

# **An introduction to functional MRI**

**Bianca de Haan & Chris Rorden**

As taken from the PhD thesis: de Haan, B. (2006). Pre-requisites of perceptual awareness: Insights from studying deficits of attention. PhD, School of Psychology, University of Nottingham, UK.

## **MEASURING BRAIN ACTIVATION.**

Functional magnetic resonance imaging (fMRI) has rapidly become a popular tool for measuring brain function. FMRI reveals which parts of the brain are activated by behavioural tasks with a spatial resolution of 2-5 millimetres, which is superior to many of the other techniques in cognitive neuroscience. This means that places of activity in the brain that are as close as 2-5 millimetres apart in the brain can theoretically still be distinguished from each other. The temporal resolution (the minimal distance in time between two data points that can still be distinguished), however, is relatively poor (5-8 seconds) (Horwitz, Friston, & Taylor, 2000; Menon, 2001).

As with any tool, the use of fMRI can be made more effective by understanding how it works and how (not) to use it. This chapter aims to provide that information. The first section explains the principles underlying MRI and how MR images are obtained. Second, attention will be given to a specific protocol of MRI used to indirectly infer the functional activity of the brain, referred to as functional Magnetic Resonance Imaging (fMRI). Specifically, the second section of this chapter describes how neural activation leads to an fMRI signal. In the third section of this chapter, attention will be given to the strengths and weaknesses of the different kinds of experimental designs that can be used in fMRI experiments. Finally, the fourth section of this chapter will be devoted to the different steps involved when analysing fMRI datasets.

## 1.1 THE BASICS OF MRI.

Magnetic Resonance Imaging, or MRI, typically measures the response of hydrogen molecules to a perturbation while in a magnetic field. Explaining the physics of MRI in full detail would be beyond the scope of this chapter. However, some basic knowledge of the principles is required to fully understand the subsequent sections. Providing this basic knowledge is the aim of this first section.

There are four basic steps involved in obtaining an MR image of the brain. The first step is to place the brain in a magnetic field. The second step is the application of a brief radiofrequency (RF) pulse. This RF-pulse perturbs the hydrogen nuclei in the brain. The third step is measuring the radio signal emitted from the hydrogen nuclei. The fourth step is transforming this signal to obtain a 3-dimensional MRI image (Buxton, 2002; Hashemi, Bradley, & Lisanti, 2004; Horowitz, 1995; Jezzard & Clare, 2001). Each of these steps will be considered in turn, concluding with a discussion of the pulse cycles that are commonly used.

### 1.1.1 THE BRAIN IN A MAGNETIC FIELD.

The first step to measuring an MRI signal is to place the brain in a magnetic field. Of the many types of nuclei in the brain, it is the hydrogen nucleus that is most commonly used in MRI (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). This is because hydrogen nuclei are abundant in the human brain and give a relatively strong MRI signal. However, in theory the following explanation applies to all nuclei that are electrically charged and spin around their axis (for example, Carbon-13 nuclei could have been used in MRI as well, but the resulting brain images would have had a very low signal to noise relative to hydrogen nuclei).

Hydrogen nuclei are positively charged particles that spin around their axis. When an electrically charged particle moves, it produces a magnetic field. This magnetic field can be

represented as a vector (a mathematical entity with both an amplitude and a direction). Generally, a vector is mathematically depicted as an arrow where the length of the arrow represents the amplitude of the vector and the direction in which the arrow is pointing reflects the direction of the vector. Therefore, each hydrogen nucleus in the brain can be seen as a vector with the vector representing the strength and direction of the magnetic field of the hydrogen nucleus produced by its spinning around its axis. This vector is also known as the Magnetic Dipole Moment (MDM) (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001).

Before the brain is placed in a magnetic field the MDM's of the hydrogen nuclei in the brain all point in random directions, they are not aligned. When the brain is placed in a magnetic field, two things happen simultaneously (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). Firstly, the MDM's of a portion of the hydrogen nuclei align themselves in the direction of the main magnetic field. How many of the MDM's align themselves in the direction of the main magnetic field depends on the strength of the magnetic field (and theoretically on the temperature of the brain, but in the present discussion that is not relevant). The stronger the magnetic field, the higher the percentage of the MDM's that align themselves to the magnetic field (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). For example, 4.5 MDM's per million align themselves in the direction of the main magnetic field in a 1.5 Tesla magnetic field (Morris, 1986). Secondly, when the brain is placed in a magnetic field the MDM's of the hydrogen nuclei start to precess (see Figure 2.1a). The frequency of this precession depends first of all on the type of nucleus. This means that the MDM of a hydrogen nucleus will have a different frequency of precession from, for instance, the MDM of a sodium nucleus in a magnetic field of equal strength. Second of all, the frequency of precession depends on the strength of the magnetic field. The frequency of precession is directly proportional to the strength of the magnetic field, so the

stronger the magnetic field the higher the frequency of precession. For example, in a magnetic field of 1.5 Tesla the frequency of precession for the MDM's of hydrogen nuclei will be 64 MHz (64,000,000 revolutions per second) and in a magnetic field of 3 Tesla the frequency of precession will be 128 MHz (Horowitz, 1995; Jezzard et al., 2001).




Figure 1.1a: Precession of the MDM; the 'tail' is stationary while the 'top' moves around in a circular motion. This motion is similar to a spinning dreidel.

### 1.1.2 APPLICATION OF THE RADIOFREQUENCY PULSE.

The second step when measuring the MRI signal is the brief application of the radiofrequency pulse. The radiofrequency (RF) pulse is typically an electromagnetic wave resulting from the brief application of an alternating current perpendicular to the direction of the main magnetic field, otherwise known as a  $90^\circ$  RF-pulse (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). The ultimate goal of this  $90^\circ$  RF-pulse is to 'tip' the MDM's of the hydrogen nuclei. Conventionally, the direction along the main magnetic field is referred to as the z-axis. The  $90^\circ$  RF-pulse then basically 'tips' the MDM's in the x-y plane (see Figure 2.1b). This will only work if the frequency of the RF-pulse equals the frequency of the precession of the MDM's. Because the MDM's of the hydrogen nuclei have their own specific frequency of precession in a given magnetic field (e.g. as mentioned earlier, 64 MHz in a 1.5 Tesla magnetic field), it is possible to selectively 'tip' the MDM's of the hydrogen nuclei (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001).




Figure 1.1b: 'Tipping' of an MDM into the x-y plane during application of the RF-pulse.

After the MDM's of the hydrogen nuclei are 'tipped', the  $90^\circ$  RF-pulse is terminated and over time the MDM's return to their original orientation. This returning to the original orientation is known as relaxation (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001).

### 1.1.3 RELAXATION.

After the  $90^\circ$  RF-pulse is terminated, the MDM's of the hydrogen nuclei return from their 'tipped' state to their original lower energy state of being aligned in the direction of the magnetic field (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). Basically, the RF-pulse poured energy into the system and this energy is released when the MDM's return to their original state. This release of energy is known as relaxation and this is the radiofrequency signal that is measured during MRI (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001).

The MDM of a hydrogen nucleus can be broken down into two components. One component of the MDM is the amplitude in the z-axis. The other component of the MDM is the amplitude in the x-y plane (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). Before application of the RF-pulse the amplitude in the z-axis is maximal while

the amplitude in the x-y plane is zero. Just after application of the RF-pulse the amplitude in the z-axis is zero while the amplitude in the x-y plane is maximal. During relaxation the amplitude in the z-axis will slowly increase while the amplitude in the x-y plane slowly decreases. Therefore, the relaxation of the MDM's of the hydrogen nuclei has two components; firstly, a re-growth along the z-axis and secondly, a decay in the x-y plane. The re-growth along the z-axis of the MDM's is referred to as T1 relaxation. The decay in the x-y plane of the MDM's is referred to as T2 relaxation (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001).

#### 1.1.4 WHEN IT ALL COMES TOGETHER.

The application of the 90° RF-pulse and the measuring of the energy released during relaxation is repeated many times in a typical MRI experiment.

Different tissues in the brain have different T1 and T2 relaxation rates (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). This means that at each moment after termination of the RF-pulse, the amplitude of the MDM's of the hydrogen nuclei in the z-axis and the amplitude of the MDM's in the x-y plane will be different for different tissues. If now the MRI signal is measured at a point after termination of the RF-pulse where either the relative difference between the amplitudes of the MDM's of the hydrogen nuclei of different tissues in the z-axis is maximized or the relative difference between the amplitudes of the MDM's of different tissues in the x-y plane is maximized a maximum contrast between different tissues will be obtained (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001).

When the MRI signal is measured at a point when the relative difference between the amplitudes of MDM's of different tissues in the z-axis is maximized, the signal is known as a T1 weighted signal. Alternatively, when the MRI signal is measured at a point when the

relative difference between the amplitudes of the MDM's of different tissues in the x-y plane is maximized, the signal is known as a T2 weighted signal (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). By changing certain scanner parameters either a T1 weighted signal or a T2 weighted signal can be acquired (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). If the time from RF-pulse to measurement of the signal (TE) is kept short, while simultaneously the time between two successive RF-pulses (TR) is also kept short, the difference in T1 for the different tissues is maximized and the acquired scan is called a T1 weighted scan. Typically, T1 weighted scans are acquired with a TE of approximately 20ms and a TR of approximately 500ms. T1 weighted scans are also known as anatomical scans, because they particularly show good contrast between grey and white matter. Alternatively, if the TE is long while the TR at the same time is also long, the difference in T2 for the different tissues is maximized and the acquired scan is called a T2 weighted scan. T2 weighted scans are typically acquired with a TE of approximately 80ms and a TR of approximately 2000ms. T2 weighted scans are sometimes referred to as pathological scans, because lesions appear very bright (Horowitz, 1995; Jezzard et al., 2001).

#### 1.1.5 THE IMAGE IN SPACE.

Now that it has been explained how the MRI signals are obtained, it is time to discuss how the MRI signals are transformed into a 3D image of the brain. Remember from the previous paragraphs that the MRI signal is obtained by placing the brain in a magnetic field, applying a RF-pulse and measuring the relaxation of the MDM's. However, at this stage the MRI signal still contains information about the entire region activated by the RF-pulse. There is no way of telling what the strength of the MRI signal for each unique part of the brain is. Therefore, some additional steps need to be performed to determine the individual

contribution of each part of the brain to the MRI signal. This paragraph will describe how this spatial location information is obtained.

To determine the unique contribution of each part of the brain to the MRI signal, three magnetic gradients are used (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). A magnetic gradient is a change in the strength of the magnetic field over a certain spatial distance. In other words, as the spatial distance between two points changes, the difference in the strength of the magnetic field between these two points changes as well. In MRI, linear magnetic gradients are used. This means that there is a linear relationship between the change in spatial distance and the change in the strength of the magnetic field. The three magnetic gradients that are used are known as the slice select gradient, the readout or frequency-encoding gradient and the phase-encoding gradient. These three magnetic gradients are oriented orthogonal to each other along the z-axis, the x-axis and the y-axis. The z-axis runs from head to foot, the x-axis runs from ear to ear and the y-axis runs from nose to back of the head (see Figure 2.1c). Each magnetic gradient can be oriented along any of these axes, but to keep things simple it will be assumed that the slice select gradient is oriented along the z-axis, the frequency-encoding gradient is oriented along the x-axis and the phase-encoding gradient is oriented along the y-axis (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). First the slice select gradient will be considered before moving on to the frequency-encoding gradient and phase-encoding gradient.




Figure 1.1c: A schematic diagram of a subject with the 3 coordinate axes.

While the brain is in a magnetic field, the slice select gradient is briefly switched on at the same moment in time as the RF-pulse is applied (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). The result of the slice select gradient is that the magnetic field is no longer uniform, but varies linearly as we move along the z-axis. The example will be used here that the magnetic field is weaker towards the feet and stronger towards the top of the head. Remember from paragraph 2.1.1 that the frequency of precession of the MDM's is directly proportional to the strength of the magnetic field. The result of this gradient is therefore that the frequency of precession of the MDM's will be different for different positions along the z-axis. The MDM's closer to the feet will have a lower frequency of precession than MDM's towards the top of the head. Remember also from paragraph 2.1.2 that the frequency of the RF-pulse needs to match the frequency of precession of the MDM's in order to tip them into the x-y plane. Taken together, this means that by selecting an RF-pulse with a specific range of frequencies, it is possible to selectively tip the MDM's at a specific spatial range along the z-axis. In other words, it is possible to selectively only measure the relaxation of MDM's in a specific two-dimensional slice of the brain. By increasing the frequency of the RF-pulse while keeping the range of frequencies equal, each slice of the brain can be selected successively until the entire brain is measured. By changing

either the steepness of the gradient or the range of frequencies in the RF-pulse, the thickness of the slices can be set (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001).

After application of the slice select gradient, the MRI signal still contains information about the entire slice. To determine the contribution of each individual point in a slice (commonly referred to as a pixel or voxel), the frequency-encoding gradient and phase-encoding gradients are used. The application of the frequency-encoding and phase-encoding gradients results in each voxel being represented by a unique combination of phase and frequency. Before explaining this however, it is necessary to first discuss the basics of Fourier analysis (Fourier, 1822).

Every mathematical function that varies with time can be decomposed into a unique combination of sine and cosine waves (Buxton, 2002; Fourier, 1822; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). Each sine and cosine wave has a frequency and an amplitude. After the mathematical function has been decomposed into a combination of sine and cosine waves, it is possible to describe the amplitudes of these waveforms as a function of their frequency. For instance, imagine a mathematical function that can be decomposed into a sine wave with an amplitude of 10 and a frequency of 10 Hz and a sine wave with an amplitude of 50 and a frequency of 20 Hz. It is now possible to represent these amplitudes as a function of frequency with a peak of 10 at 10 Hz and another peak of 50 at 20 Hz (see Figure 2.1d).




Figure 1.1d: Waveforms as a function of time can also be represented as a function of frequency. On the left two sine waves with frequencies of 10 (red) and 20 (blue) Hz and amplitudes of 10 and 50 respectively are shown. On the right these same sine waves are shown as a function of frequency. The right figure is known as a Fourier spectrum.

Transforming a mathematical function that varies with time into an amplitude function that varies with frequency is known as Fourier analysis. A different way to say the same thing would be that a Fourier transform has been applied to the mathematical function and the result can be represented as a Fourier spectrum. The neat thing about Fourier analysis is that if a Fourier transform is applied to the Fourier transform of a mathematical function, the original mathematical function is again obtained (Buxton, 2002; Fourier, 1822; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001).

Now that the basics of Fourier analysis have been explained, the workings of the frequency-encoding and phase-encoding gradients can be discussed. As mentioned, to determine the contribution of each individual voxel in a slice to the MRI signal of that slice, the frequency-encoding gradient and phase-encoding gradients are used. The frequency-encoding gradient is applied along the x-axis during the measurement of the MRI signal. This is why this gradient is also referred to as the readout gradient (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). Imagine a slice with 3 columns in the x direction and 3 rows in the y direction (i.e. a slice with 9 voxels). Before the frequency-encoding gradient is applied, the MDM's in all columns will precess at the same frequency. In a Fourier spectrum the MRI signal from the entire slice would appear as one peak. The result of the frequency-encoding gradient is a different magnetic field strength for the different columns in the slice. The result is that the MDM's in the different columns start precessing at different frequencies. The Fourier spectrum of the MRI signal during application of the frequency-encoding gradient will therefore now show three peaks, one for each frequency of precession. If the strength of the frequency-encoding gradient increases as we move in the positive

direction on the x-axis, this means that the leftmost column will have the lowest magnetic field strength and the rightmost column will have the highest magnetic field strength. In this case, in the Fourier spectrum the lowest frequency peak reflects the MRI signal from the leftmost column and the highest frequency peak reflects the MRI signal from the rightmost column. Therefore, the Fourier spectrum reflects both location and strength of the MRI signal for each column in the slice. To summarise, the frequency-encoding gradient is applied during measurement of the signal and the result is that each column is now represented by a unique frequency (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001).

After application of the frequency-encoding gradient, however, the MRI signal for each column still reflects the contribution of the 3 rows. To achieve the goal of obtaining a unique MRI signal for each individual voxel the phase-encoding gradient is used (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). To understand the workings of the phase-encoding gradient, the concept of phase must first be explained.

Remember that the MDM's of hydrogen nuclei in a magnetic field of a certain field strength all precess at the same frequency. However, even though the MDM's all precess at the same frequency, they are not necessarily precessing in the same phase (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). To clarify this, an analogy using clocks can be helpful. Imagine a number of clocks that are synchronized. At any given moment, the times that these clocks indicate will be the same. In this case, the hands of the clocks rotate both at the same frequency (an hour lasts equally long for all clocks) and in the same phase. Now imagine a number of clocks that are not synchronized. Even though the hands of the different clocks still rotate at the same frequency (an hour still lasts equally long for all clocks), at a given moment the times these clocks indicate are not the same i.e. the clocks are not in phase. When applied to the precession of the MDM's in a magnetic field, this means that even though the MDM's precess at the same frequency, they can either be all

at the same position in their cycle or they can all be at different position in their cycle at a given point in time. In the first case they are precessing in phase, in the latter case they are precessing out of phase (see Figure 2.1e) (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001).




Figure 1.1e: The top row displays one moment in the cycle of precession of 3 MDM's that are in phase. The bottom row displays one moment in the cycle of precession of 3 MDM's that are not in phase.

Now that the concept of phase has been explained, the workings of the phase-encoding gradient can be examined. The phase-encoding gradient is briefly applied along the y-axis at a moment in time between the RF-pulse and measurement of the MRI signal. At the instant when the RF-pulse is applied, the MDM's are tipped in the x-y plane and forced to precess in phase (as the brief RF-pulse tips all the MDM's simultaneously). When the phase-encoding gradient is subsequently applied along the y-axis the MDM's in the different rows will experience different magnetic field strengths. Therefore, the MDM's in the different rows will start precessing at different frequencies. When the phase-encoding gradient is switched off, the MDM's in the different rows will again all be precessing at the same frequency, but they will no longer be in phase. Each row will have a different phase (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001).

After application of the frequency-encoding and phase-encoding gradients, the MRI signal for each voxel will be represented by a unique combination of frequency and phase (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). In principle therefore, one frequency-encoding and phase-encoding step is enough to obtain unique information for each pixel in the slice of the brain. This would, however, lead to an image of the brain with a very poor signal to noise ratio. Generally, as many phase-encoding (and therefore frequency-encoding) steps as there are rows in the slice are performed.

#### 1.1.6 T2 AND THE SPIN-ECHO PULSE CYCLE.

The application of the RF-pulse, usage of the magnetic gradients and subsequent measurement of the unique contribution of each pixel in the brain to the MRI signal is known as a pulse sequence (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). In the previous paragraphs, the most basic pulse cycle has been described, namely a pulse cycle consisting of one  $90^\circ$  RF-pulse followed by one measurement of the MRI signal. A single run of this basic pulse cycle would result in the measurement of the MRI signal after one phase-encoding and frequency-encoding step for one slice of the brain. Multiple phase-encoding and frequency-encoding steps are needed to measure the MRI signal in a single slice and the brain contains multiple slices. This means that using this basic pulse cycle, many repetitions of this pulse cycle would be needed to measure the MRI signal in the entire brain with a sufficient signal to noise ratio. This would make the duration of an MRI experiment way too long (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). This section discusses some more efficient pulse cycles. To understand these pulse cycles, however, a diversion must first be made here concerning the difference between  $T2^*$  decay and  $T2$  decay. The reason for this diversion is that these more efficient pulse cycles are strongly related to this difference between  $T2^*$  decay and  $T2$  decay.

As mentioned, before the application of the RF-pulse, the MDM's are not precessing in the same phase. It is the application of the RF-pulse that forces all the MDM's to precess in the same phase. The result is that the MDM signals are additive and this gives a strong signal in the x-y plane. After termination of the 90° RF-pulse, however, the precession of the MDM's will dephase again. The MDM signals are now no longer additive, but cancel each other out and the signal decays (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). There are two reasons why the MDM's dephase after termination of the 90° RF-pulse and the signal decays. The first one is magnetic field inhomogeneity (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). A process known as shimming is usually applied at the start of the MRI experiment to make the magnetic field as homogeneous as possible. However, despite shimming, the strength of the magnetic field is not entirely uniform. Since the frequency of precession of the MDM's depends on the strength of the magnetic field, it follows that MDM's in different positions in the magnetic field will precess at a different frequency and therefore these precessions will dephase. The second reason why the MDM's dephase is because of spin-spin interaction (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). Different hydrogen nuclei are surrounded by different other nuclei. These other nuclei affect the frequency of precession of the MDM's of the hydrogen nuclei. The frequency of precession of each MDM will be differently affected by the surrounding nuclei. This, again, results in different frequencies of precession for different MDM's and hence dephasing occurs (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001).

To summarize, initially the precession of the MDM's is dephased. When the 90° RF-pulse is applied, the MDM's are forced to precess in phase, resulting in a signal in the x-y plane. After termination of the 90° RF-pulse, the precession of the MDM's will dephase again

due to the inhomogeneities in the magnetic field and spin-spin interactions and the signal in the x-y plane decays.

In the previous paragraphs of this section, it was implied that the decay of the MDM's in the x-y plane after termination of the  $90^\circ$  RF-pulse equals the T2 relaxation signal. This is, however, a simplification. True T2 decay is actually a lot slower than the decay of the MDM's in the x-y plane after termination of the  $90^\circ$  RF-pulse. The decay of the MDM's in the x-y plane is more accurately described as T2\* decay (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). So how can the T2 signal be obtained?

As explained, the reason why the MDM's decay in the x-y plane (the T2\* signal) is essentially due to dephasing (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). The trick is that the dephasing due to the inhomogeneity of the magnetic field is correctable and by correcting for this source of dephasing the true T2 signal is obtained (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). Directly after application of the  $90^\circ$  RF-pulse the precession of the MDM's is in phase. This means that at any given time, all the MDM's will be at the same point in their cycle of precession. After termination of the  $90^\circ$  RF-pulse the precession of the MDM's will slowly dephase. Basically, at a given time, different MDM's will no longer be at the same point in their cycle of precession. If an RF-pulse is now applied from the opposite direction ( $180^\circ$ ) to that from which the original RF-pulse was applied ( $90^\circ$ ), the direction of rotation of the precession of the MDM's is reversed. After the same amount of time has elapsed following the  $180^\circ$  RF-pulse as the amount of time between the  $90^\circ$  RF-pulse and the  $180^\circ$  RF-pulse, the MDM's will be in phase again (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001).

An analogy might be helpful. Imagine two cars leaving from the same point in the same direction. One is moving at a speed of 100 km/h and the other one is moving at a speed of 80 km/h. After a while the faster car will be further away from the point of origin than the

slower car. After, for example, half an hour both cars turn around and head towards the point of origin again. If both cars keep moving at their same respective speed, they will both arrive at the point of origin again at the same time. To relate this analogy back, directly after application of the  $90^\circ$  RF-pulse all the MDM's are at the same point in their cycle of precession. After termination of the RF-pulse some MDM's will rotate faster (have higher frequencies of precession) than others, the MDM's dephase. After a while the  $180^\circ$  RF-pulse is applied and this makes the direction of precession of all the MDM's reverse. The MDM's however, all keep their own frequency of precession and will arrive at their starting point (the point where they were directly after application of the  $90^\circ$  RF-pulse) in the cycle of precession at the same time; they will be in phase again.

The important thing is that the time between the  $180^\circ$  refocusing RF-pulse and measurement of the signal must be the same as the time between the  $90^\circ$  RF-pulse and the  $180^\circ$  RF-pulse. A measurement cycle where one  $90^\circ$  RF-pulse is followed by one or more  $180^\circ$  RF-pulses with a measurement after each  $180^\circ$  RF-pulse is known as a spin-echo pulse cycle. Even though the spin-echo pulse cycle corrects for the decay in the signal caused by the inhomogeneities in the magnetic field, the signal still does eventually decay because of the dephasing due to spin-spin interactions. The decay is now, however, a lot slower and this is the true T2 decay (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001).

At the beginning of this paragraph, the basic pulse cycle was mentioned, where after one  $90^\circ$  RF-pulse the MRI signal is measured once for one phase-encoding step in one slice. In a spin-echo pulse cycle, following one  $90^\circ$  RF-pulse the MRI signal is measured after each  $180^\circ$  refocusing pulse (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001). However, the phase-encoding gradient is only switched on once. This means that the same phase-encoding step for the same slice is measured after each  $180^\circ$  refocusing pulse.

Basically, the same information is collected multiple times, which increases the signal to noise ratio (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jezzard et al., 2001).

At the moment, the fastest and most commonly used scanning technique to collect functional MRI (see section 2.2) is echo planar imaging (EPI) (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jones, Brookes, & Moonen, 2001; Mansfield, 1977). The EPI technique is based on the spin-echo pulse cycle. However, instead of only switching on the phase-encoding gradient once in the pulse cycle, the phase-encoding gradient is switched on after each  $180^\circ$  refocusing pulse while the steepness of this gradient increases with each  $180^\circ$  refocusing pulse. This means that instead of only collecting one phase-encoding step for one slice in a single pulse cycle, as many phase-encoding steps as there are  $180^\circ$  refocusing pulses are collected for that slice. Imagine for example a slice where 64 phase-encoding steps are needed to measure the MRI signal. In a conventional spin-echo pulse cycle with four  $180^\circ$  refocusing pulses, 64 pulse cycles (so 64  $90^\circ$  RF-pulses) would be needed and the MRI signal of the entire slice would be measured four times. Using EPI, however, only 16 pulse cycles would be needed, but the MRI signal of the entire slice would only be measured once. This means that EPI is very fast, but the trade-off is a lower signal to noise ratio (Buxton, 2002; Hashemi et al., 2004; Horowitz, 1995; Jones et al., 2001; Mansfield, 1977).

## 1.2 FROM NEURAL ACTIVATION TO THE FMRI SIGNAL.

Besides looking at structural scans of the brain, MRI can also be used to look at functional activity of the brain, in which case it is referred to as functional magnetic resonance imaging (fMRI). The type of scanning technique most commonly used to obtain fMRI images is echo planar imaging (EPI). The contrast technique most commonly used in fMRI is the so-called BOLD (blood oxygenation level-dependent contrast) technique. The BOLD technique is based on the fact that, under normal circumstances, neuronal activity and haemodynamics (regulation of blood flow and oxygenation) are linked in the brain (Heeger & Ress, 2002; Ogawa et al., 1992). This section starts by explaining the mechanisms underlying the BOLD fMRI technique. The second part of this section considers the neural signal that BOLD fMRI is thought to reflect. Finally, some issues concerning the spatial and temporal resolution of BOLD fMRI will be discussed.

### 1.2.1 THE BOLD FMRI RESPONSE.

As explained in the previous section, the amount of energy released by the hydrogen molecules after the termination of the RF-pulse gradually decays over time. The rate of this decay differs for different tissues and this makes the distinction between different types of tissue possible. One reason for the decay of the fMRI signal is the dephasing of the precession of the MDM's of the hydrogen nuclei due to inhomogeneities in the magnetic field. The larger the inhomogeneity of the magnetic field, the faster the precessions of the MDM's will dephase and the faster the fMRI signal decays. Since the fMRI signal is measured at a predetermined point in time after termination of the RF-pulse, the magnitude of the fMRI signal will be smaller at that time of measurement when the signal decays faster as compared to when the signal decays slower (see Figure 2.2a). Therefore, the larger the inhomogeneity of

the magnetic field, the smaller the fMRI signal at the time of measurement (Heeger et al., 2002; Horowitz, 1995).




Figure 1.2a: Decay of the fMRI signal over time after termination of the RF-pulse. The blue curve represents the decay in fMRI signal in a magnetic field with a larger inhomogeneity. The red curve represents the decay in fMRI signal in a magnetic field with a smaller inhomogeneity. Therefore, the blue curve decays faster than the red curve. The green arrow indicates the point in time at which the signal intensity is measured. It can be seen that when the fMRI signal decays faster (the blue curve), the signal intensity at the moment of measurement is lower than when the fMRI signal decays slower (the red curve).

Usually, this inhomogeneity of the magnetic field is considered an artefact and every attempt is made to make the overall magnetic field as uniform as possible. However, even though great care is taken to make the overall magnetic field as uniform as possible, small local inhomogeneities in the magnetic field still occur. In the BOLD fMRI technique these little inhomogeneities are used to indirectly measure neuronal activity (Detre & Wang, 2002; Heeger et al., 2002; Ogawa et al., 1992; Ogawa, Lee, Kay, & Tank, 1990a; Ogawa, Lee, Nayak & Glynn, 1990b). Since the T2 signal is relatively insensitive to inhomogeneities in the magnetic field (remember from the previous section that the T2 signal is obtained by correcting for the effects of the inhomogeneity of the magnetic field) it follows that the signal most commonly measured in BOLD fMRI is the T2\* signal.

The BOLD fMRI technique basically measures changes in the inhomogeneity of the magnetic field, which are a result of changes in the level of oxygen present in the blood (blood oxygenation) (Detre et al., 2002; Heeger et al., 2002; Ogawa et al., 1992; Ogawa et al.,

1990a; Ogawa et al., 1990b). Deoxyhaemoglobin (a haemoglobin protein contained within the red blood cells without an oxygen molecule attached to it) has magnetic properties and will cause an inhomogeneity in the magnetic field surrounding it. Oxyhaemoglobin (a haemoglobin protein contained within the red blood cells with an oxygen molecule attached to it) has hardly any magnetic properties and therefore has very little effect on the magnetic field surrounding it. Therefore, a high level of deoxyhaemoglobin in the blood will result in a greater field inhomogeneity and therefore in a decrease of the fMRI signal (Detre et al., 2002; Heeger et al., 2002; Ogawa et al., 1992; Ogawa et al., 1990a; Ogawa et al., 1990b).

The function of the BOLD fMRI signal against time in response to a temporary increase in neuronal activity is known as the haemodynamic response function (HRF) (Heeger et al., 2002) (see Figure 2.2b). After a transient increase in neuronal activity the BOLD fMRI signal initially decreases because the active neurons use oxygen thereby increasing the relative level of deoxyhaemoglobin in the blood (Heeger et al., 2002; Vanzetta & Grinvald, 1999). This decrease, however, is tiny and is not always found (Detre et al., 2002; Ugurbil, Toth, & Kim, 2003). Following this initial decrease, there is a large increase in the BOLD fMRI signal which reaches its maximum after approximately 6 seconds (Fox, Raichle, Mintun & Dence, 1988; Heeger et al., 2002). This increase is due to a massive oversupply of oxygen-rich blood. The result of this oversupply of oxygen is a large decrease in the relative level of deoxyhaemoglobin, which in turn causes the increase in the BOLD fMRI signal. Finally, the level of deoxyhaemoglobin slowly returns to normal and the BOLD fMRI signal decays until it has reached its original baseline level after an initial undershoot after approximately 24 seconds (Heeger et al., 2002). This function of the BOLD fMRI signal against time is also known as the BOLD response.




Figure 1.2b: Time course of the HRF in response to a short-lasting increase in neuronal activity at time = 0. Note that the signal shows a small dip immediately after the transient increase in neuronal activity (initially, the level of deoxyhaemoglobin increases). Following this initial dip, however, the signal becomes much stronger due to the increase in oxygenation (and therefore the decrease in the level of deoxygenation).

The massive overcompensation for the amount of oxygen used by the neurons is the main reason there is a BOLD fMRI signal to measure in the first place. However, the reasons for the oversupply in oxygen-rich blood are not entirely clear. There are two main hypotheses regarding the reason for this disproportionate increase in oxygenated blood. The first theory states that the increase in blood flow does compensate for the oxygen being used by the active neurons (Buxton & Frank, 1997; Heeger et al., 2002; Logothetis & Wandell, 2004b). According to this theory, the increase in blood flow has to be this large since oxygen is less efficiently extracted from the blood when the blood flows fast. In this view, the amount of oxygen extracted does actually match the oxygen demands of the active neurons, even though the amount of oxygen in the blood seems to massively overcompensate for the oxygen used (Buxton et al., 1997; Heeger et al., 2002; Logothetis et al., 2004b). The second theory poses that the increase in blood flow serves to match the glucose need of the brain. This view seems to be supported by the observation that the increase in blood flow roughly corresponds to the glucose being used by the active neurons (Arthurs & Boniface, 2002; Attwell & Ladekola, 2002; Heeger et al., 2002).

## 1.2.2 WHAT NEURAL SIGNAL DOES BOLD FMRI REFLECT?

An area of hot debate concerns the type of neural signal that the BOLD response reflects. Does the BOLD response reflect the firing rates of neurons, or does the BOLD response instead follow from synaptic processing. This is an important issue for the interpretation of the BOLD fMRI signal. If the BOLD response reflects the firing rates of neurons, it should be interpreted as a measure of (relative) long-distance communication between different parts of the brain. If, on the other hand, the BOLD response is associated with synaptic processing, it should be interpreted as a measure of the input to and local processing in a particular area of the brain (Logothetis et al., 2004b).

One of the ways in which studies have tried to tackle this question, has been to look at the energy consumption of the brain. The idea is that the BOLD response is most likely to reflect the type of neural processing that demands most energy. An impressive amount of studies indicate that in the brain, most of the energy is needed for post-synaptic processes (Arthurs et al., 2002; Attwell et al., 2002; Attwell & Laughlin, 2001; Logothetis & Pfeuffer, 2004a; Logothetis et al., 2004b). This would suggest that the BOLD response is likely to be strongly related to synaptic processing. However, recent studies suggest that even though the BOLD response might be correlated with the neural processes that demand most energy, the BOLD response might not actually be driven by this energy use (Attwell et al., 2002).

Converging evidence for the association between synaptic processing and the BOLD response, however, comes from experiments combining intracortical extracellular recording techniques and BOLD fMRI measurements (Lauritzen & Gold, 2003; Logothetis, Pauls, Augath, Trinath, & Oeltermann, 2001; Logothetis et al., 2004a; Logothetis et al., 2004b). When an electrode is placed intracortically in the extracellular space, a so-called mean extracellular field potential (mEFP) can be measured (Logothetis et al., 2004a; Logothetis et al., 2004b). By using highpass and lowpass filters this mEFP can be separated into a low-

frequency component and a high-frequency component. The low-frequency component is known as the local field potential (LFP) and reflects mainly synaptic activity. The high-frequency component is known as the multiple-unit spiking activity (MUA) and reflects regional neuronal spiking (Logothetis et al., 2004a; Logothetis et al., 2004b). In a series of groundbreaking experiments Logothetis and colleagues found that BOLD responses could be more accurately predicted from LFP's than from MUA's. This suggests that the BOLD response reflects local processing and subthreshold activation rather than axonal firing. Furthermore, when LFP's and MUA's were dissociated only LFP's were able to predict the BOLD response (Logothetis et al., 2001).

Taken together, these studies seem to suggest that the BOLD fMRI response reflects mainly synaptic processing and therefore the input to and local processing in a given area of the brain. In this light it is also interesting to note that the density of the vasculature in any particular area of the brain is highly correlated with the number of synapses in that area of the brain and not necessarily with the number of neurons (Logothetis et al., 2004b).

### 1.2.3 SPATIAL AND TEMPORAL RESOLUTION OF BOLD FMRI.

BOLD fMRI reveals which parts of the brain are active in certain tasks with a spatial resolution of 2-5 millimetres, which is superior to many of the other techniques in cognitive neuroscience. The temporal resolution, however, is relatively poor (5-8 seconds) (Horwitz et al., 2000; Menon, 2001). This paragraph will discuss some issues surrounding the spatial and temporal resolution of BOLD fMRI.

In an ideal world, the measured BOLD signal would reflect only the changes in the capillaries close to the source of the neuronal activity. These local changes offer the most accurate information about the spatial location of the neural activity. However, the BOLD fMRI signal is sensitive to contaminations of large veins and arteries in the brain (Menon &

Goodyear, 2001). The large arteries supplying oxygen to the brain are fully oxygenated. This means that there is little room for a change in the oxygenation of the blood in these arteries. Since the BOLD fMRI signal measures the change in the relative level of oxygenation these arteries are therefore not likely to contribute hugely to a change in the BOLD fMRI signal (Menon & Kim, 1999). However, draining veins are only 70% oxygenated. The relative decrease in deoxyhaemoglobin is larger in these large veins than in the small capillaries close to the source of the neural activation. This means that the location of the maximum BOLD fMRI signal could be displaced towards the large draining veins that can be a few millimetres away from the site of neural activation. Furthermore, these large veins usually drain an area of the brain that is larger than the area of neural activation. This results in the area of activation appearing larger than it really is, a problem otherwise known as spatial blurring (Menon et al., 2001; Menon et al., 1999). Using a high magnetic field (4 Tesla or more) in combination with changes to the scanning sequence (by suppressing signals associated with high flow rates) and experimental design (by using short duration stimulus presentation) seems to go some way towards resolving this issue (Menon et al., 2001; Menon et al., 1999). Because of this susceptibility of the BOLD fMRI signal to the influence of draining veins, however, some have suggested that the initial dip in the BOLD response might provide better spatial resolution. The idea here is that this initial decrease in oxygenation is better co-localized with the neural activity since it is thought to be a more direct consequence of this neural activity. However, as mentioned previously, the existence of this initial dip is still somewhat controversial in humans (Menon, 2001; Menon et al., 2001; Menon et al., 1999; Detre et al., 2002).

The temporal resolution of BOLD fMRI is generally 5 to 8 seconds. The reason for this is that the haemodynamic response is inherently much slower than the underlying neuronal activity. Basically, the BOLD fMRI signal can be seen as a smoothed function of the

underlying neuronal activity. Because of this slowness of the haemodynamic response, fMRI has a relatively poor temporal resolution when compared to methods that more directly measure neuronal activity such as EEG (Detre et al., 2002). Nevertheless, in high magnetic fields some attempts have been made to obtain temporal information about neural processing using BOLD fMRI (Menon, 2001; Menon et al., 2001; Menon et al., 1999; Detre et al., 2002). It appears that under some circumstances the onset and the width of the BOLD response can be used to obtain information about the order in which certain brain areas are activated and the duration of the neural activation. There are, however, a good number of limitations and assumptions involved in using this time-resolved event-related fMRI. Furthermore, while this opens up avenues to explore the relative timing between areas of the brain, determining the absolute timing between these areas of the brain is currently not possible (Menon et al., 2001; Menon et al., 1999; Detre et al., 2002). Despite all these problems though, further development of time-resolved event-related fMRI could provide promising information about temporal processing in the brain.

### **1.3 FMRI TASK DESIGN.**

Over the course of an experiment, the brain is continuously active. This means that the level of oxygenation of the blood varies continuously. However, typically only 1 to 10 percent of this variation in oxygenation is actually related to the task at hand and constitutes the signal of interest. Therefore, great care must be taken when designing an fMRI experiment. The optimal experimental design maximizes the possibility of finding a reliable answer to the research question posed. In other words, the optimal design maximizes both statistical power and the power to draw inferences. This section will describe two different types of experimental design used in fMRI research, namely block designs and event related designs and will explore their strengths and weaknesses regarding their ability to answer certain research questions.

#### **1.3.1 BLOCK DESIGNS.**

The first type of experimental design is the so-called block design, also known as the boxcar (Aguirre & D'Esposito, 2000; Donaldson & Buckner, 2001). This is still the most commonly used experimental design in neuroimaging. In a block design, two or more conditions are alternated in blocks. Each block will have a duration of a certain number of fMRI scans and within each block only one condition is presented. By making the conditions differ in only the cognitive process of interest, the fMRI signal that differentiates the conditions should represent this cognitive process of interest. This is the so-called subtraction paradigm (Aguirre et al., 2000; Donaldson et al., 2001).

Using a block design has one main advantage. The increase in fMRI signal in response to a stimulus is additive. This means that the amplitude of the HRF increases when multiple stimuli are presented in rapid succession. When each block is alternated with a rest condition in which the HRF has enough time to return to baseline, a maximum amount of variability is

introduced in the signal (see Figure 2.3a). Therefore, block designs offer considerable statistical power (Aguirre et al., 2000; Donaldson et al., 2001).




Figure 1.3a: The HRF function for a block design. The x-axis displays the number of volumes. One volume is collected every 3 seconds. The y-axis displays the relative signal change in percentages. The tick marks on the top of the figure show the stimulus onsets. In this example, 8 stimuli spaced 1500 milliseconds apart are presented in each task block. Therefore, each task block lasts for 12 seconds alternated with 12 seconds rest allowing the HRF to return to baseline levels. Especially note the large signal change (up to 8%).

However, because within each block only one condition is presented, randomization of stimulus types is not possible within a block. This makes the type of stimulus within each block very predictable. As a consequence, participants may become aware of the order of the events (Aguirre et al., 2000; Donaldson et al., 2001). Another problem with block designs is related to the use of the subtraction paradigm (Aguirre et al., 2000). The main assumption of the subtraction paradigm is the idea of pure insertion. This means that a cognitive process can be added to a set of already active cognitive processes without affecting them. If this assumption fails, the difference in the fMRI signal between two conditions that is supposed to reflect the cognitive process of interest will in reality reflect the interaction between the cognitive process of interest and the already active cognitive processes. It is because of these problems associated with block designs that it is often very hard to draw solid conclusions from fMRI experiments using block designs. So while block designs have high statistical power, they have a low power to draw inferences (Aguirre et al., 2000).

### 1.3.2 EVENT RELATED DESIGNS.

The second type of experimental design is the so-called event related design (Aguirre et al., 2000; Donaldson et al., 2001). In an event related design the course of the HRF following each stimulus presentation is estimated. The multiple HRF's following a single type of stimulus can be averaged. This allows more real world testing, however, the statistical power of event related designs is inherently low, because the signal change in the BOLD fMRI signal following a single stimulus presentation is small (see Figure 2.3b) (Aguirre et al., 2000; Donaldson et al., 2001).




Figure 1.3b: The HRF function for an event related design. The x-axis displays the number of volumes. One volume is collected every 3 seconds. The y-axis displays the relative signal change in percentages. The tick marks on the top of the figure show the stimulus onsets. In this example, 1 stimulus is presented every 24 seconds. Especially note the small signal change (less than 2%), compared to block designs.

A number of features have been critical in the development of event related designs.

The first important feature is the invention of fast scanning techniques. In event related designs the course of the HRF following a single stimulus presentation is estimated. For this estimation of the HRF to be accurate the BOLD fMRI signal must be measured at many different time points after presentation of the stimulus. This can only be done by fast scanning techniques (Donaldson et al., 2001).

The second important feature is the discovery that the BOLD fMRI signal is very sensitive. Neuronal activation lasting as short as 34 milliseconds still produces a measurable

change in the BOLD fMRI response. This means that even stimuli of very short duration can be used in an event related design (Donaldson et al., 2001).

The third and most important feature is the discovery that the BOLD fMRI shows a roughly linear response to repeated presentations of stimuli. In other words, when two or more stimuli are presented in relatively rapid succession, the BOLD fMRI response increases roughly linearly (Dale & Buckner, 1997; Donaldson et al., 2001). This approximate linearity is important in event-related designs where stimuli are often presented in rapid succession. When stimuli are closely spaced together there is an overlap in the BOLD fMRI responses. The discovery that the BOLD fMRI response increases roughly linearly in response to successive presentations of stimuli, means that the HRF for an individual stimulus can be estimated despite the overlap in HRF's for successive stimuli (Aguirre et al., 2000; Donaldson et al., 2001; Dale et al., 1997). There are studies that seem to suggest though that when stimuli are presented less than 2 seconds apart, this linearity does not completely hold (Friston, Josephs, Rees & Turner, 1998; Pfeuffer, McCullough, van de Moortele, Ugurbil, & Hu, 2003; Ugurbil et al., 2003; Wager, Vazquez, Hernandez, & Noll, 2005).

The approximate linearity of the BOLD response in response to successive presentations of stimuli, allows stimuli to be presented in relative quick succession in an event-related fMRI design. A logical follow-up issue is the question of the optimal spacing in time between the onsets of successive stimulus presentations that maximizes the statistical power in an event related design. The answer to this question depends strongly on the type of event-related design and the effect of interest. When the effect of interest is a difference effect between the effects of multiple stimulus condition (for example, comparing condition A with condition B), then maximal power is obtained when the order of the stimulus conditions is randomized and the spacing in time between stimuli is minimized (as long as the assumption of linear summation of BOLD responses still holds) (Joseph & Henson, 1999). When the

effect of interest is a main effect of a stimulus condition, then as for the difference effect maximal power is obtained when the order of the stimulus conditions is randomized and the spacing in time between stimuli is minimized. However, for main effects this only holds when some measure is taken to add additional predicted variability into the BOLD signal. The problem is that when stimuli of the the same condition are presented rapidly and closely spaced in time, there is almost no variability in the HRF (the BOLD response saturates). Therefore, statistical power is low. However, the power can be increased by adding predicted variability to the HRF. This additional predicted variability can be added either by introducing a latency jitter (Dale, 1999) or by inclusion of null events (Friston, Zarahn, Josephs, Henson & Dale, 1999; Joseph et al., 1999).

A latency jitter can be introduced by employing a variable inter stimulus interval (ISI) instead of a fixed ISI. When a fixed ISI is used the time between the onsets of successive stimuli is always the same. When a variable ISI is used the time between the onsets of successive stimuli varies. The crucial point is randomizing the duration of the ISI according to an exponential distribution (Dale, 1999). This means that many brief ISI's are alternated with a few long ISI's in which the HRF can return to baseline. When a randomized variable ISI is used the statistical power increases with a decreasing ISI (as long as the assumption of linear summation of the BOLD responses still holds) (Dale, 1999). The second way to add predicted variability to the HRF is by inclusion of null events (Friston et al., 1999; Joseph et al., 1999). In the case of a null event, no stimulus presentation occurs. Like the introduction of a laterncy jitter, the inclusion of an occasional null event adds additional predicted variability to the HRF by allowing the HRF to return (partly) to baseline. It is important to note here however, that even though the statistical power of event related designs can be improved by the above mentioned procedures, it is still inherently low compared to the statistical power of block designs.

Despite the inherent low power of event related designs, there are a number of advantages in using this technique. The main advantage is that it allows for randomization of trials, since it is no longer necessary to group trials of the same type together as in a block design (Aguirre et al., 2000; Donaldson et al., 2001). The importance of randomization is that potential confounds like habituation, anticipation and strategy effects are minimised. This increases the power to draw solid conclusions from an experiment and so increases the power to draw meaningful inferences. Another advantage of event related designs over block designs is that they allow for removal of certain trials (Donaldson et al., 2001). For example, if a subject has to make a response to a stimulus, it can be desirable to be able to remove the HRF following stimuli associated with the wrong responses post hoc. Finally, some experimental questions cannot be answered using a block design (Donaldson et al., 2001). If, for example, the response to an infrequently occurring stimulus appearing in a series of frequently occurring stimuli is the focus of the experiment (so-called odd-ball experiments), using a block design is impossible. When stimuli have to be presented in blocks, it is by definition impossible to have an infrequently occurring stimulus. These kinds of experimental questions have to be addressed using an event related design (Donaldson et al., 2001).

## **1.4 FMRI ANALYSIS.**

In a typical fMRI experiment a measurement of the entire brain, known as a volume, is collected every 2 to 4 seconds, resulting in hundreds of collected brain volumes per experiment for each subject. The ultimate goal is determining what areas of the brain show a significant increase in the BOLD fMRI response under certain task conditions. To achieve this goal, a number of steps must be performed on the fMRI dataset. These steps can roughly be subdivided into a step where the dataset is temporally adjusted, steps where the dataset is spatially adjusted and a stage of statistical analysis. The step where the dataset is temporally adjusted is known as slice timing correction. The steps where the dataset is spatially adjusted are spatial realignment, spatial coregistration, spatial normalization and spatial smoothing. The last step is the stage of statistical analysis (Frackowiak, Friston, Frith, Dolan, & Mazziotta, 1997). The following paragraphs will discuss each of these steps and their consequences for data interpretation.

### **1.4.1 SLICE TIMING CORRECTION.**

In an fMRI dataset each volume of the entire brain consists of a number of slices and each slice consists of a number of voxels. Generally, the different 2D slices making up a 3D volume are not all collected at the exact same time (see also paragraph 2.1.6 of this chapter). However, in the statistical analysis the assumption is made that the entire volume is collected at one point in time, so each voxel in a volume is assumed to represent the same moment in time. The result is that it might seem as though the same change in the BOLD fMRI response starts at an earlier time for slices that are acquired later in time than for slices that are acquired earlier in time (see Figure 2.4a).




Figure 1.4a: The BOLD fMRI response over time. The red arrow represents the time of acquisition for an early slice in the volume, the blue arrow represents the time of acquisition for a later slice in the volume. When the assumption is made that the entire volume is collected at one point in time (as is done in the statistical analysis), it seems as though the BOLD fMRI response begins earlier for the slice that is acquired at a later time (blue curve) than for the slice that is collected at an earlier time (red curve).

To solve this, the individual slices of a volume must be adjusted in the temporal domain. This is done by performing a temporal correction for the differences in acquisition time between the slices. This is referred to as slice timing correction (Smith, 2001). First, a decision must be made as to which slice is going to be the reference slice. In other words, which time of acquisition is going to be taken as the point in time that the entire volume of the brain was collected? Usually, either the first slice or the middle slice of the image of the brain is taken as the reference slice. After a certain slice has been chosen as the reference slice, all the other slices in the image of the entire volume of the brain are shifted in time using an interpolation method. The slices of a volume that are collected earlier in time than the reference slice are weighted with the same slice in the subsequent volume. The slices of a

volume that are collected later in time than the reference slice are weighted with the same slice from the previous volume. For example, if a volume consists of 10 slices with the fifth slice being the reference slice, the second slice would be weighted with the second slice of the subsequent volume and the eighth slice would be weighted with the eighth slice from the previous volume. The reference slice is the only slice that is not shifted in time. The end result is that each voxel in each slice will approximately have the value that it would have had, had it been obtained at the same point in time as the voxels in the reference slice. In other words, the same change in the BOLD fMRI response now starts at the same time for each voxel in each slice in a volume (Smith, 2001).

Which slice to choose as the reference slice depends on where in the brain interesting activations are expected to occur. As the difference in time of the phase shifted slice from the reference slice increases, the artefacts introduced by this phase shift also increase. This is basically because the interpolation method used to shift the slices of a volume in time is not perfect. The further in time slices have to be shifted, the larger the errors in the interpolation. Therefore, it is important to choose the reference slice as close as possible to the region where the interesting activations are expected to occur, so that the interesting activations are not confounded by errors of interpolation.

#### 1.4.2 SPATIAL REALIGNMENT.

During a fMRI experiment, hundreds of fMRI brain volumes are collected per subject in a single time series. Even though subjects are usually instructed to move as little as possible inside the scanner, some head movement is unavoidable. The main result of head movements is that the same voxel does not necessarily represent the same location in the brain throughout time (Ashburner & Friston, 2003a; Ashburner & Friston, 2000; Brammer,

2001). The statistical analysis, however, assumes that the same voxel does represent the same location in the brain throughout time.

When the same voxel over time ‘moves’ from a location of the brain with a low fMRI signal to a location of the brain with a higher fMRI signal, while the statistical analysis assumes that this voxel represents the same location in the brain throughout time, it appears as though there was an increase in the fMRI signal for that voxel over time while in reality there was no increase in the fMRI signal. This has two potential consequences. Firstly, when movements are correlated with task performance (for example, if the head movement always occurs at a certain time after stimulus presentation) these false increases in the fMRI signal ultimately appear as false activations in the brain after the statistical analysis (Ashburner et al., 2003a; Ashburner et al., 2000; Brammer, 2001). Secondly, even when these intensity changes in the voxel caused by the movement of the head are not correlated with task performance they will add to the noise in the signal, thereby worsening the signal to noise ratio, which decreases the statistical power (Ashburner et al., 2003a; Ashburner et al., 2000; Brammer, 2001).

The removal of movement effects is done for each subject separately and is referred to as spatial realignment or, alternatively, as within modality rigid registration (Ashburner et al., 2003a; Ashburner et al., 2000; Brammer, 2001). The first step is choosing a reference volume. This can be the first image in the time series, but the mean image of the time series is also regularly used. All the other volumes in the same time series are then repositioned until they are in the same position as the reference volume. Only the position of the brain is changed and not the size or shape. Therefore, this repositioning treats the head as a rigid object and is therefore also known as a rigid body transformation. For this realignment a minimization algorithm for the least mean square difference between the volumes in the time series and the reference volume is used. This means that for each volume the squared difference between the

volume in the time series and the reference volume is minimized (see Figure 2.4b). The new value of the fMRI signal for a voxel after the realignment is estimated by interpolation from the values of the fMRI signal of neighbouring voxels (Ashburner et al., 2003a; Ashburner et al., 2000; Brammer, 2001).




Figure 1.4b: Illustration of the squared difference between the volume in the time series and the reference volume. **A:** The reference volume. **B:** the volume that has to be adjusted. **C:** The difference between A and B. **D:** The squared difference between A and B.

It should be noted that realignment not only adjusts for actual head movement, but also for apparent movement. As the fMRI scanner heats up during a session it appears as though the head drifts slightly. This is an artefact arising from the scanner and is also corrected for by realignment.

Unfortunately, there are a number of limitations to realignment. The first limitation is that, when the head movements are too large, the minimization algorithm might get stuck in a local minimum (Ashburner et al., 2003a; Ashburner et al., 2000; Brammer, 2001). Although, some new approaches to minimize the least mean square difference between the volumes in the time series and the reference volume have been proposed that severely reduce this risk (Jenkinson, Bannister, Brady, & Smith, 2002). The second limitation is that the brain is not rigid. Heart-beat and respiration cause variations in the shape and the size of the brain. Usually, this source of movement is ignored (Brammer, 2001). The third and main limitation is that even perfect realignment will not remove all movement related variance (Ashburner et al., 2003a; Ashburner et al., 2000; Brammer, 2001). The problem is that, in case of a

movement, the image not only moves, but also fundamentally changes. This will be explained below.

Firstly, the time it takes to acquire one image of the brain is usually shorter than the time it takes for the hydrogen nuclei to relax after the radiofrequency pulse. This means that, by the time the next radiofrequency pulse occurs, there is still a residual effect from the previous radiofrequency pulse on the hydrogen nuclei. This is not really a problem when the head does not move in the scanner, because in that case the residual effect from the previous radiofrequency pulse is predictable. However, when the head does move, the residual effect from the previous radiofrequency pulse on the hydrogen nuclei becomes unpredictable (Ashburner et al., 2003a; Ashburner et al., 2000; Brammer, 2001).

Secondly, remember from the second section in this chapter that inhomogeneities in the magnetic field result in a decrease in the fMRI signal. Although every attempt is made to make the overall magnetic field as uniform as possible, the magnetic field will still contain small inhomogeneities. These inhomogeneities cause distortions (also known as susceptibility artefacts) in the image of the brain, particularly on the tissue-air and tissue-bone boundaries. When the head now moves these distortions change in a nonlinear fashion. This artefact is known as a susceptibility-by-movement interaction (Andersson, Hutton, Ashburner, Turner, & Friston, 2001; Jezzard & Clare, 1999). This susceptibility-by-movement interaction can be largely dealt with by ‘unwarping’ the fMRI dataset. After realignment, the amount of movement is known for each subject. If it is also known how the magnetic field changes when the subject moves, the resulting distortions in the image can be estimated. The variance in the dataset that is a consequence of this susceptibility-by-movement interaction can therefore be estimated and removed (Andersson et al., 2001).

One final way to remove all movement related variance is to use a linear regression to remove any variance in the signal that is correlated with both movements during the present

scan and movements during the preceding scan. However, when the movements are correlated with task performance, there is a risk of accidental removal of interesting activations (Ashburner et al., 2003a; Ashburner et al., 2000; Brammer, 2001). Ultimately, it seems obvious that head movement should be minimized, for example, by using a head constraint.

### 1.4.3 SPATIAL NORMALIZATION.

During a fMRI experiment, datasets are usually collected for several subjects. However, all brains differ in orientation, size and shape. Therefore, the orientation, size and shape of the brains of individual subjects is changed to match the orientation, size and shape of a standard brain (Ashburner & Friston, 2003b; Ashburner et al., 2000). There are a number of reasons why making these different brains more alike in size and shape is desirable. Firstly, when the same voxels in the brain of each subject represent the same anatomical location, comparisons over and between different subjects are possible. Secondly, when different brains are mapped to a certain standard brain, communication of the anatomical loci of interesting effects between different research groups becomes less arbitrary. This matching of individual brains to a standard brain is known as spatial normalization (Ashburner et al., 2003b; Ashburner et al., 2000; Jenkinson, 2001).

The goal of spatial normalization is the matching of the orientation, size and shape of each individual brain to the orientation, size and shape of a standard brain. Obviously, choosing the right standard brain is important. The standard brain most commonly used is the MNI152 template which is based on the coordinate system described by Talairach & Tournoux (1988). However, different scanners have unique characteristics and therefore research groups often create their own standard brain. Even though this improves the match of the individual brain to the standard brain for the research groups, it makes comparisons of results between research groups more tricky (Ashburner et al., 2003b; Ashburner et al., 2000).

The matching of the orientation, size and shape of each individual to the orientation, size and shape of the template is done using a number of linear parameters. Each of these linear parameters changes the entire brain in the same way and these are also known as affine parameters. In addition, a set of nonlinear cosine basis functions can be used. These enable local changes in the brain, also known as warping, so that it can be even better matched to the template (Ashburner et al., 2003b; Ashburner et al., 2000). A minimisation algorithm utilising the least mean square difference between the brain volumes of the subject and the template volume is used and the new value of the fMRI signal for a voxel after the normalization is estimated by interpolation from the values of the fMRI signal of neighbouring voxels (Ashburner et al., 2003b; Ashburner et al., 2000). Obviously, some constraints have to be imposed on this minimisation procedure, since a large difference in the amount of warping between adjacent areas would lead to unrealistic results. The constraint usually imposed is that the overall warp of the brain has to be smooth. Practically, this means minimising the difference in the amount of warping between adjacent brain areas. The result is that normalization matches the overall orientation, size and shape of the brain of each subject to the template, but not the individual sulci (Ashburner et al., 2003b; Ashburner et al., 2000). Usually, this overall matching of normalization is not considered a problem, since precise matching of each individual brain to the template would be unrealistic due to inter-subject differences (Ashburner et al., 2003b; Ashburner et al., 2000).

#### 1.4.4 SPATIAL CO-REGISTRATION.

During spatial realignment, different scans within the same modality are aligned together. This means that EPI fMRI scans are repositioned until they are in the same position as a reference EPI fMRI scan. Sometimes, however, it can be desirable to align scans from the same subject, but from different modalities together. For example, sometimes it can be

advantageous to align a low resolution EPI fMRI scan to the same individual's high resolution T1-weighted scan. This alignment of scans from different modalities is known as spatial co-registration or alternatively as between-modality rigid registration (Ashburner et al., 2003a; Jenkinson, 2001).

There are a number of reasons to perform a co-registration. Firstly, to accurately determine the location of a statistical increase in the fMRI signal for a subject, it is helpful to be able to overlay the statistical map on a high resolution anatomical scan of that same subject. This means that the fMRI dataset and the high resolution anatomical scan have to be in alignment. Secondly, co-registration can sometimes help with spatial normalization. The high resolution anatomical scan is more detailed than the fMRI scans, so normalization of the high resolution anatomical scan to a standard brain often leads to better results than matching the fMRI scans to the standard brain. If the high resolution anatomical scan and the fMRI scans are first co-registered, the parameters that are used to match the anatomical to the standard brain can then be applied to the fMRI dataset (Ashburner et al., 2003a; Jenkinson, 2001).

It would seem intuitive to just apply the same algorithms as used during spatial realignment when registering images from different modalities. However, two images from the same modality will have similar intensities for specific tissue types (for example, water is always dark on a T1 scan). Two images from different modalities will usually not have similar intensities for specific tissue types (for example, water is dark on a T1 scan, but bright on a T2 scan) so the matching of the intensities of the different images (as is done in spatial realignment) is not possible. The problem here is that even when the alignment is perfect, the minimum squared difference in intensities between the two images would still be large. One way to align two images from different modalities is by using intermodal similarity functions (Ashburner et al., 2003a; Jenkinson, 2001). First, each modality is normalized to a standard

brain of the same modality. After this normalization, one of the images is segmented into separate areas for each tissue type. The result is a number of segmented areas, where each segmented area contains only the locations in the brain with a specific intensity range. Co-registration is subsequently performed by mapping each of these segmented areas onto the second image. The idea is that, if both images are in alignment, the intensity within a segmented area should be constant for the second image, since it reflects the same tissue type for both images. For each segmented area, co-registration with the second modality is again performed using an algorithm that minimizes the squared differences (Ashburner et al., 2003a; Jenkinson, 2001).

Another way to align two images from different modalities is by using mutual information theory (Ashburner et al., 2003a; Jenkinson, 2001). Mutual information theory is based on the principle that even though the same tissue appears different in different modalities, intensity differences between different tissues are preserved in each modality. So even though water appears dark on a T1 scan and bright on a T2 scan, in both these modalities water has a different intensity from for example bone tissue (Ashburner et al., 2003a; Jenkinson, 2001). If now the intensities in each image are represented as a histogram, a so-called joint histogram can be created (see Figure 2.4c). This histogram has the intensities of one modality on the x-axis and the intensity of the other modality on the y-axis. Now, for each voxel the value it has on the x-axis and the value it has on the y-axis are determined. If the images are in perfect alignment, the histogram will show a 'blob' for each tissue type. The idea is that each tissue type has a constant value on both the x-axis and the y-axis, which is different from the other tissue types in each image. As alignment becomes less perfect, the amount of noise in the joint histogram increases. Between modality registration basically tries to minimize the amount of noise in this joint histogram (Ashburner et al., 2003a; Jenkinson, 2001).




Figure 1.4c: Between modality registration based on mutual information theory works by minimizing the noise in a joint histogram. The picture on the left shows two scans that are perfectly aligned. Even though different tissues will have different intensities in the different modalities, the joint histogram shows a distinct ‘blob’ for each tissue type. The picture on the right shows two scans that are not aligned. The joint histogram now shows a lot of noise in addition to the ‘blobs’ for each tissue type.

#### 1.4.5 SPATIAL SMOOTHING.

At a first glance it seems counterintuitive to spatially blur images when spatial precision is a factor of primary interest. However, there are a number of very good reasons to perform smoothing. Firstly, smoothing the dataset increases the signal to noise ratio in the fMRI signal by removing the high spatial frequencies. In fMRI the effects of interest are produced by changes in blood flow. Changes in blood flow are predominantly expressed in low spatial frequency bands (with wavelengths in the order of several millimetres). Noise is usually present in higher spatial frequency bands. Secondly, smoothing removes small frequency differences, so comparisons across subjects are made easier. Thirdly, smoothing the dataset helps to satisfy the requirements for applying Gaussian Field Theory to correct for multiple comparisons in the ensuing statistical analysis by making the data more normally distributed (Smith, 2001).

Smoothing is generally performed by convolving the 3D volume with a 3D Gaussian kernel. Basically this means that every data point is multiplied by a curve in the shape of a 3D normal distribution. The shape of the 3D smoothing curve should match the spatial shape of the signal of interest, so that frequencies matching the frequencies of the 3D smoothing curve are emphasized and frequencies not matching the frequencies of the 3D smoothing curve are filtered out. The shape of the smoothing curve is defined by the Full Width Half Maximum (FWHM). This is the width of the curve at half of the maximum and is usually defined in millimetres. The FWHM chosen for the smoothing curve is typically two or three times the voxel size. However, it is important to keep in mind that the amount of smoothing will influence the size of the areas where a significant increase in the BOLD fMRI signal can be detected in the statistical analysis. Using a smoothing curve with a FWHM of 8mm will result in not being able to detect areas of significant increase in the BOLD fMRI signal that are smaller than 8mm (Smith, 2001).

#### 1.4.6 STATISTICAL ANALYSIS.

The goal of the statistical analysis is to determine what areas of the brain are significantly activated under certain task conditions. Most commonly the General Linear Model (GLM) is used for this statistical analysis. The statistical analysis consists of a number of steps. First the BOLD fMRI response must be modelled. Secondly, the parameters of this model must be estimated. Finally, it must be determined whether there is any evidence for a statistically significant increase in the BOLD fMRI response in response to a task condition. These steps will be explained below. First, however, two more preprocessing steps must be discussed; intensity normalization and temporal filtering. Since these preprocessing steps are usually part of the statistical analysis, they are discussed in this paragraph.

Intensity normalization is the re-scaling of the intensities to a common mean intensity value (Kiebel & Holmes, 2003; Smith, 2001). One way to perform intensity normalization is within a continuous fMRI dataset. First, the mean intensity of a fMRI volume is determined. Subsequently, all voxels in that volume are re-scaled to a certain pre-determined mean intensity value. The idea is that re-scaling to a common mean will get rid of global variations in intensity between volumes that are a result of scanner properties (for example heating) (Kiebel et al., 2003; Smith, 2001). There is one huge problem with this method of intensity normalization however. Imagine certain voxels in a volume showing a strong increase in response to the task condition. Intensity normalization will now re-scale all voxels in the whole volume to a certain mean intensity value. However, the result is also that the voxels showing less increase in the fMRI signal in that same volume might be re-scaled to such an extent that they appear to show a decrease in the fMRI response. This is one of the reasons that intensity normalisation within a continuous fMRI dataset is usually not performed. Global variations in intensity that are not part of the actual signal can generally be dealt with using highpass temporal filtering (Smith, 2001).

There is still the problem though that some sort of intensity normalization between different fMRI datasets (obtained from for example different subjects) must be performed to make comparisons over and between these different fMRI datasets valid. This form of intensity normalization is commonly referred to as grand mean scaling (Kiebel et al., 2003; Smith, 2001). In this case each continuous fMRI dataset is re-scaled to a pre-specified common mean for the dataset. The result is that all the different fMRI datasets have the same mean intensity value so that comparisons over and between these datasets are meaningful (Kiebel et al., 2003; Smith, 2001).

Temporal filtering is the removal of certain unwanted temporal frequencies from the dataset (Kiebel et al., 2003; Smith, 2001). Low frequencies in a BOLD fMRI time series often

reflect physiological effects like breathing and heartbeat or scanner drift (as the scanner heats up in the course of an experiment, the signal will appear to drift). These low frequencies can be dealt with by using a highpass temporal filter. A highpass filter removes all frequencies lower than a specified cut-off frequency from the dataset (Smith, 2001). By reducing these low frequency drifts from the time series of each individual voxel, temporal filtering increases the signal to noise ratio (Kiebel et al., 2003; Smith, 2001).

After intensity normalization and temporal filtering have been applied, the dataset is ready for the statistical analysis. As mentioned, most commonly the General Linear Model (GLM) is used (Kiebel et al., 2003; Friston et al., 1995; Lange, 2000; Worsley, 2001). The basic formula used in the general linear model is  $Y(j) = \beta * x(j) + c + E(j)$ . In this formula  $j$  indexes the volume,  $Y$  is the observed value for a voxel in volume  $j$ ,  $\beta$  is the slope of the linear regression line,  $x$  is the value of the predictor variable for a voxel in volume  $j$ ,  $c$  is the intercept of the linear regression line and  $E$  is the error in volume  $j$ . In the statistical analysis of fMRI this formula is in essence solved for each voxel in each volume in the fMRI time series separately. In other words, the statistical analysis of fMRI datasets is massively univariate (Kiebel et al., 2003; Friston et al., 1995; Lange, 2000; Worsley, 2001).

When the formula for the general linear model is re-formulated in terms of matrices it becomes  $\mathbf{Y} = \mathbf{X} * \mathbf{B} + \mathbf{E}$  (see Figure 2.4d). Now  $\mathbf{Y}$  is a matrix containing all the observed data with a column for each voxel and a row for each fMRI volume,  $\mathbf{X}$  is a matrix containing the predicted data (often known as the so-called design matrix) with a column for each predictor variable and a row for each fMRI volume,  $\mathbf{B}$  is a matrix containing the slopes and intercepts (also known as a parameter matrix) and  $\mathbf{E}$  is a matrix with the normally distributed error terms (Kiebel et al., 2003; Friston et al., 1995; Lange, 2000; Worsley, 2001). As mentioned before, in the statistical analysis the BOLD response must first be modelled. This modelled BOLD response is contained in the design matrix. Subsequently the parameters (slope and intercept

of the regression line) of this model must be estimated. This is the estimation of the parameter matrix. Finally, this parameter matrix is compared to the error matrix for each voxel individually, culminating in a test statistic (t-value and finally Z-value) for each voxel.

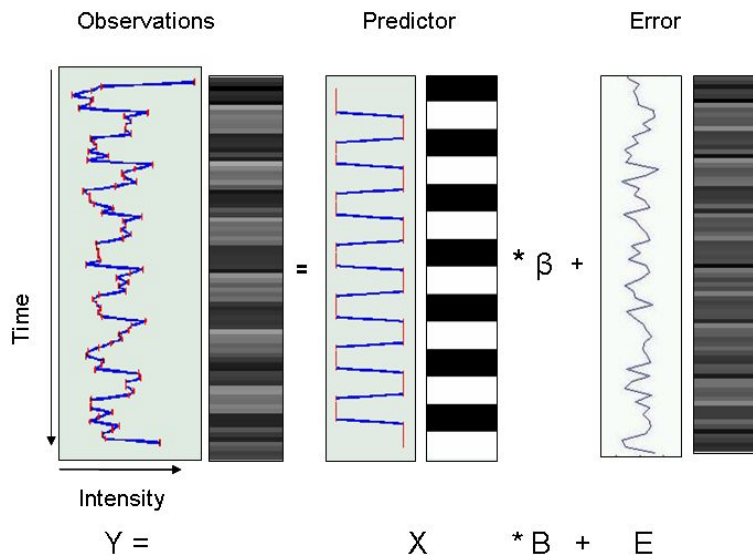


Figure 1.4d: Graphical representation of the GLM as used during the statistical analysis of fMRI data. The figure shows a GLM for a single voxel with one predictor variable. For each of the terms in the GLM (except the to be estimated B) both a graphic of the time course (relative left) and a representation of the time course as a column in a matrix (relative right) are shown.

In the design matrix the general shape of the BOLD response over time is modelled. On the most simple level each column (remember that each column reflects a predictor variable) in the design matrix reflects a certain stimulus function (often an on/off block curve where the amplitude of the stimulus function is either maximal or 0; when the stimulus is present/on the amplitude of the stimulus function is maximal, when the stimulus is absent/off the amplitude of the stimulus function is 0) smoothed with the BOLD fMRI response function (see Figure 2.4d). However, other predictor variables can also be modelled in the design matrix for example, covariates like task performance and confounds like slow drifts (Kiebel et al., 2003; Friston et al., 1995; Lange, 2000; Worsley, 2001). In addition to the predictor variables, derivatives of these main predictor variables are often modelled. The most commonly used derivatives are the temporal derivative and dispersion derivative. The

temporal derivative allows for variations in the onset of the BOLD response. The dispersion derivative allows for variations in the width of the BOLD response.

After the design matrix and therefore the global shape of the BOLD response following a stimulus presentation has been modelled, this model has to be fit separately to the time course of each voxel. This is the estimation of the parameter matrix using linear regression. For each predictor variable a parameter is obtained that optimises the fit of the overall BOLD response modelled in the column of the design matrix to the observed data for each voxel. To clarify, the design matrix predicts for each predictor variable the overall shape of the BOLD response over time and the parameter matrix is estimated to fit this overall shape of the BOLD response to the specific time course of each voxel (Kiebel et al., 2003; Friston et al., 1995; Lange, 2000; Worsley, 2001). It would be intuitive to think that the error matrix is now the variance in the model that is left unexplained after estimating the parameter matrix. However, in BOLD fMRI datasets there is the problem of temporal autocorrelation. In an fMRI time series the successive data points are not independent. The time it takes to acquire a single data point is usually shorter than the time it takes for HRF to return to baseline. This means that, successive data point acquisitions (partly) sample the same HRF. The result is that successive data points are correlated in time, which makes the estimation of the error variance incorrect. This temporal autocorrelation therefore has to be modelled to avoid inaccurate estimates of the error. Only after compensation for this temporal autocorrelation does the error matrix reflect the unexplained variance in the model (Kiebel et al., 2003; Friston et al., 1995; Lange, 2000; Worsley, 2001).

Now that the parameters and error have been estimated, the next step is to use a t-test to determine for each voxel separately whether a specific linear combination of estimated parameters is significantly different from zero (Kiebel et al., 2003; Friston et al., 1995; Lange, 2000; Worsley, 2001). In other words, does any linear combination of the predictor variables

defined in the design matrix explain a significant amount of variance in the BOLD fMRI dataset. Basically, the explained variance given by a linear combination of parameters in the parameter matrix is statistically compared to the unexplained variance in the error matrix. Imagine a design matrix with 2 columns. The first column in the design matrix contains the predicted HRF for the stimulus function of a visual stimulus and the second column in the design matrix contains the predicted HRF for the stimulus function of an auditory stimulus. After the estimation of the parameter matrix a parameter for both columns in this design matrix is obtained together with an estimate of the error. To determine for each voxel whether the response to visual stimuli is for example significantly larger than the response to auditory stimuli it is necessary to statistically compare the difference in the estimated parameters of the two columns to the error. The question here is whether in any voxel the subtraction of the estimated parameter for the second column in the design matrix from the estimated parameter for the first column explains a significant amount of variance in the data when compared to the error variance. In principle any linear combination of predictor variables can be statistically assessed in this manner. The result is a t-score for each voxel and for each linear combination of predictor variables reflecting how well the observed BOLD fMRI timecourse is explained (Kiebel et al., 2003; Friston et al., 1995; Lange, 2000; Worsley, 2001). As a final step often these t-scores are converted to Z-scores.

#### 1.4.7 CHOOSING A STATISTICAL THRESHOLD.

After the previous statistics the end result is a Z-score for each voxel and for each linear combination of predictor variables. Usually, the probability threshold set in a statistical test to determine whether a specific linear combination of columns in the design matrix differs significantly from zero is 5%, i.e. it has a 5% or less probability of occurring by chance alone. However, this statistical test is performed for each voxel separately. If twenty thousand voxels

are tested at a probability threshold of 5%, a thousand voxels will show erroneous significant activation by chance (i.e. a statistical type I error). These ‘fake’ activations are known as false positives. The problem is that the probability of finding false positives for all univariate tests combined (or the family-wise error) is actually a lot higher than the probability of finding false positives for each individual test. The desired solution is to choose a statistical threshold that keeps the family-wise error below 5%. One way to achieve this is by applying a Bonferroni correction. In a Bonferroni correction the probability threshold is divided by the number of univariate tests (one for each voxel in this case). The probability threshold for one voxel then becomes  $5/20000 = 0.00025\%$ . The end result is that the probability threshold for all univariate tests combined is the desired 5%. However, a Bonferroni correction assumes that the different univariate tests are independent. In fMRI the signal in adjacent voxels is often correlated so the amount of independent tests is actually lower than the amount of univariate tests performed. Therefore a Bonferroni correction will be too stringent a correction. As a consequence, applying a Bonferroni correction will not only result in a low probability of detecting false positives, but also in missing many ‘real’ activations (i.e. a statistical type II error). In other words, a Bonferroni correction dramatically reduces statistical power (Nichols & Hayasaka, 2003).

Ideally, the statistical threshold should be chosen to keep the family-wise error below 5% while preserving as much statistical power as possible. There are a number of methods that aim to do just that. These solutions include using a minimum cluster size, utilising Gaussian random field theory, false discovery rate correction and region of interest analysis. These methods will each be discussed in turn.

The first solution is using a minimum cluster size. No correction for multiple comparisons is applied; instead a reasonably strict uncorrected probability threshold is used together with setting a lower limit to the cluster size. For example, only clusters of at least 30

adjacent significantly activated voxels are accepted as ‘true’ activations. The logic is that it is highly unlikely that 30 adjacent voxels would all appear significantly activated purely by chance. The drawback of this method is that it is by definition impossible to detect small clusters of real significant activation. Furthermore, if spatial smoothing has been applied (see paragraph 2.4.5 of this chapter), adjacent voxels are not really independent tests: smoothing (as well as interpolation during other preprocessing steps) results in neighbouring voxels often having similar statistical values.

A very popular solution aimed to keep the family-wise error below 5% while preserving as much statistical power as possible uses the Gaussian random field theory (Worsley, Evans, Marrett, & Neelin, 1992). The Gaussian random field theory assumes that the BOLD fMRI signal at each voxel has a normal spatial distribution. When the data has been spatially smoothed with a Gaussian kernel (see paragraph 2.4.5) this assumption will generally be valid. In this method first the number of independent observations, known as resels, is determined. This number of resels is calculated using the FWHM used during smoothing of the data. So, if the dataset was smoothed using an 8mm FWHM smoothing kernel, each resel will be  $8\text{mm}^3$  in size. Next, the so-called expected Euler characteristic is determined. The Euler characteristic is the number of significantly activated clusters at a given Z threshold. As the critical Z-value increases the Euler characteristic will decrease. It turns out that if the number of resels in a volume is known, it is possible to calculate the expected Euler characteristic given a certain Z threshold. The statistical threshold that is now used is when the expected Euler characteristic is less than 0.05. This basically means that the chance of one or more significantly activated clusters being false positives is less than 0.05 (Brett, Penny, & Kiebel, 2003; Nichols et al., 2003; Worsley, 2001; Worsley et al., 1992). This method is far less stringent than a Bonferroni correction while still resulting in an acceptable probability of false positives.

A few years ago, a novel way of choosing a statistical threshold that controls for family-wise error while preserving as much statistical power as possible has been proposed. This method is known as the false discovery rate correction (FDR) (Benjamini & Hochberg, 1995; Genovese, Lazar, & Nichols, 2002). The previously discussed methods of the Bonferroni correction and the theory of Gaussian random fields both strive to keep the probability of a false positive for all univariate tests combined below a certain statistical threshold (usually 5%). The FDR, on the other hand, only attempts to keep the expected proportion of false positives amongst real positives below a certain statistical threshold. The practical consequence is that using a Bonferroni correction or a Gaussian random field method with a statistical threshold of 5% will give a 5% chance of a single false positive. Using the FDR with a statistical threshold of 5% on the other hand, means accepting that 5% of the significant activations in the dataset will be false positives. The statistical threshold determined with the FDR is sensitive to the amount of signal in the dataset. If there is no real signal present in the dataset, the FDR will be just as conservative as the Gaussian random field method. If however, there is a signal present in the dataset, the FDR will be far more sensitive in detecting it than the Gaussian random field method (see Figure 2.4e). The idea is that when there is no signal present in the dataset, the Z-scores of all voxels will be distributed according to a normal distribution. If there is now a signal present, the distribution of the Z-scores will no longer be according to a single normal distribution, but instead will show a second peak reflecting the signal shifted with respect to the first peak reflecting the noise. The distribution of the Z-scores, therefore, can be used as a measure of how much signal is present in the dataset. Depending on the amount of signal in the dataset, the FDR threshold will be more or less conservative (Benjamini et al., 1995; Genovese et al., 2002).

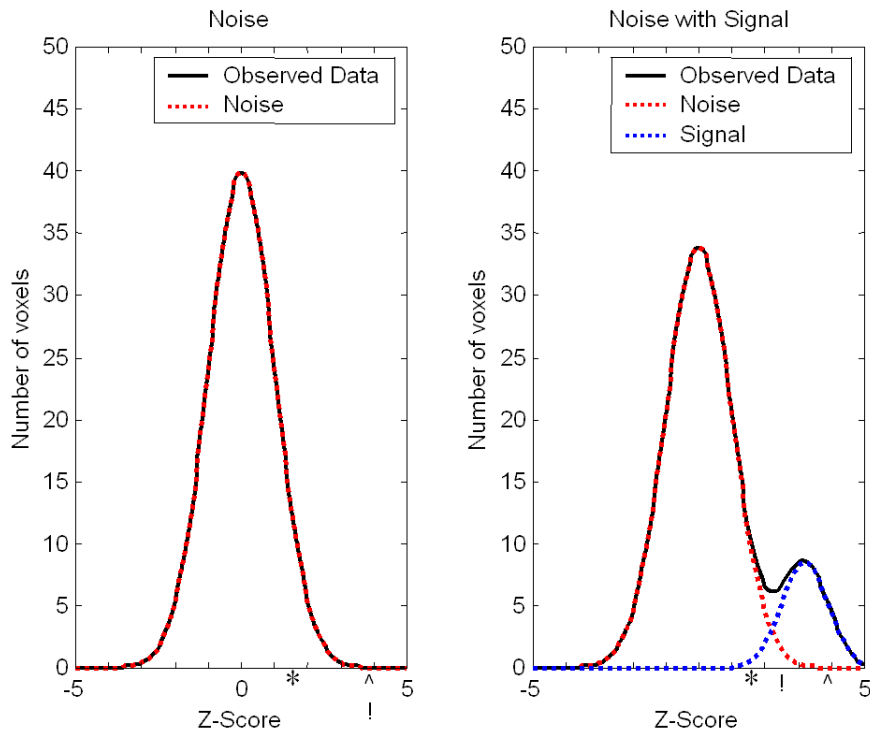


Figure 1.4e: A graphical illustration of the FDR. The left figure shows the distribution of Z-scores when there is no signal present in the dataset. The right figure shows the distribution of Z-scores when there is (in addition to the noise) a signal present in the dataset. When there is no signal present, the Bonferroni corrected threshold (indicated with  $\wedge$ ) and the FDR corrected threshold (indicated with  $!$ ) will be equally conservative. However, when there is a signal present in the data, the Bonferroni corrected threshold remains unchanged, while the FDR corrected threshold shows sensitivity to the amount of signal in the data and is now more lenient. The threshold for a t-test without correction for multiple comparisons is indicated with  $*$ .

The final method to keep the family-wise error below 5% while preserving as much statistical power as possible is an extension of the Gaussian random field theory. In this method however, in addition to using the Gaussian random field theory, the number of multiple comparisons is reduced by only selecting a specific subset of voxels for analysis, thus increasing the statistical power even further. This approach is known as a region of interest (ROI) analysis. There are two ways to perform a ROI analysis. Either a ROI can be defined where only the voxels in that ROI are statistically assessed. This is sometimes also known as a Small Volume Correction (SVC). Alternatively, the average activity within a ROI can be calculated. In this case only this average activity is statistically assessed, so only one

statistical comparison is made. A ROI analysis can be very useful when there are precise predictions about the locations of the areas of significant increase in the BOLD fMRI response. One way to define a ROI is by running a so-called localizer experiment. Areas of the brain where the BOLD fMRI activation is significantly increased in this localizer experiment can then be used to create a ROI for the analysis of a subsequent experiment. Another way to define a ROI is by selection of certain anatomical coordinates (for example, based on previous literature) or landmarks. An important pre-requisite for carrying out ROI analyses, however, is that the dataset used to define the ROI is not the same as the dataset on which the ROI analysis is performed.

## REFERENCES.

- Aguirre, G. K. & D'Esposito, M. (2000). Experimental design for brain fMRI. In C. Moonen & T. W. Bandettini (Eds.), *Functional MRI* (pp. 369-380). Heidelberg: Springer-Verlag Berlin.
- Andersson, J. L. R., Hutton, C., Ashburner, J., Turner, R., & Friston, K. (2001). Modelling geometric deformations in EPI time series. *NeuroImage*, *13*, 903-919.
- Arthurs, O. J. & Boniface, S. (2002). How well do we understand the neural origins of the fMRI BOLD signal? *Trends in Neurosciences*, *25*, 27-31.
- Ashburner, J. (2002). Another MRI bias correction approach. Presented at the 8th national conference on functional mapping of the human brain, June 2-6, 2002, Sendai, Japan. Available on CD-ROM in NeuroImage 16 (2).
- Ashburner, J. & Friston, K. J. (2000). Image registration. In C.T.W. Moonen & P. A. Bandettini (Eds.), *Functional MRI* (pp. 285-299). Heidelberg: Springer-Verlag Berlin.
- Ashburner, J. & Friston, K. J. (2003a). Rigid body registration. In R.S.J. Frackowiak, K. J. Friston, C. Frith, R. Dolan, C. J. Price, S. Zeki, J. Ashburner, & W. D. Penny (Eds.), *Human brain function* (2nd ed., San Diego: Academic Press.
- Ashburner, J. & Friston, K. J. (2003b). Spatial normalization using basis functions. In R.S.J. Frackowiak, K. J. Friston, C. Frith, R. Dolan, C. J. Price, S. Zeki, J. Ashburner, & W. D. Penny (Eds.), *Human brain function* (2nd ed., San Diego: Academic Press.
- Attwell, D. & Ladecola, C. (2002). The neural basis of functional brain imaging signals. *Trends in Neurosciences*, *25*, 621-625.
- Attwell, D. & Laughlin, S. B. (2001). An energy budget for signalling in the grey matter of the brain. *Journal of Cerebral Blood Flow and Metabolism*, *21*, 1133-1145.
- Beckmann, C., Jenkinson, M., & Smith, S. M. (2003). General multi-level linear modelling for group analysis in fMRI. *NeuroImage*, *20*, 1052-1063.
- Benjamini, Y. & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, *57*, 289-300.
- Brammer, M. J. (2001). Head motion and its correction. In P. Jezzard, P. M. Matthews, & S. M. Smith (Eds.), *Functional MRI: An introduction to methods* (pp. 243-250). New York: Oxford University Press Inc.
- Brett, M., Anton, J.-L., Valabregue, R., & Poline, J.-P. (2002). Region of interest analysis using an SPM toolbox. Presented at the 8th international conference on functional mapping of the human brain, June 2-6, 2002, Sendai, Japan. Available on CD-ROM in NeuroImage 16 (2).
- Brett, M., Penny, W., & Kiebel, S. (2003). An introduction to random field theory. In R.S.J. Frackowiak, K. J. Friston, C. Frith, R. Dolan, C. J. Price, S. Zeki, J. Ashburner, & W. D. Penny (Eds.), *Human brain function* (2nd ed., San Diego: Academic Press.
- Buxton, R. B. (2002). *Introduction to functional magnetic resonance imaging: Principles and techniques*. Cambridge: Cambridge University Press.
- Buxton, R. B. & Frank, L. R. (1997). A model for the coupling between cerebral blood flow and oxygen metabolism during neural stimulation. *Journal of Cerebral Blood Flow and Metabolism*, *17*, 64-72.
- Dale, A. M. (1999). Optimal experimental design for event-related fMRI. *Human Brain Mapping*, *8*, 109-114.
- Dale, A. M. & Buckner, R. L. (1997). Selective averaging of rapidly presented individual trials using fMRI. *Human Brain Mapping*, *5*, 329-340.

- Detre, J. A. & Wang, J. (2002). Technical aspects and utility of fMRI using BOLD and ASL. *Clinical Neurophysiology*, *113*, 621-634.
- Donaldson, D. L. & Buckner, R. L. (2001). Effective paradigm design. In P. Jezzard, P. M. Matthews, & S. M. Smith (Eds.), *Functional MRI: An introduction to methods* (pp. 177-195). New York: Oxford University Press Inc.
- Frackowiak, R. S. J., Friston, K. J., Frith, C. D., Dolan, R. J., & Mazziotta, J. C. (1997). *Human brain function*. San Diego: Academic Press.
- Friston, K. J., Holmes, A. P., Price, C. J., Büchel, C., & Worsley, K. J. (1999). Multisubject fMRI studies and conjunction analyses. *NeuroImage*, *10*, 385-396.
- Friston, K. J., Holmes, A. P., Worsley, K. J., Poline, J.-P., Frith, C. D., & Frackowiak, R. S. J. (1995). Statistical parametric maps in functional imaging: A general linear approach. *Human Brain Mapping*, *2*, 189-210.
- Friston, K. J., Josephs, O., Rees, G. & Turner, R. (1998). Nonlinear event-related responses in fMRI. *Magnetic Research in Medicine*, *39*, 41-52.
- Friston, K. J., Penny, W. D., & Glaser, D. E. (2005). Conjunction revisited. *NeuroImage*, *25*, 661-667.
- Friston, K. J., Zarahn, E., Josephs, O., Henson, R. N. A. & Dale, A. M. (1999). Stochastic designs in event-related fMRI. *NeuroImage*, *10*, 607-619.
- Fourier, J. (1822). *The analytical theory of heat*. (translated by Freeman, A. 1878, re-released 2003) New York: Dover Publications.
- Fox, P. T., Raichle, M. E., Mintun, M. A. & Dence, C. (1988). Nonoxidative glucose consumption during focal physiologic neural activity. *Science*, *241*, 462-464.
- Genovese, C. R., Lazar, N. A., & Nichols, T. (2002). Thresholding of statistical maps in functional neuroimaging using the False Discovery Rate. *NeuroImage*, *15*, 870-878.
- Hashemi, R. H., Bradley, W. G., & Lisanti, C. J. (2004). *MRI the basics*. (2nd ed.) Philadelphia: Lippincott Williams & Wilkins.
- Heeger, D. J., Huk, A. C., Geisler, W. S. & Albrecht, D. G. (2000). Spikes versus BOLD: What does neuroimaging tell us about neuronal activity. *Nature Neuroscience*, *3*, 631-633.
- Heeger, D. J. & Ress, D. (2002). What does fMRI tell us about neuronal activity? *Nature Reviews Neuroscience*, *3*, 142-151.
- Holmes, A. P. & Friston, K. J. (1998). Generalisability, random effects and population inference. *NeuroImage*, *7*, S754.
- Horowitz, A. L. (1995). *MRI physics for radiologists*. (3rd ed.) New York: Springer-Verlag.
- Horowitz, B., Friston, K. J., & Taylor, J. G. (2000). Neural modeling and functional brain imaging: an overview. *Neural Networks*, *13*, 829-846.
- Jenkinson, M. (2001). Registration, atlases and cortical flattening. In P. Jezzard, P. M. Matthews, & S. M. Smith (Eds.), *Functional MRI: An introduction to methods* (pp. 271-293). New York: Oxford University Press Inc.
- Jenkinson, M., Bannister, P., Brady, M., & Smith, S. (2002). Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, *17*, 825-841.
- Jenkinson, M. & Smith, S. M. (2001). A global optimisation method for robust affine registration of brain images. *Medical Image Analysis*, *5*, 143-156.
- Jezzard, P. & Clare, S. (1999). Sources of distortion in functional MRI data. *Human Brain Mapping*, *8*, 80-85.
- Jezzard, P. & Clare, S. (2001). Principles of nuclear magnetic resonance and MRI. In P. Jezzard, P. M. Matthews, & S. M. Smith (Eds.), *Functional MRI: an introduction to methods*. (pp. 67-92). New York: Oxford University Press Inc.

- Jones, R. A., Brookes, J. A., & Moonen, C. T. W. (2001). Ultra-fast MRI. In P. Jezzard & S. Clare (Eds.), *Functional MRI: An introduction to methods*. (pp. 93-108). New York: Oxford University Press Inc.
- Josephs, O. & Henson, R. N. A. (1999). Event-related functional magnetic resonance imaging: Modelling, inference and optimization. *Philosophical Transactions of the Royal Society of London Series B-Biological Sciences*, 354, 1215-1228.
- Kiebel, S. J. & Holmes, A. P. (2003). The general linear model. In R.S.J. Frackowiak, K. J. Friston, C. Frith, R. Dolan, C. J. Price, S. Zeki, J. Ashburner, & W. D. Penny (Eds.), *Human brain function* (2nd ed., San Diego: Academic Press.
- Lange, N. (2000). Statistical procedures for functional MRI. In C.T.W. Moonen & P. A. Bandettini (Eds.), *Functional MRI* (pp. 301-335). Heidelberg: Springer-Verlag Berlin.
- Lauritzen, M. & Gold, M. (2003). Brain function and neurophysiological correlates of signals used in functional neuroimaging. *The Journal of Neuroscience*, 23, 3972-3980.
- Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., & Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*, 412, 150-157.
- Logothetis, N. K. & Pfeuffer, J. (2004a). On the nature of the BOLD fMRI contrast mechanism. *Magnetic Resonance Imaging*, 22, 1517-1531.
- Logothetis, N. K. & Wandell, B. A. (2004b). Interpreting the BOLD signal. *Annual Review of Physiology*, 66, 735-769.
- Mansfield, P. (1977). Multiplanar image formation using NMR spin echoes. *The Journal of Physical Chemistry*, 10, L55-L58
- Menon, R. S. (2001). Imaging function in the working brain with fMRI. *Current Opinion in Neurobiology*, 11, 630-636.
- Menon, R. S. & Goodyear, B. G. (2001). Spatial and temporal resolution in fMRI. In P. Jezzard, P. M. Matthews, & S. M. Smith (Eds.), *Functional MRI: An Introduction to methods* (pp. 145-158). New York: Oxford University Press Inc.
- Menon, R. S. & Kim, S.-G. (1999). Spatial and temporal limits in cognitive neuroimaging with fMRI. *Trends in Cognitive Sciences*, 3, 207-216.
- Morris, P. G. (1986). *Nuclear magnetic resonance imaging in medicine and biology*. Oxford: Clarendon Press.
- Nichols, T., Brett, M., Andersson, J., Wager, T., & Poline, J.-P. (2005). Valid conjunction inference with the minimum statistic. *NeuroImage*, 25, 653-660.
- Nichols, T. & Hayasaka, S. (2003). Controlling the familywise error rate in functional neuroimaging: A comparative review. *Statistical Methods in Medical Research*, 12, 419-446.
- Ogawa, S., Lee, T. M., Kay, A. R., & Tank, D. W. (1990a). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences of the United States of America*, 87, 9868-9872.
- Ogawa, S., Lee, T. M., Nayak, A. S. & Glynn, P. (1990b). Oxygenation-sensitive contrast in magnetic resonance imaging of rodent brain at high magnetic fields. *Magnetic Resonance in Medicine*, 14, 68-78.
- Ogawa, S., Tank, D. W., Menon, R., Ellermann, J. M., Kim, S.-G., Merkle, H. & Ugurbil, K. (1992). Intrinsic signal changes accompanying sensory stimulation: Functional brain mapping with magnetic resonance imaging. *Proceedings of the National Academy of Sciences of the United States of America*, 89, 5951-5955.
- Pfeuffer, J., McCullough, J. C., van de Moortele, P.-F., Ugurbil, K., & Hu, X. (2003). Spatial dependence of the nonlinear BOLD response at short stimulus duration. *NeuroImage*, 18, 990-1000.

- Smith, S. M. (2001). Preparing fMRI data for statistical analysis. In P. Jezzard, P. M. Matthews, & S. M. Smith (Eds.), *Functional MRI: An introduction to methods* (pp. 229-241). New York: Oxford University Press Inc.
- Talairach, J. & Tournoux, P. (1988). *Co-planar stereotaxic atlas of the human brain: 3-dimensional proportional system: An approach to cerebral imaging*. New York: Thieme Medical Publishers.
- Ugurbil, K., Toth, L., & Kim, D.-S. (2003). How accurate is magnetic resonance imaging of brain function? *Trends in Neurosciences*, *26*, 108-114.
- Vanzetta, I. & Grinvald, A. (1999). Increased cortical oxidative metabolism due to sensory stimulation: Implications for functional brain imaging. *Science*, *286*, 1555-1558.
- Wager, T. D., Vazquez, A., Hernandez, L., & Noll, D. C. (2005). Accounting for nonlinear BOLD effects in fMRI: Parameter estimates and a model for prediction in rapid event-related studies. *NeuroImage*, *25*, 206-218.
- Woolrich, M. W., Behrens, T. E. J., Beckmann, C. F., Jenkinson, M., & Smith, S. M. (2004). Multi-level linear modelling for fMRI group analysis using Bayesian inference. *NeuroImage*, *21*, 1732-1747.
- Woolrich, M. W., Ripley, B. D., Brady, J. M., & Smith, S. M. (2001). Temporal autocorrelation in univariate linear modelling of fMRI data. *NeuroImage*, *14*, 1370-1386.
- Worsley, K. J. (2001). Statistical analysis of activation images. In P. Jezzard, P. M. Matthews, & S. M. Smith (Eds.), *Functional MRI: An introduction to methods* (pp. 251-270). New York: Oxford University Press Inc.
- Worsley, K. J., Evans, A. C., Marrett, S., & Neelin, P. (1992). A three-dimensional statistical analysis for CBF activation studies in human brain. *Journal of Cerebral Blood Flow and Metabolism*, *12*, 900-912.