in some circumstances a blocked experimental design is more appropriate. For example, requiring patients with frontal lobe damage to constantly alternate between tasks on a trial by trial basis may simply interfere with their ability to perform the tasks at all.

Unfortunately, it is not possible to simply produce an absolute index of which regions are involved in performing a given cognitive process – regardless of whether blocked or

event-related procedures are employed. As with all neuroimaging techniques, information is provided by examining relative changes in activity – either in terms of differences between pairs (or a series) of tasks, or in correlations with behavioral measures.

11.6 Introducing mixed 'event-related and blocked' designs

A relatively subtle facet of paradigm design, and one that has been largely ignored to date, is the sensitivity of different kinds of paradigms to signal changes that evolve on distinct time scales. By definition, the event-related designs discussed above are sensitive to transient changes in activity that are time-locked to events of interest. Blocked paradigms are also sensitive to this form of signal change – but average over the events within the block. That is, a series of transient events will sum into a single extended signal change in standard blocked paradigms. However, as noted above, blocked paradigms are additionally sensitive to sustained changes in activity that exist across extended periods of time and that are not necessarily modulated on an event-by-event basis. In this regard, blocked and event-related procedures may be employed to investigate quite different questions.

The contrast between the use of event-related and blocked designs to investigate signal changes that develop over very different time scales is clear in two recent studies by Henson and colleagues, investigating different aspects of memory retrieval. In one case (Henson et al. 1999a) event-related procedures were used because the focus of interest lay at the level of individual test items, i.e., in differences between individual recognition responses associated with trial-by-trial difference in experience. In a second study (Henson et al. 1999b) blocked procedures were employed because interest lay in defining any differences that span across all test items. That is, differences between task blocks that differed in the retrieval instructions given to subjects – differences that might be expected to yield a sustained activity change reflecting the cognitive set (or mode) that a subject was engaged in.

The two experiments by Henson and colleges described above are interesting because they elucidate an important feature of experimental tasks; namely that performance on a task usually involves a combination of set-related processes and stimulus-related processes. Subjects are given an instructional set, and then use that set to dictate how they respond on each individual trial. A reasonable hypothesis is that certain components of the set-related responses are sustained while the set is maintained, while transient responses may modulate in response to the individual stimulus trials. That is, processes may exist that are elicited by the general situation or task setting (state), alongside processes that are elicited by the specific stimuli that are presented (item). Figure 11.6 provides an example of a paradigm in which both transient and sustained signal changes may be present.

As discussed above, blocked trial paradigms generally confound these two types of response (although for an interesting exception using Event-Related Potentials see Duzel et al. 1999). In a typical blocked design fMRI experiment it is not possible to distinguish sustained changes in activity (effects that are tonic, lasting throughout a task period) from transient changes in activity (effects that are stimulus-specific and associated with current mental processes). Thus, any comparisons made between blocks will confound sustained and transient changes in activity, unless explicit attempts are made to control one or other type of response across blocks. By contrast, event-related paradigms are typically designed to be sensitive only to transient signal changes. Any sustained (tonic or task-related) activity that is present in all of the event-related responses during a task will therefore be subtracted out when comparisons are made between the responses to different conditions.

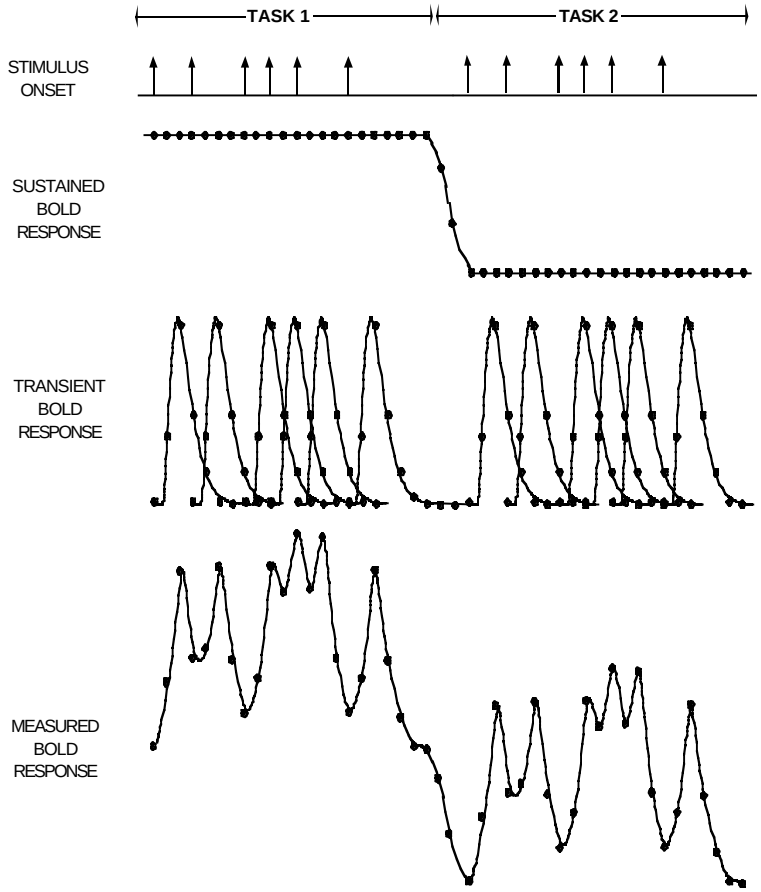

Fig. 11.6: Paradigm design for fMRI. A representation of the signal measured in a mixed 'blocked and event-related' paradigm. Stimuli are presented in two separate blocks (task 1 and task 2). Two distinct kinds of correlates of activity may occur, a sustained BOLD response that is modulated across the two tasks, and a transient BOLD response that is associated with the individual stimuli. The measured BOLD response will be a linear summation of these different signal changes. If blocked analyses were employed then a single aggregate measure of transient and sustained activity would be provided over each task period, confounding transient and sustained signal changes. By contrast, if event-related analyses were employed then only the transient signal changes would be examined. However, the general linear model can be used to provide independent estimates of i) sustained activity associated with the different tasks, and ii) transient signal changes associated with individual stimuli. This kind of paradigm design can be employed to distinguished between transient and sustained changes in activity, for example allowing different task states such as full versus divided attention or deep versus shallow encoding to be compared whilst holding the individual items constant.

Of course, in many circumstances it may be that one is interested in neural activity associated with the processing of experimentally presented stimuli, and that sustained, ongoing set-related activity is essentially unwanted noise. However, in other settings, where one is explicitly manipulating the presentation of individual stimuli and the context within which this is done, the distinction between sustained and transient effects can be crucial. For example, in studies of memory retrieval, neural activity associated with being in a 'retrieval mode' (related to the general requirement to perform a memory test) may be distinct from that associated with attempting or succeeding to retrieve information from memory (related to individual test items). Similarly, in the study of retrieval monitoring by Henson et al. (1999b) the effect of task states would ideally have been separated from the effects of the individual items presented during each task period. A thorough examination of the neural basis of performance thus necessitates investigations that allow a hybrid or combination of blocked and event-related procedures to be employed, allowing sustained or tonic neural effects to be distinguished from transient or item specific effects.

One of the exciting aspects of recent approaches to analysing fMRI data is that mixed 'event-related and blocked' analyses can be easily implemented. For example, the GLM approach allows multiple independent factors to be coded into the analytical model that is used to represent the experimental design. Thus, as long as an experiment is designed such that individual trials are separated into distinct blocks, transient (i.e., stimulus related) and sustained (i.e., set-related) effects can be simultaneously incorporated into the GLM. Figure

11.6 provides one example of how an experiment can be designed to incorporate both blocked and event-related features.

When analyzed using the GLM this kind of design allows sustained signal changes associated with the cognitive set to be estimated alongside independent estimates of the transient responses to individual stimuli. Moreover, with this type of paradigm design it should be possible to investigate the extent to which tonic or sustained neural activity contributes to, or interacts with, transient stimulus specific effects. For example, it may be that the pattern of transient activity associated with a given stimulus type is modulated (increased or decreased) dependent upon the subjects cognitive set. It is worth remembering however, that it is the introduction of event-related procedures that makes such a goal achievable. Thus, a detail that could be easily overlooked is that in a mixed blocked and event-related design, the trial types within the blocks must be jittered just as they would be in a purely event-related design. If that constraint is met, then solving for the independent sustained and transient components of the signal is simply the simultaneous solution to a set of effects that includes both the individual events and the blocks. Only with ‘mixed’ designs however, is it possible to investigate the role of both transient and sustained activity in supporting cognitive processing, and any potential interplay between the two.

References

Aguirre, G. K., Zarahn, E., and D'Esposito, M. (1998). The variability of human BOLD hemodynamic responses. NeuroImage, 8, 360-9.

Bandettini, P. A. (1993). MRI studies of brain activation: Dynamic characteristics. In Functional MRI of the Brain, pp.144-51. Society of Magnetic Resonance in Medicine, Berkeley.

Bandettini, P. A. (1999). The temporal resolution of MRI. In Functional MRI (eds. C. T. W. Moonen and P. Bandettini), pp 205-20. Springer-Verlag, Mauer, Germany.

Bandettini, P. A., and Cox, R. W. (1998). Contrast in single trail fMRI: Interstimulus interval dependency and comparison with blocked strategies. Proceedings of the International Society for Magnetic Resonance in Imaging, 6th Meeting, 1, 161.

Bandettini, P. A., Wong, E. C., Hinks, R. S., Tikofsky, R. S., and Hyde, J. S., (1992). Time course of EPI of human brain function during task activation. Magnetic Resonance in Medicine, 25, 390-7.

Binder, J. R., and Rao, S. M. (1994). Human brain mapping with functional magnetic resonance imaging. In Localization and Neuroimaging in Neuropsychology, (ed. A. Kertesz), pp. 185-212. San Diego: Academic Press.

Binder, J. R., Frost, J. A., Hammeke, T. A., Bellgowan, P. S. F., Rao, S. M., and Cox, R. W. (1999). Conceptual processing during the conscious resting state. A functional MRI study. Journal of Cognitive Neuroscience, 11, 80-95.

Blamire, A. M., Ogawa, S., Ugurbil, K., Rothman, D., McCarthy, G., Ellerman, J. M., Hyder, F., Rattner, Z., and Shulman, R.G. (1992). Dynamic mapping of the human visual cortex by high-speed magnetic resonance imaging. Proceeding of the National Academy of Science, USA 89, 11069-73.

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. Journal of Neuroscience, 16, 4207-21.

Braver, T. S., Cohen, J. D., Nystrom, L. E., Jonides, J., Smith, E. E., and Noll, D. C. (1997). A parametric study of prefrontal cortex involvement in human working memory. NeuroImage, 5, 49-62.

Brewer, J. B., Zhao, Z., Desmond, J. E., Glover, G. H., and Gabrieli, J. D. (1998). Making memories: brain activity that predicts how well visual experience will be remembered. Science, 281, 1185-7.

Buckner, R. L. (1996). Beyond HERA: Contributions of specific prefrontal brain areas to long-term memory retrieval. Psychonomic Bulletin and Review, 3, 149-58.

Buckner, R. L., and Koutstaal, W. (1998). Functional neuroimaging studies of encoding, priming, and explicit memory retrieval. Proceedings of the National Academy of Science, USA, 95, 891-8.

Buckner, R. L., and Logan, J. M. (in press). Functional neuroimaging methods: PET and fMRI. In Functional Neuroimaging (ed. R. Cabeza). Cambridge: MIT Press.

Buckner, R. L., Bandettini, P. A., O'Craven, K. M., Savoy, R. L., Petersen, S. E., Raichle, M. E., and Rosen, B. R. (1996). Detection of cortical activation during averaged single trials of a cognitive task using functional magnetic resonance imaging. Proceedings of the National Academy of Science, USA, 93, 14878-83.

Buckner, R. L., Koutstaal, W., Schacter, D. L., Dale, A. M., Rotte, M., and Rosen, B. R. (1998). Functional-anatomic study of episodic retrieval. II. Selective averaging of event-related fMRI trials to test the retrieval success hypothesis. NeuroImage, 7, 163-75.

Burock, M. A., Buckner, R. L., Woldorff, M. G., Rosen, B. R., and Dale, A. M. (1998). Randomized event-related experimental designs allow for extremely rapid presentation rates using functional MRI. NeuroReport, 9, 3735-9.

Buxton, R. B., Wong, E. C., and Frank, L. R. (1998). Dynamics of blood flow and oxygenation changes during brain activation: the balloon model. Magnetic Resonance in Medicine, 39, 855-64.

Carter, C. S., Braver, T. S., Barch, D. M., Botvinick, M. M., Noll, D.C., and Cohen, J. D. (1998). Anterior cingulate cortex, error detection, and the online monitoring of performance. Science, 280, 747-9.

Casey, B. J., Cohen, J. D., O'Craven, K., Davidson, R. J., Irwin, W., Nelson, C. A., Noll, D. C., Hu, X., Lowe, M. J., Rosen, B. R., Truwitt, C.L., and Turski, P. A. (1998). Reproducibility of fMRI results across four institutions using a spatial working memory task. NeuroImage, 8, 249-61.

Clark, V. P., Keil, K., Maisog, J. M., Courtney, S., Ungerleider, L. G., and Haxby, J. V. (1996). Functional magnetic resonance imaging of human visual cortex during face matching: A comparison with positron emission tomography. NeuroImage, 4, 1-15.

Clark, V. P., Maisog, J. M., and Haxby, J. V. (1998). fMRI study of face perception and memory using random stimulus sequences. Journal of Neurophysiology, 79, 3257-65.

Cohen, J. D., Perlstein, W. M., Braver, T. S., Nystrom, L. E., Noll, D. C., Jonides, J., and Smith, E. E. (1997). Temporal dynamics of brain activation during a working memory task. Nature, 386, 604-7.

Cox, R. W., Jesmanowicz, A., and Hyde, J. S. (1995). Real-time functional magnetic resonance imaging. Magnetic Resonance in Medicine, 33, 230-36.

Dale, A. M., and Buckner, R. L. (1997). Selective averaging of rapidly presented individual trials using fMRI. Human Brain Mapping, 5, 329-40.

Demb, J. B., Desmond, J. E., Wagner, A. D., Vaidya, C. J., Glover, G. H., and Gabrieli, J. D. (1995). Semantic encoding and retrieval in the left inferior prefrontal cortex: A functional MRI study of task difficulty and process specificity. Journal of Neuroscience, 15, 5870-8.

D'Esposito, M., Zarahn, E., and Aguirre, G. K. (1999). Event-related functional MRI: Implications for cognitive psychology. Psychological Bulletin, 125, 155-64.

DeYoe, E. A., Carman, G. J., Bandettini, P., Glickman, S., Wieser, J., Cox, R., Miller, D., and Neitz, J. (1996). Mapping striate and extrastriate visual areas in human cerebral cortex. Proceeding of the National Academy of Science, USA, 93, 2382-6.

Duzel, E., Cabeza, R., Picton, T. W., Yonelinas, A. P., Scheich, H., Heinze, H. J., and Tulving, E. (1999). Task-related and item-related brain processes of memory retrieval. Proceeding of the National Academy of Science, USA, 96, 1794-9.

Engel, S. A., Glover, G. H., and Wandell, B. A. (1997). Retinotopic organization in human visual cortex and the spatial precision of functional MRI. Cerebral Cortex, 7, 181-92.

Frackowiak, R. S. J., and Friston, K. J. (1995). Methodology of activation paradigms. In Handbook of Neuropsychology, Volume 10. (eds. F. Boller and J. Grafman), pp. 282-369. Amsterdam, Elsevier.

Fransson, P., Kruger, G., Merboldt, K. D., and Frahm, J. (1998). Physiologic aspects of event related paradigms in magnetic resonance functional neuroimaging. NeuroReport, 9, 2001-5.

Fransson, P., Kruger, G., Merboldt, K. D., and Frahm, J. (1999). Temporal and spatial MRI responses to subsecond visual activation. Magnetic Resonance Imaging, 17, 1-7.

Friston, K. J., Jezzard, P., and Turner, R. (1994). The analysis of functional MRI time series. Human Brain Mapping, 1, 153-71.

Friston, K. J., Buechel, C., Fink, G. R., Morris, J., Rolls, E., and Dolan, R. J. (1997). Psychophysiological and modulatory interactions in neuroimaging. Neuroimage, 6, 218-29.

Friston, K. J., Fletcher, P., Josephs, O., Holmes, A., Rugg, M. D., and Turner, R. (1998a). Event-related fMRI: Characterizing differential responses. NeuroImage, 7, 30-40.

Friston, K. J., Josephs, O., Rees, G., and Turner, R. (1998b). Nonlinear event-related responses in fMRI. Magnetic Resonance in Medicine, 39, 41-52.

Henson, R. N., Rugg, M. D., Shallice, T., Josephs, O., and Dolan, R. J. (1999a). Recollection and familiarity in recognition memory: An event-related functional magnetic resonance imaging study. Journal of Neuroscience, 15, 3962-72.

Henson, R. N., Shallice, T., and Dolan, R. J. (1999b). Right prefrontal cortex and episodic memory retrieval: A functional MRI test of the monitoring hypothesis. Brain, 122, 1367-81.

Jezzard, P., and Song, A. W. (1996). Technical foundations and pitfalls of clinical fMRI. NeuroImage, 4, S63-75.

Josephs, O., Turner, R. and Friston, K. (1997). Event-related fMRI. Human Brain Mapping, 5, 243-8.

Kim, S. G., Richter, W., and Ugurbil, K. (1997). Limitations of temporal resolution in functional MRI. Magnetic Resonance in Medicine, 37, 631-6.

Konishi, S., Yoneyama, R., Itagaki, H., Uchida, I., Nakajima, K., Kato, H., Okajima, K., Koizumi, H., and Miyashita, Y. (1996). Transient brain activity used in magnetic resonance imaging to detect functional areas. NeuroReport, 8, 19-23.

Kwong, K. K., Belliveau, J. W., Chesler, D. A., Goldberg, I. E., Weisskoff, R. M., Poncelet, B. P., Kennedy, D. N., Hoppel, B. E., Cohen, M. S., and Turner, R. (1992). Dynamic magnetic

resonance imaging of human brain activity during primary sensory stimulation. Proceeding of the National Academy of Science, USA, 89, 5675-9.

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. Magnetic Resonance in Medicine, 33, 745-54.

Logan, W. J. (1999). Functional magnetic resonance imaging in children. Seminars in Pediatric Neurology, 6, 78-86.

McCarthy, G., Luby, M., Gore, J., and Goldman-Rakic, P. (1997). Infrequent events transiently activate human prefrontal and parietal cortex as measured by functional MRI. Journal of Neurophysiology, 77, 1630-4.

Menon, R. S., Gati, J. S., Goodyear, B. G., Luknowsky, D. C., and Thomas, C. G. (1998). Spatial and temporal resolution of functional magnetic resonance imaging. Biochemistry and Cell Biology, 76, 560-71.

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

Ogawa, S., Lee, T., Nayak, A., and Glynn, P. (1990). Oxygenation-sensitive contrast in magnetic resonance image of rodent brain at high magnetic fields. Magnetic Resonance in Medicine, 14, 68-78.

Ogawa, S., Tank, D. W., Menon, R., Ellerman, J. M., Kim, S. G., Merkle, H., and Ugurbil, K. (1992). Intrinsic signal changes accompanying sensory stimulation: Functional brain mapping with magnetic resonance imaging. Proceeding of the National Academy of Science, USA, 89, 5951-5.

Ojemann, J. G., Buckner, R. L., Akbudak, E., Snyder, A. Z., Ollinger, J. M., McKinstry, R. C., Rosen, B. R., Petersen, S. E., Raichle, M. E., and Conturo, T. E. (1998). Functional MRI studies of word-stem completion: Reliability across laboratories and comparison to blood flow imaging with PET. Human Brain Mapping, 6, 203-15.

Posse, S., Schor, S., Gembris, D., Muller, E., Peyerl., M., Krocker., R., Grosse-Ruken, M. L., Elgghwaghi, B. and Taylor, J. G. (1998). Real-time fMRI on a clinical whole body scanner: Single-event detection of sensorimotor stimulation. NeuroImage, 7, S567.

Price, C. J., Veltman, D. J., Ashburner, J., Josephs, O., and Friston, K. J. (1999). The critical relationship between the timing of stimulus presentation and data acquisition in blocked designs with fMRI. NeuroImage, 10, 36-44.

Raichle, M. E. (in press). A brief history of functional brain imaging. In Brain Mapping: The Applications. (eds. A. Toga and J. Mazziotta), Academic Press, San Diego.

Richter, W., Georgopoulos, A. P., Ugurbil, K., and Kim, S. G. (1997). Detection of brain activity during mental rotation in a single trial by fMRI. NeuroImage, 5, S49.

Robson, M. D., Dorosz, J. L., and Gore, J. C. (1998). Measurements of the temporal fMRI response of the human auditory cortex to trains of tones. NeuroImage, 7, 185-98.

Robinson, D. L., and Rugg, M. D. (1988). Latencies of visually responsive neurons in various regions of the rhesus monkey brain and their relation to human visual responses. Biological Psychology, 26, 111-6.

Rosen, B. R., Buckner, R. L., and Dale, A. M. (1998). Event related fMRI: Past, present, and future. Proceeding of the National Academy of Science, USA, 95, 773-80.

Rugg, M. D., Fletcher, P. C., Frith, C. D., Frackowiak, R. S., and Dolan, R. J. (1996). Differential activation of the prefrontal cortex in successful and unsuccessful memory retrieval. Brain, 119, 2073-83.

Savoy, R. L., Bandettini, P. A., O'Craven, K. M., Kwong, K. K., Davis, T. L., Baker, J. R., Weiskoff, R. M., and Rosen, B. R. (1995). Pushing the temporal resolution of fMRI: Studies of very brief visual stimuli, onset variability and asynchrony, and stimulus-correlated changes in noise. Proceedings of the Society of Magnetic Resonance Third Scientific Meeting and Exhibition, 2, 450.

Schacter, D. L., Buckner, R. L., Koutstaal, W., Dale, A. M., and Rosen, B. R. (1997). Late onset of anterior prefrontal activity during true and false recognition: An event-related fMRI study. NeuroImage 6, 259-69.

Sereno, M. I., Dale, A. M., Reppas, J. B., Kwong, K. K., Belliveau, J. W., Brady, T. J., Rosen, B. R., and Tootell, R. B. (1995). Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. Science, 268, 889-93.

Shulman, G. L., Fiez, J. A., Corbetta, M., Buckner, R. L., Miezin, F. M., Raichle, M. E., and Petersen, S. E. (1997). Common blood flow changes across visual tasks: II. Decreases in cerebral cortex. Journal of Cognitive Neuroscience, 9, 648-63.

Thomas, K. M., King, S. W., Franzen, P. L., Welsh, T. F., Berkowitz, A. L., Noll, D. C., Birmaher, V., and Casey, B. J. (1999). A developmental functional MRI study of spatial working memory. NeuroImage, 10, 327-38.

Vasquez, A. L., and Noll, D. C. (1998). Nonlinear aspects of the BOLD response in functional MRI. NeuroImage, 7, 108-18.

Voyvodic, J. T. (1999). Real-time fMRI paradigm control, physiology, and behavior combined with near real-time statistical analysis. NeuroImage, 10, 91-106.

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, 1188-91.

Worsley, K. J., and Friston, K. J. (1995). Analysis of fMRI time-series revisited – again. NeuroImage, 2, 173-81.

Zarahn, E., Aguirre, G., and D'Esposito, M. (1997). A trial-based experimental design for fMRI. NeuroImage, 6, 122-38.