An introduction to functional MRI

Bianca de Haan and Chris Rorden

Introduction1					
1 The basics of MRI					
	1.1	The brain in a magnetic field	1		
1.2		Application of the radiofrequency pulse	3		
	1.3	Relaxation	4		
1.4 When it all comes together		When it all comes together	4		
	1.5	T2* and the spin-echo pulse cycle	5		
2	Fre	om neural activation to the fMRI signal	8		
	2.1	The T2* signal	8		
	2.2	Optimizing the BOLD signal	11		
3 fMRI task desi		RI task design	.13		
	3.1	Block designs	.14		
	3.2	Event related designs	15		
	3.3	Summary of fMRI Design	18		
4	fM	RI analysis	.18		
	4.1	Slice timing correction	.19		
	4.2	Spatial realignment	21		
	4.3	Spatial normalization	.24		
	4.4	Spatial coregistration	.26		
	4.5	Spatial smoothing	.28		
	4.6	Statistical analysis	.29		
	4.7	Choosing a statistical threshold	31		
5	Lit	erature	.38		

Introduction

Functional Magnetic Resonance Imaging (fMRI) has rapidly become a widely used tool for measuring brain function. fMRI reveals which parts of the brain are active in certain tasks with a spatial resolution of 2-5 millimetres, which is superior to any of the other noninvasive 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 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 et al., 2000). Given this superior spatial resolution but poor temporal resolution, fMRI is best seen as a complementary tool for understanding brain function.

To start with, some of the principles underlying MRI will be considered. How is the MRI image obtained? This will be explained in the first section of this report. In cognitive neuroscience, MRI is used to indirectly infer the functional activity of the brain, in which case it is referred to as functional Magnetic Resonance Imaging (fMRI). How neural activation leads to an fMRI signal will be discussed in the second section of this report. In the third section of this report, attention will be focused towards issues relating to the different kinds of experimental design that can be used in fMRI experiments. The fourth section of this report will be discussed in the second section.

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 report. However, some basic knowledge of the principles is required to fully understand the subsequent sections.

There are three basic steps involved in measuring the MRI signal. The first step is to place the brain in a magnetic field. The second step is the application of a brief radiofrequency (RF) pulse and the third step is measuring the relaxation (Horowitz, 1995). Each of these steps will in turn be considered in the following paragraphs.

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. This will cause the atomic nuclei to align with the magnetic field. This occurs to all nuclei that are

electrically charged and spin around their axis. Of the many types of nuclei in the brain, it is the hydrogen nucleus that is most commonly measured in MRI (Horowitz, 1995). This is because hydrogen nuclei are abundant in the human brain and give a strong MRI signal. 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. Since each hydrogen nucleus produces a magnetic field, which can be represented 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) (Horowitz, 1995; Jezzard & Clare, 2001).

Before the brain is placed in a magnetic field, the MDM's of each hydrogen nucleus points in a random direction: the nuclei are not aligned. When the brain is placed in a magnetic field, two things happen simultaneously (Horowitz, 1995; Jezzard & Clare, 2001). Firstly, the MDM's of many 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 this magnetic field. The stronger the magnetic field, the higher the percentage of the MDM's that align themselves to the magnetic field (Horowitz, 1995; Jezzard & Clare, 2001). Secondly, when the brain is placed in a magnetic field the MDM's of the hydrogen nuclei start to precess (see Fig. 1.1). 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 sodium nuclei in a certain magnetic field. 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 (64000000 revolutions per second) and in a magnetic field of 3 Tesla the frequency of precession will be 128 MHz (Horowitz, 1995; Jezzard & Clare, 2001).



Fig. 1.1: Precession of the MDM; the 'tail' is stationary while the 'top' moves around in a circular motion.

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° RF-pulse (Horowitz, 1995; Jezzard & Clare, 2001). The ultimate goal of this 90° 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° RF-pulse then basically 'tips' the MDM's in the x-y plane (see fig. 1.2). 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, it is possible to selectively 'tip' the MDM's of the hydrogen nuclei (Horowitz, 1995; Jezzard & Clare, 2001).



Fig 1.2: '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° RF-pulse is terminated and the MDM's return to their original orientation. This returning to the original orientation is known as relaxation (Horowitz, 1995; Jezzard & Clare, 2001).

1.3 Relaxation

After the 90° RF-pulse is terminated, the MDM's of the hydrogen nuclei will return from their 'tipped' state to their original lower energy state of being aligned in the direction of the magnetic field (Horowitz, 1995; Jezzard & Clare, 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 is the signal that is measured during MRI (Horowitz, 1995; Jezzard & Clare, 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 (Horowitz, 1995; Jezzard & Clare, 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 (Horowitz, 1995; Jezzard & Clare, 2001).

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 over a vast amount of times in a typical MRI experiment. Different tissues in the brain have different T1 and T2 relaxation rates (Horowitz, 1995; Jezzard & Clare, 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 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 (Horowitz, 1995; Jezzard & Clare, 2001). At first, it sounds a bit counterintuitive that not the absolute difference in amplitude is maximized but instead the relative difference between the amplitudes is maximized. An analogy might clarify things a bit. Imagine a person

A who makes 20,000 dollars a year and a person B who makes 22,000 dollars a year. The absolute difference is 2,000 dollars but the relative difference is quite small (10%). As a consequence, person B doesn't really feel a lot richer than person A because the contrast between the two incomes is not very big. If now, on the other hand, person A makes 3,000 dollars a year and person B makes 4,700 dollars a year, even though the absolute difference is now smaller (1,700 dollars), the relative difference is a lot bigger. In the second case, even though the absolute difference is smaller, the relative difference is a lot larger and person B ends up feeling a lot better off than person A because the contrast between the two incomes is now a lot bigger. To return to contrasts between different tissues in the brain, it can now be seen that to maximize contrast between different tissues the MRI signal must be measured at a time when the relative difference in amplitudes of the MDM's is maximized and not the absolute difference.

When the MRI signal is measured at a point when the ratio of 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 ratio of 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 (Horowitz, 1995; Jezzard & Clare, 2001). By changing certain scanner parameters either a T1 weighted signal or a T2 weighted signal can be acquired (Horowitz, 1995; Jezzard & Clare, 2001). When the time from RF-pulse to measurement of the signal (TE) is kept short, while at the same time 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. T1 weighted scans are also known as anatomical scans, because they particularly show good contrast between grey and white matter. On the other hand, when the TE is long while at the same time the TR is also long, the difference in T2 for the different tissues is maximized and the acquired scan. T2 weighted scans are also known as pathological scans, because lesions appear very bright (Horowitz, 1995; Jezzard & Clare, 2001).

1.5 T2* and the spin-echo pulse cycle

In the previous sections, it was implied that the decay of the MDM's in the x-y plane after termination of the 90° RF-pulse equals the T2 relaxation signal. This is, however, a simplification and to really understand the use of MRI we will need to explain T2 in more detail. 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° RF-pulse. The decay of the MDM's in the x-y plane is more

accurately described as T2* decay (Horowitz, 1995; Jezzard & Clare, 2001). So, the question remains; what is T2 decay?

The reason why the MDM's decay in the x-y plane (the T2* signal) is essentially due to dephasing (Horowitz, 1995; Jezzard & Clare, 2001). Remember that the MDM's of hydrogen nuclei in a magnetic field of a certain field strength all precess at the same frequency. However, before the application of the 90° RF-pulse, they are not precessing in the same phase. To clarify this, an analogy with a clock can be helpful. Imagine a number of clocks that are not synchronized. Even though the hands of the different clocks rotate at the same frequency (an hour lasts equally long for all clocks), at a given moment the times the 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 will each be at a different position in their cycle at a given point in time i.e. they are precessing in a different phase (see fig. 1.3) (Horowitz, 1995; Jezzard & Clare, 2001).



Fig 1.3: 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.

At the moment when the 90° RF-pulse is applied, the MDM's are 'tipped' in the x-y plane and are forced to precess in phase (as the brief RF-pulse 'tips' all the MDM's simultaneously). Now the MDM signals are additive and therefore result in a strong signal in the x-y plane. After termination of the 90° RF-pulse, however, the precession of the MDM's will gradually dephase. The MDM signals are now no longer additive, but cancel each other out and the signal decays. This is the T2* decay (Horowitz, 1995; Jezzard & Clare, 2001).

There are two reasons why the MDM's dephase after termination of the 90° RF-pulse and the signal decays (Horowitz, 1995; Jezzard & Clare, 2001). The first one is magnetic field inhomogeneity. The strength of the magnetic field is not uniform and since the frequency of precession of the MDM's depends on the strength of the magnetic field it follows that different MDM's will precess at a different frequency and therefore this precession will get out of phase. The second reason why the MDM's dephase is because of spin-spin interaction. Different hydrogen nuclei are surrounded by different other nuclei. These other nuclei affect the frequency of precession of the MDM of the hydrogen nuclei. The frequency of precession of each MDM's will be differently affected by the surrounding nuclei. This, again, results in different frequencies of precession for different MDM's and hence dephasing occurs (Horowitz, 1995; Jezzard & Clare, 2001).

To summarize, initially the precession of the MDM's is dephased. When the 90° RFpulse 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. This is T2* decay.

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 (Horowitz, 1995; Jezzard & Clare, 2001). Directly after application of the 90° RF-pulse the precession of the MDM's is in phase. This means that at a given time, all the MDM's will be at the same point in their cycle of precession. After termination of the 90° RF-pulse the precession of the MDM's will slowly dephase. Basically, at a given time, different MDM's will not longer be at the same point in their cycle of precession. If an RF-pulse is now applied from the opposite direction (180°) as the direction from which the original RF-pulse was applied (90°), the direction of rotation of the precession of the MDM's is reversed. After the same amount of time has elapsed following the 180° RF-pulse as the amount of time between the 90° RF-pulse and the 180° RF-pulse, the MDM's will be in phase again (Horowitz, 1995; Jezzard & Clare, 2001).

An analogy might, again, 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° 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° 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° 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° RF-pulse and measurement of the signal must be the same as the time between the 90° RF-pulse and the 180° RF-pulse. A measurement cycle where one 90° RF-pulse is followed by one or more 180° RF-pulses with a measurement after each 180° 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 (Horowitz, 1995; Jezzard & Clare, 2001).

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), which is basically a technique that allows for fast measurement of the signal. The technique most commonly used in fMRI is the so-called BOLD (blood oxygenation level-dependent contrast) technique (Horwitz et al., 2000; Uğurbil, Toth & Kim, 2003). 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; Horwitz et al., 2000; Matthews, 2001; Uğurbil, Toth & Kim, 2003). In this section the relationship between neuronal activity and haemodynamics will be explained.

2.1 The T2* signal

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 speed 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 fig. 2.1). Therefore, the larger the inhomogeneity of the magnetic field, the smaller the fMRI signal at the time of measurement (Heeger & Ress, 2002; Horowitz, 1995; Matthews, 2001).



Fig. 2.1: 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 through a process known as 'shimming' 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. 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). Deoxyhaemoglobin (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 (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 (Heeger & Ress, 2002; Matthews, 2001).

The function of the fMRI signal against time in response to a temporary increase in neuronal activity is known as the Haemodynamic Response Function (HRF). The HRF goes through three stages (see fig. 2.2) (Heeger & Ress, 2002).



Fig 2.2: Time course of the HRF in response to a short-lasting increase in neuronal activity at time = 0.

The fMRI signal initially decreases, because the active neurons use oxygen thereby increasing the relative level of deoxyhaemoglobin in the blood. This decrease, however, is tiny and is not always found (Uğurbil, Toth & Kim, 2003). Following this initial decrease, there is a large increase in the fMRI signal which reaches its maximum after approximately 6 seconds. This increase is due to a massive oversupply of oxygen-rich blood. There are two main hypotheses regarding the reason for this increase in blood flow. The first theory states that the increase in blood flow compensates for the oxygen being used by the active neurons. However, the supply in oxygen by the increase in blood flow is much larger than the amount of oxygen used by the active neurons. Instead, the increase in blood flow is proportional to the amount of glucose being used by the active neurons. Therefore, the second theory states that the increase in blood flow instead compensates for the amount of glucose being used by the neurons and not the amount of oxygen (Heeger & Ress, 2002). In any case, 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 fMRI signal. Finally, the last stage of the HRF is a slow return to the normal level of deoxyhaemoglobin and a decay of the fMRI signal until it has reached its original baseline level after an initial undershoot after approximately 24 seconds (Heeger & Ress, 2002).

It is important to note, first of all, that the BOLD fMRI signal is an indirect measure of the underlying neuronal activity and therefore strongly relies on the assumption that neuronal activity and haemodynamics are indeed linked. Secondly, the fMRI signal reflects the sum of the activity of a large group of neurons (Heeger & Ress, 2002). This first of all means that an increase of the fMRI signal can both be caused by a large increase in activation of a small group of neurons and by a small increase in activation of a large group of neurons (Heeger & Ress, 2002). Inherent, however, is also the deeper assumption that the neurons responsible for the same function will be grouped together in the brain. Thirdly, the BOLD fMRI signal is sensitive to contaminations of large veins in the brain (Lai, Glover & Haacke, 2000; Menon & Goodyear, 2001; Uğurbil, Toth & Kim, 2003). The relative decrease in deoxyhaemoglobin is larger in large veins than in small veins. This means that the maximum BOLD fMRI signal is often obtained in the large veins that can be a few millimetres away from the site of neural activation. Fourthly and related to the third point, the increase in blood flow (and the associated decrease in deoxyhaemoglobin) is not very specific to the area of neuronal activation (Lai, Glover & Haacke, 2000; Matthews, 2001; Menon & Goodyear, 2001; Uğurbil, Toth & Kim, 2003). Usually, veins supply blood to a larger area of the brain than just the area of neuronal activation. This leads to spatial blurring, where the area that shows an increase in fMRI signal is larger than the area of neuronal activation. Finally, it is also important to note that the haemodynamic response is inherently much slower than the underlying neuronal activity. Basically, the fMRI signal can be seen as a smoothed function of the underlying neuronal activity (Horwitz et al., 2000). 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 (Heeger & Ress, 2002). Note that this poor temporal resolution is inherent in the BOLD fMRI technique. While the spatial resolution might be improved by further technical innovations, the temporal resolution can not be improved because of this inherent slowness of the haemodynamic response. The only way to improve the temporal resolution would be by measuring a more direct consequence of neuronal activity instead of the indirect consequence that is the haemodynamic response.

2.2 Optimizing the BOLD signal

As noted in the previous section, the magnitude of the BOLD effects measured during fMRI is very small. As Figure 2.2 depicts, a visual flash may result in a 2% signal change in the occipital cortex (the first cortical stage of visual processing). This low signal-to-noise is a

severe limit on our statistical power (our ability to find real effects). Broadly speaking, there are three ways we can maximize statistical power. First, we can optimise the behavioural task used (e.g. using a block design), as discussed in the next section. Second, we can run very long fMRI sessions. Unfortunately, this is expensive and often results in participants feeling fatigue. A final method to maximize fMRI signal is to optimise the hardware, which we will discuss in this section. In practice, combining all three techniques is often required to detect subtle differences in the BOLD signal.

In particular, we identify four common practices that influence the BOLD signal. First, we discuss echo planar imaging and spiral imaging to collect images of the brain quickly. Next, we note that the strength of the T2* signal is greatly dependent on the echo time (TE). Third, we discuss the relationship between magnet field strength and signal. Fourth, we discuss how the matrix size, voxel size and receiver coil all critically influence the measured signal.

Optimal BOLD effects will be found when the echo time (TE) matches the T2* of blood. In theory, this means that the BOLD effect will be maximum with a TE of around 65ms at a field strength of 1.5T and around 40ms at 3.0T (Hennig et al., 2003). In practice, most labs use slightly shorter echo times. This probably reflects a trade-off: by using a shorter TE, the BOLD effect in each image is reduced but more images are collected in a fixed amount of time. Therefore, while the signal in each sample may be reduced, the larger number of samples results in improved statistical power. Furthermore, the shorter TE can result in slightly reduced image distortion. In our own research, we use echo planar imaging with a TE of 60ms at 1.5T and a TE of 30ms at 3.0T.

As noted earlier, the numbers of hydrogen nuclei that align with the magnetic field depend directly on the field strength of the magnet (and the temperature of the sample, about 37 Celcius for the human brain). In a 1.5 Tesla, about 4.5 nuclei per million are aligned with the field than against the field, while at 3.0 Tesla the number of rises to about 9 nuclei per million. Therefore, at higher field strength, there is theoretically more signal to measure. Also, as noted in the previous paragraph, the T2* of blood reduces at higher field strengths, allowing a shorter time to echo time. Therefore, we can collect images of the brain more rapidly at higher field strengths. Despite these advantages, higher field magnets also have limitations. First of all, increasing field strength drives up the installation and service cost for a MRI system, and higher field systems tend to be acoustically much louder (due to more powerful gradients required). Second, at higher field strengths inhomogeneity effects become more severe. Images from 3.0 Tesla magnets typically show much greater inhomogeity

artefacts (with some regions of the brain appearing brighter than other regions), and exhibit regions of signal dropout than images from 1.5 Tesla systems. Clearly, studies in single individuals do reveal stronger fMRI signals with increasing field strength (Hennig et al., 2003). However, images from higher field scanners may not necessarily yield better signal in group studies (as image distortion results in poor normalization of images, as discussed later) and may result in an inability to measure some regions of the brain (such as the medial temporal lobe and ventral frontal lobe).

With MRI scanning, one can trade-off image resolution for signal quality. By increasing the size of each voxel, we can effectively measuring more hydrogen nuclei. For example, consider axial slices where we have a 64x64 matrix (i.e. the image has 64 rows of pixels and 64 columns) with a field of view of 192x192mm (pixels are separated by 3mm in plane). If the distance between slices is 3mm, each voxel will measure 27mm³, and it will require around 36 axial slices to cover the entire cortex and most of the cerebellum. On the other hand, if we collect images that are 5mm thick, each voxel will measure 45 mm³, and we only need to collect around 22 slices. By collecting fewer slices, we can image the brain more rapidly, and get more samples of a given brain area. Another common technique to increase the speed of acquisition is to collect EPI data as axial or sagittal slices instead of coronal slices (as the brain is longest in the axial-posterior direction). Another issue is interference between slices: unless a gap is added between EPI slices, there will be some reduction in image quality due to the RF pulse saturating neighbouring slices along with the target slice. One simple technique to reduce this interference is to collect the data in an interleaved fashion: all the odd slices of an image are collected first (e.g. 1,3,5..35) and then all the even slices (2,4,...36). This minimises interference between slices and allows you to sample more nuclei than if a slice gap is employed.

3 fMRI task design

The fMRI signal is very noisy. The signal of interest is typically only 1 to 10 percent of the overall signal. 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. In this section, two different types of experimental design used in fMRI research, namely block designs and event related designs will be described and their strengths and weaknesses regarding their ability to answer certain research questions will be explored.

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; Dale & Buckner, 1997; Donaldson & Buckner, 2001). This is the most commonly used experimental design. 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 & D'Esposito, 2000; Donaldson & Buckner, 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 fig. 3.1). Therefore, block designs offer considerable statistical power (Aguirre & D'Esposito, 2000).



Fig.3.1: The HRF function for a block design. The x-axis displays the number of volumes. One volume is collected every 2 seconds. The vertical axis displays the relative signal change in percentages, while the horizontal axis shows time. The tick marks along the bottom of the figure illustrate when fMRI is sampled. The red diamonds show the stimulus onsets. In this example, 5 stimuli spaced 1700ms apart are presented in each task block. Therefore, each task block lasts for 8.5 seconds alternated with 11.5 seconds rest allowing the HRF to return to baseline levels. Note the large signal change (>6%), as the signal from the successive stimuli in the block is fairly additive.

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 (Aguirre & D'Esposito, 2000). Another problem with block designs is related to the use of the subtraction paradigm (Aguirre & D'Esposito, 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 block designs have a low power to draw inferences (Aguirre & D'Esposito, 2000).

3.2 Event related designs

The second type of experimental design is the so-called event related design (Aguirre & D'Esposito, 2000; Dale, 1999; Dale & Buckner, 1997; Donaldson & Buckner, 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 (Dale & Buckner, 1997). 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 fig. 3.2) (Aguirre & D'Esposito, 2000; Donaldson & Buckner, 2001).



Fig.3.1: The HRF function for an event related design. The x-axis displays the number of volumes. One volume is collected every 2 seconds. The y-axis displays the relative signal change in percentages. The red diamonds show the stimulus onsets. In this example, 1 stimulus is presented every 20 seconds. Especially note the small signal change (less than 2%).

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 & Buckner, 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 (Dale & Buckner, 1997; Donaldson & Buckner, 2001). The third and most important feature is the discovery that the BOLD fMRI shows a roughly linear response to repeated presentations of stimuli. When two or more stimuli are presented in rapid succession, the BOLD fMRI response increases roughly linearly. This is especially important in event related designs when stimuli are closely spaced together. When stimuli are closely spaced together there is an overlap in the BOLD fMRI responses. The discovery that the BOLD fMRI response increases 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, which enables stimuli to be presented in rapid succession in an event related design (Aguirre & D'Esposito, 2000; Dale & Buckner, 1997; Donaldson & Buckner, 2001).

A big issue when maximizing the statistical power in an event related design is the question of the optimal spacing in time between the onsets of successive stimulus presentations. It has been argued that to maximize statistical power the spacing between the onsets of successive stimulus presentations, also known as InterStimulus Interval (ISI) should be at least 15 seconds and that the statistical power decreases as the ISI decreases (see Dale, 1999). However, recently it has been shown that the statistical power can increase with decreasing ISI's (Dale, 1999). The key is the distinction between fixed ISI's and variable ISI's. 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. This variable ISI is also known as a latency jitter (Aguirre & D'Esposito, 2000; Dale, 1999; Dale & Buckner, 1997; Donaldson & Buckner, 2001).



Fig. 3.2: The HRF function for an event related design. The x-axis displays the number of volumes. One volume is collected every 2 seconds. The y-axis displays the relative signal change in percentages. The red diamonds show the stimulus onsets. **Top panel** : HRF for stimulus presentation with a fixed ISI of 4 seconds. Note that there is very little predictable variability in the response. **B**: HRF for stimulus presentation with a variable ISI and a mean ISI of 4 seconds. Note the larger amount of predictable variability than in the fixed ISI condition yielding much improved statistical power.

When stimuli are presented rapidly with fixed ISI of a short duration, there is almost no variability in the HRF (see fig. 3.2.A). Therefore, statistical power is low. If, however, a latency jitter is introduced, the variability in the signal increases (see fig. 3.2 B) and therefore statistical power increases. The crucial point is randomizing the duration of the ISI according to an exponential distribution (Dale, 1999). This results in many brief ISI's which 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 (at very short ISI's this relation breaks down a bit) (Dale, 1999). It is important to note here however, that even though the statistical power of event related designs can be improved by above mentioned procedure, 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 & D'Esposito, 2000; Donaldson & Buckner, 2001). The importance of

randomization is that potential confounds like habituation, anticipation and strategy effects are cancelled out. 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 & Buckner, 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 the stimulus associated with the wrong responses post hoc. Finally, some experimental questions can not be answered using a block design (Donaldson & Buckner, 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 & Buckner, 2001).

3.3 Summary of fMRI Design

To summarize, block designs have high statistical power, but this is combined with a low power to draw inferences. On the other hand, event related designs have low statistical power, but they allow for more freedom in experimental design leading to a higher power to draw inferences. The obvious way to combine the strengths and weaknesses of both block and event related designs is to use block designs as an exploratory tool. For example, block designs can be used in a brief 'localizer' session that identifies the crucial brain regions for each individual. During subsequent event related designs, the hemodynamic responses within these regions can be identified.

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. This means that the HRF is sampled at many different points in time. Every volume consists of a number of 3D slices (or slabs). These slices are typically a few millimetres thick. Each slice consists of a number of voxels (a 3D data points) and each voxel represents a data point on the HRF. In EPI, one 3D slice is collected at a time. Alongside the EPI volumes that display areas of the brain that are active in a given task (functional activity) often a high resolution T1 or T2 anatomical volume is also obtained. This

anatomical volume is used for spatial localization of the areas of the brain that are active in a given task.

After obtaining the EPI data from the scanner, a number of stages of analysis must be performed on the dataset to finally end up with a map that displays the regions with significant functional activity. These stages can roughly be subdivided in a stage where the data is temporally adjusted, stages where the data is spatially adjusted and a stage of statistical analysis. The stage where the data is temporally adjusted is known as slice timing correction. The stages where the data is spatially adjusted are spatial realignment, spatial normalization and spatial smoothing. The last stage is the stage of statistical analysis (Frackowiak et al., 1997). Each of these stages and their consequences for data interpretation will be considered separately in the following paragraphs.

4.1 Slice timing correction

In EPI every 3D volume of the entire brain consists of a number of 3D slices and these slices are each collected at a slightly different time. However, during 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 HRF starts at an earlier time for slices that are acquired later in time than for slices that are acquired earlier in time (see fig. 4.1).



Fig. 4.1: The HRF 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 HRF 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, 2001b; Veltman & Hutton, 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 of 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 by method of interpolation. 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 they would have

had, had they been obtained at the same point in time as the voxels in the reference slice. In other words, the same change in the HRF now starts at the same time for each voxel in each slice in a volume (Smith, 2001b; Veltman & Hutton, 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 (Veltman & Hutton, 2001).

4.2 Spatial realignment

During an fMRI experiment, hundreds of EPI 3D 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 to represent the same location in the brain throughout time (Brammer, 2001; Smith, 2001b). 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 & Friston, 1997; Ashburner & Friston, 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 & Friston, 1997; Ashburner & Friston, 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 image coregistration (Ashburner &

Friston, 1997; Ashburner & Friston, 2000; Frackowiak et al., 1997). One brain volume (usually the first image in the time series) is taken as the reference volume. All the other volumes in the same time series are repositioned until they are in the same position as the reference volume. This repositioning is done using 6 parameters; x-translation, y-translation, z-translation, rotation around the x-axis, rotation around the y-axis and rotation around the z-axis (these transforms are discussed in more detail in the next section). Only the position of the brain is changed and not the size or shape. This repositioning treats the head as a rigid object and is therefore also known as a rigid body transformation.

For this realignment SPM2 uses a minimization algorithm for the least mean square difference between the volumes in the time series and the reference volume. This means that for each volume the squared difference between the volume in the time series and the reference volume is minimized (see Fig 4.2). 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 & Friston, 1997; Ashburner & Friston, 2000; Frackowiak et al., 1997).



Fig 4.2: 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 due to head movement. C: The difference between A and B. D: The squared difference(variance) 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 of realignment. The first limitation is that, when the head movements are too large, the minimization algorithm might get stuck in a local minimum (Ashburner & Friston, 2000; Brammer, 2001). Usually, the data of subjects with movements that are too large will have to be removed. The second limitation is that the brain is not rigid. Heart-beat and respiration cause variations in shape and size of the brain.

Usually, this source of movement is ignored (P. Liddle, personal communication, fall, 2002). The third limitation is that even perfect realignment will not remove all movement related variance (Ashburner & Friston, 2000; Brammer, 2001; Frackowiak et al., 1997; Veltman & Hutton, 2001). The problem is that, in case of a movement, the image not only moves, but also fundamentally changes. 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. One 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 & Friston, 2000; Brammer, 2001; Frackowiak et al., 1997; Veltman & Hutton, 2001). The fourth limitation of realignment is that a movement also changes the overall magnetic field. Remember from the second section in this report that inhomogeneities in the magnetic field result in a decrease in the fMRI signal. Through a process known as 'shimming' every attempt is made to make the overall magnetic field as uniform as possible. However, when the head moves the overall magnetic field changes. Shimming is usually only done at the beginning of a scanning session and head movements therefore create new inhomogeneities in the overall magnetic field, leading to the overall magnetic field no longer being uniform and consequently to loss of signal in some areas of the brain. The fifth limitation of realignment is that the interpolation method used to estimate the new value of the fMRI signal of a voxel after the realignment is not perfect. This leads to so-called interpolation errors in the signal (Ashburner & Friston, 2000; Veltman & Hutton, 2001).

Realignment computes a spatial transform to generate a stabilized image of the brain. As mentioned, head movements in the scanner result in changes that a simple rigid-body transformation can not correct. For example, the head movements alter the shim of the magnet, so head motion causes changes to the intensity of some regions of the image. Andersson and colleagues have added an EPI unwarping function to SPM2 that reduces this movement-by-inhomogeneity interaction. This software compensates for the changes in image brightness that result from head movements. This unwarping stage can be selected at the same time as the realignment. This stage is computationally intensive, and typically requires more processing time than all other spatial preprocessing steps combined.

Clearly, motion creates serious consequences for fMRI data analysis. As we have discussed, a number of clever techniques have been devised to minimise motion artefacts. However, these techniques will not completely eliminate motion artefacts, in particular if the head movements are large or sudden. Therefore, if possible, head movement should be minimized, for example, by using a head constraint. Furthermore, tasks should be designed to minimize the risk of task correlated movement.

4.3 Spatial normalization

During an fMRI experiment, usually, data is collected for several subjects. However, each individual's brain differs in orientation, size and shape relative to other members in the group. Usually, the orientation, size and shape of the brains of individual subjects are changed to match the orientation, size and shape of a standard brain (Ashburner & Friston, 1997; Ashburner & Friston, 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 represents the same anatomical location, comparisons between different subjects are possible (Ashburner & Friston, 1997; Ashburner & Friston, 2000; Frackowiak et al., 1997). 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 (Jenkinson, 2001). This matching of individual brains to a standard brain is known as spatial normalization (Ashburner & Friston, 1997; Ashburner & Friston, 2000; Frackowiak et al., 1997).

The orientation, size and shape of the brains of all the subjects are changed to match the orientation, size and shape of a standard brain. This standard brain is also known as a template (Frackowiak et al., 1997). Obviously, choosing the right template is important. Optimally, the template should represent the brain of most subjects. Often, the template is based on an average of many different subjects (Smith, 2001a). However, different scanners have unique characteristics and therefore research groups often create their own template. Even though this improves the match of the individual brain to the template for the research groups, it makes comparisons between research groups more tricky.

The matching of the orientation, size and shape of each individual to the orientation, size and shape of the template is done using 12 linear parameters; x-translation, y-translation, z-translation, roll, pitch, yaw, resizing (growing or shrinking) in three dimensions and three

shear (skewing) deformations. Each of these linear parameters changes the entire brain in the same way and these are also known as affine parameters (see Figure 4.3). Note that any three points that are co-linear (a straight line can be drawn through all three points) will remain co-linear after these linear transforms. In addition, each linear transform influences the entire 3D volume. As a result, linear transforms are fairly robust: small regions of signal dropout rarely have adverse influence on linear transforms as the intact portions of the image constrain the transform. However, for the same reason, linear transforms are limited in how accurately they can match two brains: while the overall size and orientation will be matched, local features will often not be precisely aligned.



Fig 4.3: Linear transforms used for realignment (motion correction), coregistration and the first stages of normalization. The top row shows images prior to the transform, with the image at the bottom showing the result of the transform. Note that each transform can be applied in any of three dimensions (e.g. the image can be translated in the anterior-posterior, superior-inferior or left-right dimension). Note that realignment does not typically implement or require shear transforms.

In sum, linear transforms offer a robust but rudimentary match between an individual's scan and a template image. Therefore, packages like SPM2 and AIR typically use linear transforms for the initial normalization and subsequently apply nonlinear transforms to offer a more accurate normalization. For example, SPM2 applies a set of nonlinear cosine basis functions to improve the normalization. Figure 4.4 shows how nonlinear functions can be used to aid normalization. These enable local changes in the brain, also known as warping,

and potentially offer a better match to the template. A minimization algorithm for 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.



Figure 4.4: Nonlinear transforms. The image on the left cannot be made to accurately match the image on the right using only linear transforms. However, applying to nonlinear basis functions the image can be accurately normalized. These transforms are shown in the middle of the figure: one function crushes information near the vertical center of the image, while the other function crushes information near the horizontal center.

Obviously, some constraints have to be imposed on this minimization procedure, since a large difference in the amount of warping between adjacent areas can lead to unrealistic results. The constraint usually imposed is that the overall warp of the brain has to be smooth. Practically, this means minimizing 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. Usually, this overall matching of normalization is not considered a problem, since precise matching of each individual brain to the template would be unrealistic (Ashburner & Friston, 1997; Ashburner & Friston, 2000; Frackowiak et al., 1997; Jenkinson, 2001).

4.4 Spatial coregistration

In image processing terms, 'coregistration' applies to any method for aligning images. By this definition, both motion correction and normalization are forms of coregistration. However, neuroimagers tend to use the term 'coregistration' to refer to alignment of images from different modalities. For example, matching the low resolution T2* fMRI scan to the high-resolution T1-weighted anatomical scan from the same individual.

In theory, intermodality coregistration is simpler than normalization: there should always be a perfect correspondence between different scans from the same individual (e.g. the same sulcal pattern should exist in both scans). However, note that algorithms implemented by SPM2 during motion correction and normalization assume a similar relationship between tissues in the different scans. For example, in normalization, a T1 scan will be matched to a template image with T1-contrast: with air and water (cerebral spinal fluid) appearing dark and white matter appearing bright. The difference-squared cost function will fail when this relationship is broken. For example, consider a T1 scan and a T2 scan from the same individual that are perfectly aligned with each other: the difference-squared cost function will indicate major difference between these images, as the air is dark in both images, but the water is bright in the T2 but dark in the T1 scan. In this case, the automated algorithms used by SPM's normalization will result in images that do not accurately match each other.

Therefore, inter-modality coregistration can not rely on a simple cost function that simply relies on the raw difference in intensity between the two images. There are two approaches that have been implemented to tackle this problem. The first technique simply normalizes each image to a corresponding template image (e.g. normalizing the T1 scan to a T1 template, and the T2* scan to a T2* template). Next, the images are segmented into grey and white matter maps, and these resulting maps are coregistered to each other using a standard difference-squared cost function. A second technique relies on mutual information theory: this simply relies on the concept that different material will have different intensities within a scan modality (e.g. air will have a consistent brightness in the scan, and this brightness will be different from some other materials, such as white matter). This method is illustrated in Figure 4.5.



Figure 4.5: Intermodality image coregistration using mutual information theory. Consider two MRI scans using different sequences (*Tz*-weighted and *Ty*-weighted images) of the same individual. If these images are accurately aligned (top left) than the resulting image histogram will show very little noise (lower left). On the other hand, if the images are not aligned (top left), than the resulting joint histogram (lower right) will show much more noise. Therefore, we can iteratively apply linear transforms and see if they influence the entropy between the two images, eventually resulting in an accurate alignment.

4.5 Spatial smoothing

An important preprocessing step is to blur the fMRI data prior to statistical analysis. At first glance, the idea of spatially smoothing data seems counterintuitive. After all, blurred data clearly degrades the spatial precision of the image. However, there are a number of important reasons for smoothing. Firstly, smoothing increases the signal to noise ratio in the fMRI signal by removing the noise present in the high spatial frequencies. In fMRI the effects of interest are produced by changes in blood flow. Changes in blood flow are expressed on a low spatial frequency of several millimetres. Noise is usually expressed on a higher spatial frequency. Secondly, smoothing removes small frequency differences, so comparisons across subjects are made easier. Thirdly, smoothing helps to satisfy the requirements for applying Gaussian Field Theory to correct for multiple comparisons in the ensuing statistical analysis. Finally, smoothing makes the data more normally distributed (Smith, 2001b).

Smoothing is performed by convoluting 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, 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 voxelsize. In our own work, we usually collect data with a resolution of around 3x3x3 mm, so we tend to choose an 8 mm FWHM for our smoothing filter. However, it should be noted that the smoothing acts as a spatial filter – a FWHM of 8 mm is tuned to detect clusters of around this size. Therefore, if you have an a priori prediction of the size of the region you hope to measure, you can set your smoothing filter to maximize signals for that size of region (for example, if you wanted to examine the tiny superior colliculi you would want to select a smaller smoothing filter than if you were interested in a larger region such as early visual cortex). The end result of this spatial smoothing is therefore that high spatial frequencies in the signal are filtered out while low spatial frequencies in the signal are emphasized (Smith, 2001b).

4.6 Statistical analysis

After the fMRI volumes have been temporally and spatially adjusted, they are ready for the statistical analysis. In the statistical analysis a test is performed to determine which voxels in the brain are significantly activated by a certain type of stimulus. In the most commonly used type of analysis, the General Linear Model (GLM), the time series of each voxel (HRF over time) is analyzed separately. This separate analysis of the time series of each voxel means that the analysis of fMRI datasets is univariate (Frackowiak et al., 1997; Holmes & Friston, 1997; Lange, 2000; Smith, 2001a; Worsley, 2001).

The first step in the statistical analysis is to create a model to predict the observed data. In the GLM, a model is created that predicts the observed data according to the following formula for linear regression:

$Y = X \cdot \beta + \epsilon$

In this model Y is a matrix which has a column for each voxel and a row for each fMRI scan. This is the matrix with the observed fMRI values. X is the design matrix (see Fig 4.4 for an example) which has a row for every scan and a column for each predictor variable. β is the parameter matrix and ε is the matrix of normally distributed error terms (Frackowiak et al., 1997; Holmes & Friston, 1997; Lange, 2000; Smith, 2001a; Worsley, 2001). In this first step

of the statistical analysis therefore two things have to be done. First the design matrix has to be specified and then the parameter matrix has to be estimated.

In the design matrix the general shape of the HRF over time is modelled. On the most simple level each column in the design matrix reflects a certain stimulus function (often an on/off block curve, stimuli are either present or they are not) smoothed with the HRF. However, other predictor variables can also be modelled in the design matrix for example, covariates like task performance and confounds like slow drifts.

After the design matrix has been modelled, it has to be fit separately to the time course of each voxel. This is the estimation of the parameter matrix. For each column in the design matrix a parameter is obtained that optimizes the fit of the overall waveform modelled in the column of the design matrix to the observed data for each voxel. To clarify, the design matrix predicts the overall shape of the HRF over time and the parameter matrix is estimated to fit this overall shape of the HRF to the specific time course of each voxel (Frackowiak et al., 1997; Holmes & Friston, 1997; Lange, 2000; Smith, 2001a; Worsley, 2001).

After the model has been created to predict the observed data, it is time for the second step in the statistical analysis. In the first step of the statistical analysis a predicted HRF (the design matrix multiplied by the parameter matrix) has been obtained for each column in the design matrix (each predictor variable) and each voxel in the brain. Now a t test is used to determine for each voxel separately whether a specific linear combination of columns in the design matrix is significantly different from zero. Which linear combination of columns in the design matrix is used is determined by a setting up contrasts. If the example is used of the first column in the design matrix containing the predicted HRF for the stimulus function of visual stimuli and the second column in the design matrix containing the predicted HRF for the brain whether the response to the visual stimuli is significantly larger than the response to the auditory stimuli (Frackowiak et al., 1997; Holmes & Friston, 1997; Lange, 2000; Smith, 2001a; Worsley, 2001).

The end result of the statistical analysis is a statistical map that shows which voxels are significantly activated given a certain linear combination of columns in the design matrix (see fig. 4.2). This statistical map can be used as an overlay for an anatomical image to display where in the brain a significant increase in activation occurs in response to a task manipulation.



Fig. 4.6: Example output of the statistical analysis using a widely used program SPM. On the top left the areas of significant activation are shown, with the red arrow indicating the selected voxel. On the bottom left the predicted (solid line) and the observed (dashed line) time course of the activation of this selected voxel is shown, especially not the good fit between the predicted and the observed HRF. On the right, the design matrix and the contrasts used to obtain the areas of significant activation are shown. In this example two stimulus types (or conditions) are compared averaged over nine subjects.

4.7 Choosing a statistical threshold

This section discusses methods for selecting a statistical threshold to apply to your data. In particular, we discuss two related problems. First, how do we decide on a statistical threshold that maximizes our chance of findings real results. Second, fMRI analysis typically results in hundreds of thousands of statistical tests, so we need to control for 'familywise error' (unless we correct for multiple comparisons, we will tend to make many false alarms when we compute many tests).

Statistics measure the probability of an event occurring by chance. Ultimately, we make a decision regarding whether an event represents simply random noise or reflects a real effect. For example, if we are testing to see if a drug treats cancer, we will conclude either that the drug has an effect or that it has no influence.

	Reality				
u		Ho is true (no difference)	Ho is false (conditions influence DV)		
sio	Reject	Wrong	Correct		
$\overline{\mathbf{O}}$	Null Ho	Туре І	Power		
De	Accept	Correct	Wrong Type II		
	NUII HO		Type II		

Figure 4.7: When we conduct statistical tests we decide either that there is no effect (i.e. we accept the null hypothesis Ho is true) or we claim that our conditions influence the dependent variable (we reject the null hypothesis). In reality, the null hypothesis is either true or false, so we can make two kinds of correct judgements: a 'hit' (where we correctly identify a that our conditions have an effect) or a correct rejection (accurately reporting that there is no effect). On the other hand, we can make two types of error, a false alarm (Type I error, where we incorrectly report an effect) and a 'miss' (Type II error, where we fail to detect a real effect). The ability to make hits is often referred to as 'power'.

There are three factors that directly influence our ability to detect effects. First, the size of the treatment effect relative to noise. Second the number of observations (e.g. number of stimuli each person sees in an fMRI study, as well as the number of participants in a group fMRI study). In other words, it is easy to detect big effects, but hard to detect subtle effects. Unfortunately, fMRI signals tend to be noisy, and the effect size is very small: moving ones hand in the scanner might increase the brightness of the motor hand area by around 5%, while the brightness varies due to measurement noise, motion artefacts and cardiac gating. Obviously, more subtle effects are harder to find (e.g. the activation seen when someone sees a bright light will most likely be larger and more intense than for a very dim stimulus). One technique to optimize the small signal is to use block designs (as discuss in section 3.1). Increasing your sample size also has a dramatic effect on your statistical power. For example, collecting lots of trials with each participant and increasing the total number of people tested in an fMRI group study. However, fMRI scans are expensive, and testing people across long sessions can induce fatigue. The final way we can influence statistical power is to adjust our criterion, often referred to as the Alpha level.

An inherent problem with applying statistics to noisy data is that we will routinely make errors. Traditionally, we can make a trade-off: we can prevent making false alarms (claiming effects that in reality are random noise) if we are willing to also miss many real effects. For example, if we set an alpha level of p < 0.1 we will often detect real effects, but we will report false alarms 10% of the time. On the other hand, if we set a more conservative alpha of p < 0.001 our false alarm rate will be low (one in a thousand), but we will often miss real effects.

Figure 4.8 illustrates how adjusting the statistical threshold will influence our conclusions from an analysis. The concept of selecting a statistical threshold is a general problem in statistics. Typically, we select conservative thresholds when we fear the consequences of false alarms (e.g. failing to detect harmful side-effects of a new drug), while liberal thresholds are applied when it is important to identify real effects (e.g. identifying if a drug treats a deadly disease).



Figure 4.8: The left panel shows how adjusting our statistical threshold will influence our conclusions. The top and bottom images show the same data, with the top image showing a liberal statistical threshold (Z>2.0), while the bottom image shows a more conservative threshold (Z>5.0). Note that the liberal test identifies more regions, but has a higher incidence of false alarms (suggesting an area is activated when in fact its activity was simply noise). The right panel illustrates the concept of statistical significance: the L/R and A/P axis show different locations in a slice (from left to right and anterior to posterior) while the color indicates the Z-score at each location. Changing our statistical threshold is analogous to changing the water level, exposing more or fewer peaks.

A very important consideration for neuroimaging data is the tremendous number of statistical tests we compute for each comparison: we typically compute a statistical test for each and every voxel in the image. As a typical brain has a volume of about 2550cc, one will compute almost 100,000 tests if the voxels are 3x3x3mm (or over 300,000 tests if the data is has a resolution of 2x2x2mm). If 100,000 voxels are tested at a probability threshold of 5%

than we should expect five thousand voxels will erroneously appear significant activation. In other words, the chance of 'familywise error' is exceedingly high. These apparent activations by chance are known as false positives. One way to correct for these false positives is by applying a Bonferroni correction. In a Bonferroni correction the probability threshold is divided by the number of tests. The probability threshold for each voxel then becomes 0.05/100,000 = 0.0000005%. While this stringent correction guards against familywise error (we only have a 5% chance of making a false alarm anywhere in the entire brain), we will often fail to detect real effects (i.e. we will have very little statistical power). For example, with an uncorrected 0.05 threshold we will detect any signal that has a z-score greater than 1.64 (this can be found with the Excel formula =NORMINV(0.05,0,1)). On the other hand, the Bonferroni correction will only detect effects with a z-score greater than 5 (=NORMINV(0.000005,0,1)). As a consequence, applying a Bonferroni correction will result in not detecting some 'real' activations.

Considering the massive number of tests performed during neuroimaging, how can we preserve the most statistical power (our ability to find real effects) while minimising our chance of making false alarms. One simple but expensive solution is to run a large number of participants. However, SPM and other brain imaging packages are beginning to implement a number of sophisticated techniques that allow us to choose reasonable statistical thresholds. Solutions include using a minimum cluster size, utilizing Gaussian field theory, false discovery rate correction and region of interest analysis.

In the days of SPM96 (around 1996), one common approach was to set an uncorrected threshold of p < 0.001 but only examine regions where a large number (e.g. 32) contiguous voxels that reached this threshold. The concept is that if a large number of neighbouring observations all show the same behaviour, then the corresponding activity was unlikely to represent noise. One problem with this method is that you will fail to identify real but small modules in the brain. A more fundamental challenge is that neighbouring voxels are not really independent samples. First of all, the smoothing that is applied to the data ensures that nearby voxels will show a similar signal. Furthermore, interpolation during all the stages of spatial preprocessing and the initial MRI reconstruction mean that neighbouring voxels are not really independent samples. Fortunately, clever statisticians have worked out that it is possible to accurately estimate the number of independent tests computed on spatially smoothed fMRI data. Hence, software like SPM2 utilizes Gaussian Random Field theory to accurately estimate the number of independent samples in your data.

The Gaussian Random Field theory assumes that the fMRI signal at each voxel has a normal spatial distribution. When the data has been spatially smoothed with a Gaussian kernel (see section 4.4) this assumption will generally be valid. In this correction the probability threshold is not divided by the total number of voxels, but instead by the number of independent resells (resolution elements) and this correction is therefore much less stringent (Frackowiak et al., 1997; Poline et al., 1997; Lange, 2000; Smith, 2001a; Worsley, 2001). For a description of how the Euler characteristic can be used to set a realistic threshold corrected for familywise error, see Matthew Brett's outstanding web page on the topic (www.mrc-cbu.cam.ac.uk/Imaging/Common/randomfields.shtml). Thankfully, when you choose 'FWE' from SPM2's results section, it will automatically compensate for Familywise Error that takes into account the influence of the resels in your data.

SPM2 introduces a new way to control for multiple comparisons. SPM labels this 'FDR', which is an abbreviation of 'False Discovery Rate'. While Bonferroni correction attempts to correct for the number of false alarms we will make in all samples, FDR attempts to control the ratio of hits to false alarms. In other words, a FDR rate of 5% will result in about 5% of 'activated' voxels being errors. What is remarkable about FDR is that it is sensitive to the amount of signal that is actually in your data. If your dataset is completely noise (with Z-scores appearing in a normal distribution), then FDR will be just as stringent as FWE. On the other hand, if your measurements have a real signal, the resulting Z-scores will not be a normal distribution – the real signal will appear as a shifted distribution. Therefore, FDR adjusts the criterion used based on the amount of signal present in your data. Just like other statistical measures, FDR will be better at detecting large, robust signals. However, FDR does offer a principled way to detect smaller signals while still controlling for false alarms. If you do use FDR, it is important to remember the implications of this correction. For example, consider a study where an FDR of 5% reports 100 activated voxels: in this example around 5 of the 'activated' voxels are in fact false alarms. On the other hand, if you computed a FWE correction at 5% and found 50 voxels active, you can be reasonably confident that all of those are real activations (in fact, there should only be a 1/20 chance of that you have made just one false alarm). FDR is a fairly radical change in the way people think about statistics. Figure 4.9 attempts to illustrate how FDR works.



Figure 4.9: A brief introduction to false discovery rate. Consider a study where 1000 voxels are collected and we compute the Z-score for each voxel. If there is no signal in the data, we can expect the the Z-scores will fall into a nice normal distribution (left panel). Compare this to a situation where 150 of the voxels show a real but small signal (right graph, with blue line showing voxels with signal). Note that due to variability, there is still considerable overlap between the signal and the noise. In both cases, the critical values for the uncorrected t-test p < 0.05 is 1.64 (shown as a *) while the Bonferroni corrected threshold is 3.89 (the ^ symbol). On the other hand, the threshold provided by the false discovery rate is sensitive to the shape of the observed data: in the simulation on the left the value will be 5.89 (! Symbol) while on the right the value will be 2.5. In this simulation, the uncorrected t-test makes 50 errors in the situation where there is no signal, while the Bonferroni and FDR tests have a rate of just 0.05 (only making an average of one error every twenty simulations). For our simulation on the right, the uncorrected test made 148 hits, but 42 false alarms. The FDR test achieved 129 hits with just 6 false alarms, and the Bonferroni correction detected only 27 hits but retained a false alarm rate of 0.05.

A final way to improve your statistical power is to reduce the number of comparisons you make. Instead of examining the entire brain, you can examine just a small region. This approach can be referred to as the 'volume of interest approach'. This is especially useful when you have an a priori reason for believing that a region of the brain is involved with a task. With SPM2, there are two approaches you can either do a 'small volume correction' (or

'SVC' in SPM2-speak), where you only examine the relatively small number of voxels in a given region, or 'region of interest' analysis as implemented by Matthew Brett's MarsBar software. With a region of interest analysis you compute the average activity within the region, creating a single sample and therefore computing just a single test. While individual voxels are noisy, pooling data across a functional module's worth of voxels can give a better estimate of that region's activation. Therefore, region of interest analysis benefits both by reducing the number of comparisons and hopefully making a cleaner estimate of actual activation. Of course, these volume of interest approaches are only as accurate at the regions you provide for analyses. One effective approach is to run a localizer scan to find the approximate location of a module in each individual's brain, and then use this region to look for more subtle effects in future scanning sessions. Another technique is to use a stereotaxic region previously identified. For example, Danish human brain project offers regions of interest for the motor hand area (hendrix.imm.dtu.dk/services/jerne/ninf/voi/hand area.html) and other common regions. Note that you should not use the same data to both functionally define the region of interest and compute a region-of-interest analysis (the region should be defined independently from the data you plan to test).

5 Literature

- Ashburner, J. & Friston, K. J. (1997). Spatial transformation of images. Retrieved February 27, 2003, from http://www.fil.ion.ucl.ac.uk/spm/course/notes97/Ch2.pdf
- Ashburner, J. & Friston, K. J. (2000). Image registration. In Moonen, C. T. W. & Bandettini, P. A. (Eds.), *Functional MRI*. (pp. 285-299). Heidelberg: Springer-Verlag Berlin.
- Aguirre, G. K. & D'Esposito, M. (2000). Experimental design for brain fMRI. In Moonen, C. T. W. & Bandettini, P. A. (Eds.), *Functional MRI*. (pp. 369-380). Heidelberg: Springer-Verlag Berlin.
- Brammer, M. J. (2001). Head motion and its correction. In Jezzard, P., Matthews, P. M. & Smith, S. M. (Eds.), *Functional MRI: an introduction to methods.* (pp. 243-250). New York: Oxford University Press Inc.
- 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.
- Donaldson, D. L & Buckner, R. L. (2001). Effective paradigm design. In Jezzard, P., Matthews, P. M. & Smith, S. M. (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. (1997). *Basic concepts and overview*. Retrieved February 27, 2003, from http://www.fil.ion.ucl.ac.uk/spm/course/notes97/Ch1.pdf
- Heeger, D. J. & Ress, D. (2002). *What does fMRI tell us about neuronal activity?* Nature Reviews Neuroscience, 3, 142-151.
- Henig, J., Speck, O., Koch, M.A., Weiller, C. (2003) Functional magnetic resonance imaging: A review of methodological and clinical applications. Journal of Magentic Resonance Imaging, 18, 1-15.
- Holmes, A. P. & Friston, K. J. (1997). *Statistical models and experimental design*. Retrieved February 27, 2003, from http://www.fil.ion.ucl.ac.uk/spm/course/notes97/Ch3.pdf
- Horowitz, A. L. (1995). MRI physics for radiologists. New York: Springer-Verlag.
- Horwitz, B., Friston, K. J. & Taylor, J. G. (2000). *Neural modelling and functional brain imaging: An overview*. Neural networks, 13, 829-846.
- Jenkinson, M (2001). Registration, atlases and cortical flattening. In Jezzard, P., Matthews, P. M. & Smith, S. M. (Eds.), *Functional MRI: an introduction to methods*. (pp. 271-293). New York: Oxford University Press Inc.
- Jezzard, P. & Clare, S. (2001). Principles of nuclear magnetic resonance and MRI. In Jezzard, P., Matthews, P. M. & Smith, S. M. (Eds.), *Functional MRI: an introduction to methods*. (pp. 67-92). New York: Oxford University Press Inc.
- Lange, N (2000). Statistical procedures for functional MRI. In Moonen, C. T. W. & Bandettini, P. A. (Eds.), Functional MRI. (pp. 301-335). Heidelberg: Springer- Verlag Berlin.
- Lai, S., Glover, G. H. & Haacke, E. M. (2000). Spatial selectivity of BOLD contrast: Effects in and around draining veins. In Moonen, C. T. W. & Bandettini, P. A. (Eds.), *Functional MRI*. (pp. 221-231). Heidelberg: Springer-Verlag Berlin.
- Matthews, P. M. (2001). An introduction to functional magnetic resonance imaging of the brain. In Jezzard, P., Matthews, P. M. & Smith, S. M. (Eds.), *Functional MRI: an introduction to methods*. (pp. 3-34). New York: Oxford University Press Inc.
- Menon, R. S. & Goodyear, B. G. (2001). Spatial and temporal resolution in fMRI. In Jezzard, P., Matthews, P. M. & Smith, S. M. (Eds.), *Functional MRI: an introduction to methods*. (pp. 145-158). New York: Oxford University Press Inc.

- Poline, J-B., Holmes, A. P., Worsley, K. J. & Friston, K. J. (1997). Statistical inference and the theory of random fields. Retrieved February 27, 2003, from http://www.fil.ion.ucl.ac.uk/spm/course/notes97/Ch4.pdf
- Smith, S. M. (2001a). Overview of fMRI analysis. In Jezzard, P., Matthews, P. M. & Smith, S. M. (Eds.), Functional MRI: an introduction to methods. (pp. 215-227). New York: Oxford University Press Inc.
- Smith, S. M. (2001b). Preparing fMRI data for statistical analysis. In Jezzard, P., Matthews, P. M. & Smith, S. M. (Eds.), *Functional MRI: an introduction to methods*. (pp. 229-241). New York: Oxford University Press Inc.
- Uğurbil, K., Toth, L. & Kim, D-S. (2003). *How accurate is magnetic resonance imaging of brain function?* Trends in neurosciences, 26(2), 108-114.
- Veltman, D. & C. Hutton (2001). *SPM99 manual*. Retrieved February 27, 2003, from http://www.fil.ion.ucl.ac.uk/spm/course/manual/manual.pdf
- Worsley, K. J. (2001). Statistical analysis of activation images. In Jezzard, P., Matthews, P. M. & Smith, S. M. (Eds.), *Functional MRI: an introduction to methods*. (pp. 251-270). New York: Oxford University Press Inc.