Fronto-Parietal Cortex in Sequential Behaviour

Ausaf Ahmed Farooqui

Fitzwilliam College

MRC – Cognition and Brain Sciences Unit

This dissertation is submitted for the degree of Doctor of Philosophy at University of Cambridge

August, 2011

Preface

I declare that this dissertation is the result of my own work and includes nothing, which is the outcome of work done in collaboration. Where reference is made to the work of others, the extent to which that work has been used is indicated and duly acknowledged in the text and bibliography.

This dissertation does not exceed 60,000 words in length.

Ausaf Ahmed Farooqui

_____

Acknowledgment

It is a pleasure to thank all those who made this thesis possible.

My supervisor, John Duncan, for his valuable advice and guidance. Working under him has been a very instructive experience. I'm thankful for his encouragement and support for my varied ideas.

My co-supervisor and friend, Russell Thompson, for patiently teaching me almost everything I know of experimental methods, and for being the most accessible and reliable help for almost everything.

The wider community at the MRC – Cognition and Brain Sciences Unit.

The Gates Cambridge Trust, for funding my stay.

To my wife, *shukriya*.

Abstract

This dissertation investigates the fronto-parietal representation of the structure of organised mental episodes by studying its effect on the representation of cognitive events occurring at various positions within it. The experiments in chapter 2 look at the completion of hierarchically organized mental (task/subtask) episodes. Multiple identical target-detection events were organized into a sequential task episode, and the individual events were connected in a means-to-end relationship. It is shown that events that are conceptualized as completing defined task episodes elicit greater activity compared to identical events lying within the episode; the magnitude of the end of episode activity depended on the hierarchical abstraction of the episode.

In chapter 3, the effect of ordinal position of the cognitive events, making up the task episode, on their representation is investigated in the context of a biphasic task episode. The design further manipulated the cognitive load of the two phases independently. This allowed for a direct comparison of the effect of phase vis-à-vis the effect of cognitive load. The results showed that fronto-parietal regions that increased their activity in response to cognitive load, also increased their activity for the later phases of the task episode, even though the cognitive load associated with the later phase was, arguably, lower than the previous phase.

Chapter 4 investigates if the characteristics of the higher-level representations, like organization of task descriptions, have a causal role in determining the structure of the ensuing mental episode. Results show this to be true. They also confirm the results of earlier chapters in a different framework. Chapter 5 shows that the effect of episode structure is not limited to the elicited activity, but also affects the information content of the representation of the events composing the episode. Specifically, the information content in many regions of later steps is higher than that of earlier steps.

Together, the results show widespread representation of the structure of organised mental episodes.

Table of Contents

CHAPTER 1

Introduction

Understanding the fronto-parietal regions presents a challenge to neuroscience. Activity in these regions can be found in almost all neuroimaging studies across almost all kinds of task conditions. However, the deficits caused by their lesions are variable and difficult to characterise and range from subtle and barely characterisable deficits to gross disorganisations in various domains of behaviour. In the following review some of the salient findings concerning these regions are presented. It is argued that they make a strong case that these regions represent all kinds of currently attended task-related representations.

In the subsequent section, the case of goal directed behaviour is then discussed; it is hypothesised that the representations organising episodes of such behaviour will also be coded in fronto-parietal regions and affect the representation of individual behavioural events occurring in the episode. Two aspects of organised behaviour are specifically reviewed – completion of discrete episodes and the effect of the ordinal position of task events within the sequence making up the episode.

Finally, a summary of subsequent chapters, looking at the various issues raised, is presented.

1.1 Fronto-Parietal Regions

1.1.1 Prefrontal regions and Control

Lesion studies suggest that the prefrontal cortex (PFC) is necessary for performance in situations that require the operation of a number of cognitive processes. Bianchi (1922) characterised the behaviour of PFC lesioned monkeys as having lost the ability to coordinate the elements of a complex activity. Patients with damage to PFC are described as 'having difficulty in grasping the whole of a

complicated state of affairs, well able to work along old routine lines. But they cannot master new types of tasks, in new situations' (Rylanders, 1939). Luria (1966) documented that such patients had problems in programming non-routine aspects of behaviour including the preliminary analysis of situations, constructing a plan of action and monitoring their performance. These and other case studies (e.g. Milner, 1964; Grafman et al., 1986) make a strong case for the conceptualisation of the PFC as an 'executive' or a 'cognitive control centre' controlling the representation and processing in other brain regions (Norman and Shallice, 1986; Desimone and Duncan, 1995; Miller and Cohen, 2001).

More structured experimental studies of prefrontal lesion patients have substantiated the view that the deficit in such patients is that of an absence of control in their thought and behaviour. Such patients are unable to abstain from doing a habitual but irrelevant action. For example, they are impaired at making eye movements diametrically away from peripherally presented stimulus, and usually end up making a saccade to the visual stimulus. Plausibly, the antisaccade requires overriding the habitual tendency of making saccades (Milea et al., 2003). Similar tendencies are also manifest in the difficulty faced by such patients in the Stroop task (Stroop, 1935) that requires them to inhibit their tendency to read the word and instead name the colour of the ink in which the words are written (Perret, 1974; Vendrell et al., 1995), and in the Wisconsin Card Sorting Test (WCST; Grant and Berg, 1948) that requires dismantling of the now irrelevant task set and the creation of a new one as per the now relevant rule (Milner, 1964). Severe prefrontal lesioned patients, at times, show frank utilisation behaviour wherein the mere presence of an object elicits its use, notwithstanding the social propriety of such an act; e.g. the presence of a toothbrush on the doctor's table elicited the act of brushing one's teeth (Lhermitte, 1983). Although utilisation behaviour involves actions that have become habitual in the course of life, some prefrontal patients also show

reappearance of primitive reflexes (De Renzi and Barbieri, 1992) that are usually present early in life and disappear with the maturation of the brain.

Prefrontal patients have been found to be impaired on a host of other processes that would be important for purposive behaviour in novel and poorly characterised situations such as – sustained attention (Luria, 1966; Chao and Knight 1995), attentional shifts (Windmann, 2006), controlling the interference from irrelevant distractors (Chao and Knight 1997), working memory (Milner, 1964; Lewinsohn et al., 1972), organising mnemonic information (Shimamura et al, 1991), temporal integration of cognitive activities (Fuster, 2008), monitoring (Swick and Turken, 2002) and inhibitory control (Aron et al., 2004).

Indeed prefrontal lesions may decrease the common cognitive capacity ('g'; Spearman, 1904) required for the performance on a wide variety of tasks (Duncan et al., 1995). This could explain the deficit seen in prefrontal patients on a wide variety of tasks. In line with this, Roca et al. (2010) found that the deficit of frontal patients on a wide variety of executive processes could be explained by their concomitant deficits on tasks of fluid intelligence. In other words, the performance across the patients and controls was equal if they were matched on the measures of fluid intelligence. It is possible that prefrontal cortices are required for some domain-general process that is required for control in all kinds of task requirements.

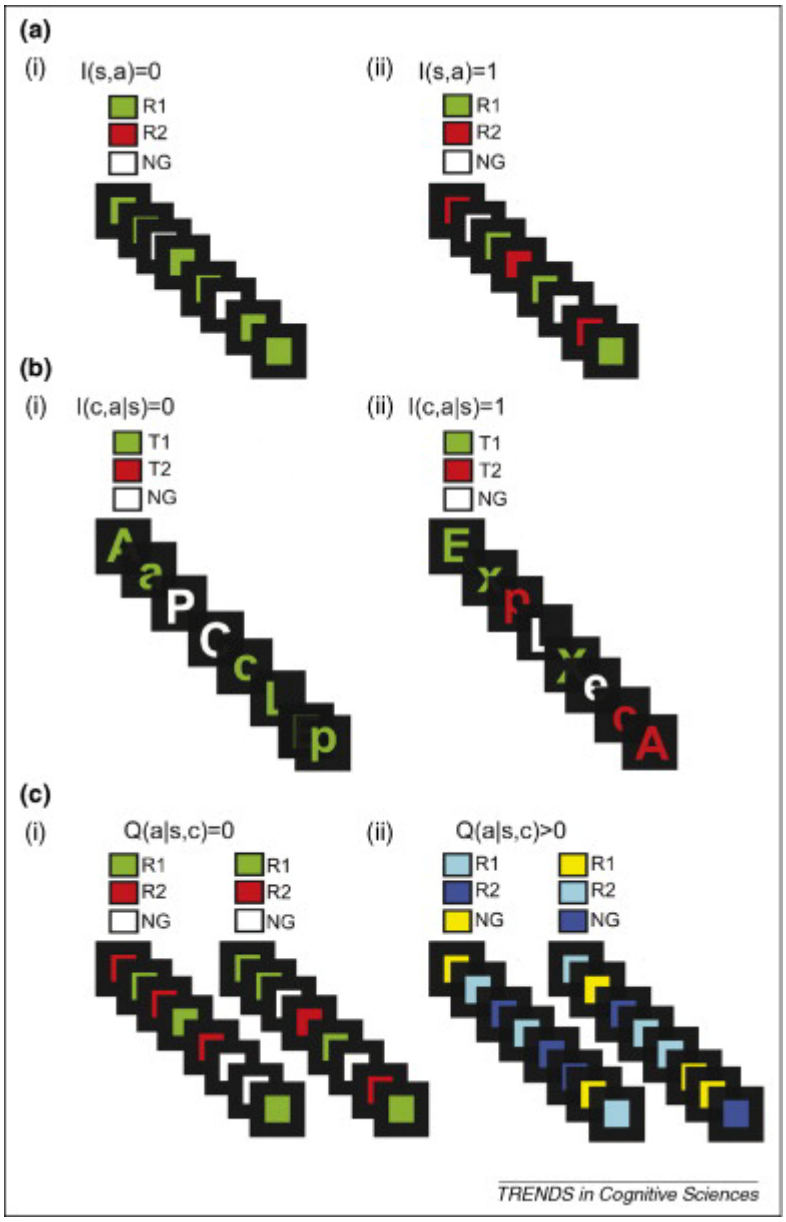

Figure 1.1. Description of the colour and letter tasks used by Koechlin et al (2003). See text for details. Reproduced from Koechlin and Summerfield (2007).

Neuroimaging studies routinely show greater activity in the PFC in conditions where lesions impair the performance (reviewed in Fuster, 2008). Specific to the discussion here, greater activity is seen when greater task contingencies are to be represented for correct task execution (Koechlin et al., 2003; Badre and D'Esposito, 2007). Koechlin et al. designed a study in which they systematically varied the quality as well as the quantity of the information required for the execution of a trial block (Fig. 1.1). Subjects responded to coloured shapes and letters (by making simple

colour and letter judgements) or withheld their responses according to the rule cued at the beginning of the block. The rules could remain the same across the blocks or vary from block to block or even trial to trial. In the blocks depicted in Fig. 1.1a, the subjects respond to green colour by their left hand and to red by their right hand. They are not supposed to respond to the white colour at all. In low information load blocks of this kind (Fig. 1.1a.i), only one stimulus type appeared and hence the same response was to be made on all trials; on high information load blocks of this category both kinds of stimuli could appear (Fig. 1.1a.ii) and so potentially two kinds of responses could be made on any trial. Thus these blocks manipulated the number of valid stimulus-response associations.

In the blocks depicted in Fig. 1.1b, higher-level information, i.e. the relevant rule, was manipulated. Green coloured letter stimuli were to be categorised as vowel or consonant, whereas the red coloured ones were to be categorised as small or capital letters. One (Fig. 1.1b.i) or both (Fig. 1.1b.ii) of these rules could be relevant on any of such blocks. Note that the rule has been referred to as higher-level information because it subsumes more than one of the previous category (stimulus-response associations); hence, to make a decision on the valid stimulus-response association a decision on the relevant rule has to be made first.

Finally in the blocks depicted in Fig. 1.1c, the colour-to-rule mapping changed across the blocks. Performance in such blocks required first resolving the valid colour-to-rule mapping, only then could the valid rule be decided about. Thus the information additionally manipulated in these blocks subsumed the rule information and hence can be regarded as belonging to an even higher level.

The relation between these three categories of manipulated information is nested, such that increase in the information of the higher category simultaneously increased the information of the lower categories e.g. increasing the number of rules

linking the stimulus to the response would also increase the number of stimulus-response combinations.

The authors found that blocks with higher content of relevant information led to higher prefrontal activity. They reported a hierarchical pattern of results wherein an increase in the magnitude of contingencies related to stimulus-response combinations (Fig. 1.1a) increased the activity only at the posterior loci of the PFC, while increasing the number of rules (Fig. 1.1b) increased the activity in middle and posterior prefrontal regions, and increasing the cue-to-rule contingencies (Fig. 1.1c) additionally increased the activity at anterior prefrontal regions. These results, however, cannot be considered conclusive about the relation between the locus on the antero-posterior axis and the qualitative type of manipulated information load. For example, it is not clear if the anterior prefrontal cortex is so specific to the cue-to-rule contingency that *any* amount of information load related to rules relevant on a trial would not modulate its activity. Nonetheless, these results do show that, as a general rule, increasing the information load of task-related contingencies led to greater and more widespread PFC activity.

Figure 1.2. Sample trial from a Goal Neglect task. Trials began with an instruction to attend to the left/right side of the fixation spot. The letters were to be read aloud, pairs of numbers were to be added, asterisks were to be ignored. '<'/'>' signalled the side to be attended for the succeeding stimuli. Reproduced from Duncan et al. (2008).

Another direct demonstration of the role of PFC in the creation and utilisation of relevant task contingencies for the control of behaviour comes from the goal neglect studies of Duncan et al. (1996, 2008) (see also Dumontheil et al., 2010). They found that in prefrontal patients and healthy individuals with low fluid intelligence, instructions may be understood and remembered but yet don't exert control over behaviour. In a sample task (Fig. 1.2), subjects are shown a stream of pairs of stimuli (letters or numbers) presented simultaneously on the two sides of the centre of a

computer screen. The last three pairs of stimuli of the stream were preceded by a presentation of '>' or '<'. The subjects are to follow a series of instructions in a particular sequence – attend to left (or right), read aloud the letters on the attended side, add the pairs of numbers, but ignore the asterisks, after the appearance of '>' read aloud the letters on the right, if '<' appears, read aloud the letters on the left side.

They found that the last part of such instructions (in the above example, those pertaining to '>' or '<') was sometimes ignored by patients with lesions in their PFC, or by subjects with low fluid intelligence scores. They further found that the frequency of such neglect of a particular instruction was dependent on the complexity and load of task instructions given prior to the instruction in question, even when the rules relevant on the particular block were themselves not complex. Further, the ignored instructions, though neglected, were nonetheless remembered. These findings make a strong case that in prefrontal patients the capacity to construct task control representations is limited. Moreover, although particular information may be present in the mind, they are not incorporated in the relevant control structure, if the latter has reached a sufficient degree of complexity.

In an fMRI study of a related design, the authors discovered that the presentation of each rule phasically activated multiple fronto-parietal regions and tonically increased the baseline activity in those regions (Dumontheil et al., 2010), as if the incorporation of a rule into the task control structure is related to an increase in the baseline activity in fronto-parietal regions.

Goal neglect studies have also looked at the complexity of task control mental structures that can be created. Prefrontal lesion patients are also impaired at carrying out prospective memory tasks that required them to execute intended actions after a delay period filled with irrelevant tasks (Volle et al., 2011). While the experiments reviewed above (Duncan et al., 2008; Dumontheil et al., 2010) required

an immediate creation of a task control structure, prospective memory tasks require that some information be held till a future time when, plausibly, it will be used to create a control structure. Prefrontal representations thus seem important not only for immediate control of behaviour but also for control requirements that will come into play in the future (see also Fuster, 2008).

Finally, stimulation of prefrontal regions can enhance performance specifically in situations that require greater cognitive control. Stuphorn and Schall (2006) trained monkeys on a stop signal task that required them to make saccades to peripherally presented stimuli on most trials (called *go* trials), but on a fraction of trials (called *stop* trials) after the appearance of the *go* stimulus, they were signalled to cancel the prepared saccade. As can be expected, slowing down on the *go* action in general would improve performance since more time would be available for the *stop* action to come into play and cancel the *go* action that as yet has not been fully prepared. On the other hand, if the *go* action is executed quickly, the *stop* action will more often lose the race.

The authors found that stimulating the supplementary eye fields delayed saccade initiation and hence improved the performance on the stop signal task. Importantly, initiation of simple visually guided saccades was not delayed when that was the only task to be done. This shows that the activation did not lead to a general slowing down of saccadic initiation, but enhanced cognitive control.

1.1.2 Parietal Regions and Control

While the above account focussed on the PFC, similar evidence exists for the role of the parietal regions in the control of behaviour. Parietal lesions have been extensively documented to cause attentional dysfunctions both spatial (Posner et al., 1984; Duncan et al., 1997; Driver and Mattingley, 1998; Smania et al., 1998; Mesulam, 1999) and non-spatial (Husain and Rorden, 2003; Husain and Nachev,

2007) as well as a reduction in attentional capacity (Duncan et al., 1999) and a reduction in capacity for sustained attention (Robertson et al., 1997). Beyond attention, both spatial (Husain et al., 2001) and non-spatial (Koenigs et al., 2009) working memory deficits have been associated with these lesions. Classical tests of frontal dysfunction also show impairments in parietal lesioned patients, e.g. WCST (Anderson et al., 1991), Stroop (Pujol et al., 2001). Case reports of frank utilisation behaviour after parietal lesions also exist (Mizobuchi et al., 2011). In an interesting study Desmurget et al. (2009) stimulated inferior parietal regions in tumour patients undergoing surgery. They found that in the context of motor action, this elicited a feeling of intention for action but not the action itself. Intention being an abstract feeling has usually been associated with activity in the prefrontal regions (Libet et al., 1983; Lau et al., 2004; Haynes et al., 2007).

It is likely that frontal and parietal regions are part of the larger brain network that allows for behavioural control. Such a view gets strong support from neuroimaging studies, most of which find a co-occurrence of frontal and parietal activity (reviewed in Cabeza and Nyberg, 2000; Duncan and Owen, 2000; Corbetta and Shulman, 2002; Dosenbach et al., 2006; Duncan, 2006). While both frontal and parietal regions seem to be necessary for the controlled cognition, meta-analyses of the neuroimaging studies showing activation in frontal and parietal regions suggest that some regions within the fronto-parietal cortices are more important for control.

Duncan (2006) reviewed 20 studies that manipulated what would be considered as different forms of control – response conflict, task novelty, number of elements in the working memory, working memory delay and perceptual difficulty. He found that reported peaks of a contrast between high versus low control conditions clustered in common regions of the fronto-parietal cortex (Fig 1.3a) – inferior frontal sulcus (IFS), anterior insula extending into frontal operculum (AI/FO), supplementary motor area/anterior cingulate (SMA/ACC) and intraparietal sulcus

(IPS). A number of other studies investigating as disparate forms or control as would manifest in visual attention (Hon et al., 2006; Jiang et al., 2000), auditory discrimination (Holcomb et al., 1998), self paced response production (Jahanshahi et al., 1997), response conflict (Bush et al., 1998), working memory (Rypma et al., 1994), perceptual difficulty (Woolgar et al., 2011), task switching (Dove et al., 2000), episodic memory (Duncan and Owen, 2000), conscious perception (Dehaene and Changeaux, 2011), complex response mapping (Jiang and Kanwisher, 2003a), semantic (Thompson-Schill et al., 1997) and syntactic processing (Jiang and Kanwisher, 2003b), find activity clustered in these regions of the fronto-parietal cortices. On this basis these regions have been referred to variously as 'multiple demands' (Duncan, 2010), 'task positive' (Fox et al., 2005), 'task-activation ensemble' (Seelay et al., 2007) or a 'task control' network (Dosenbach et al., 2006) and (Cole and Schneider, 2007).

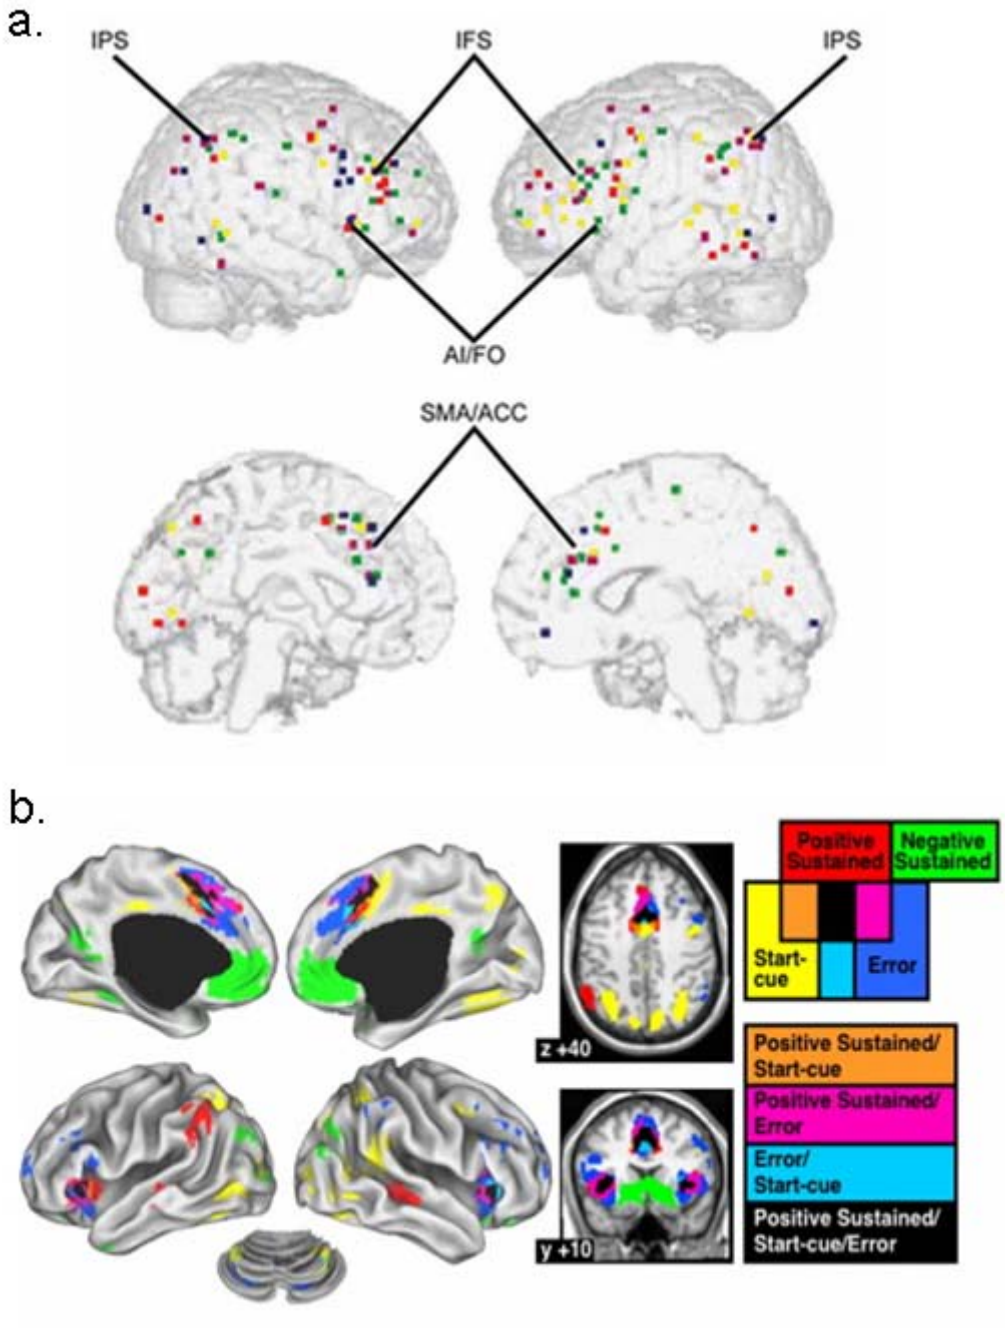

Figure 1.3 (a) Clustering of reported peaks of activation of studies investigating various forms of control; response conflict (green), task novelty (purple), number of elements in working memory (yellow), working memory delay (red) and perceptual difficulty (blue). IFS: inferior frontal sulcus. AI/FO: anterior insula/frontal operculum. SMA/ACC: supplementary motor area/anterior cingulate. IPS: intraparietal sulcus. Reproduced from Duncan, 2006. (b) Regions with activity common across different tasks in Dosenbach et al., 2006. Reproduced from Dosenbach et al., 2006.

Conceptually similar results were found by Dosenbach et al (2006), who conjointly analysed data from experiments using ten very different tasks letter identification, verb generation and reading, object naming, matching symbols and letters, living/nonliving judgements, reading aloud, motor timing, visual detection, abstract/concrete and physical/semantic judgement on nouns. They found that irrespective of task, discrete aspects of trial blocks were associated with common patterns of activity across the different tasks (Fig 1.3b). The starting cue in all tasks involved modulation of activity in ACC, AI/FO, intraparietal sulcus (IPS), precuneus, posterior cingulate cortex (PCC) and fusiform regions. The duration of the trial blocks across the tasks were associated with sustained change of activity in ACC, AI/FO, anterior prefrontal cortex (APFC), intraparietal lobule (IPL), middle and posterior temporal regions, PCC and the ventromedial prefrontal cortex.

Evidence for a causal relation between these specific areas of the fronto-parietal cortices and diverse forms of cognitive control was recently provided by Woolgar et al. (2010). They looked at the relation between these multiple demands regions and the general factor *g* that has been hypothesised to contribute to performance in diverse kinds of cognitive activity (Spearman, 1927), and found that damage to each of the above regions predicted loss in *g*, whereas damage to fronto-parietal regions outside these multiple demands regions was not predictive.

These findings make interesting suggestions. A set of fronto-parietal regions are always co-active in response to seemingly diverse conditions, and plausibly, these form a core component of diverse forms of cognitive control. Secondly, and as a corollary to the previous point, the seemingly diverse forms of control do not appear to be anatomically diverse. It is possible that the absence of expected diversity is due to the poor spatio-temporal resolution of fMRI. However, it is also possible that these regions are sensitive to some core component that is common to all of the above reviewed control processes, or perhaps that the apparently diverse control

processes are actually not diverse, but are the manifestations of a more abstract form of control.

1.1.3 Representations in the Fronto-Parietal regions

The above review demonstrates the role of fronto-parietal regions in behavioural control. However, the mechanism by which these regions lead to control is far from clear. A good starting point in this regard is to investigate what kind of mental/behavioural events are represented in these regions.

Neurophysiological studies show that neurons in multiple fronto-parietal regions respond to all behavioural and mental events that are relevant and consequential in the context of the task at hand, irrespective of the actual identity of the task (Duncan, 2010). Kusunoki et al (2009) trained monkeys to covertly monitor a stream of different pictures presented on the right or left (varying across trials) of the fixation spot for a specific target, and to specify its occurrence by making a saccade to the relevant side at the target stimulus offset (Fig 1.4a). Across the task variants the identity of the target could remain constant across trials (called fixed target task) or could vary in case of the cued target trials. In the latter version of the task, monkeys learnt three cue-target associations. In such trials, the identity of the target stimulus varied across trials in such a way that the target on one trial could be a non-target on the other. Many of the non-target stimuli on any trial had served as targets on previous trials at some point during stimulus presentation. The firing profile of more than half of all neurons recorded from the lateral PFC discriminated the target from the non-target events (Fig 1.4b). Very few neurons discriminated individual non-targets, including those that had been targets in earlier trials.

Similar coding of task-relevant events has been found in a number of studies, wherein the neurons have been shown to code for task rules (Asaad et al., 1998; Hoshi et al., 1998; White and Wise, 1999; Asaad et al., 2000; Wallis et al., 2001;

Wallis and Miller, 2003), relevant response (Quintana and Fuster, 1992; Asaad et al., 1998), cues, rewards and cue-reward association (Watanabe, 1990, 1992, 1996). With regard to stimulus features, fronto-parietal neurons seem to code for any feature that is linked to a behaviourally relevant event – stimulus identity (di Pellegrino and Wise, 1991; Rao et al., 1997), location (Azuma and Suzuki, 1984; Rao et al., 1997), colour (Quintana and Fuster, 1992), orientation (Mikami et al., 1982), vibration frequency (Romo et al., 1999), category membership (Freedman et al., 2001) and so forth.

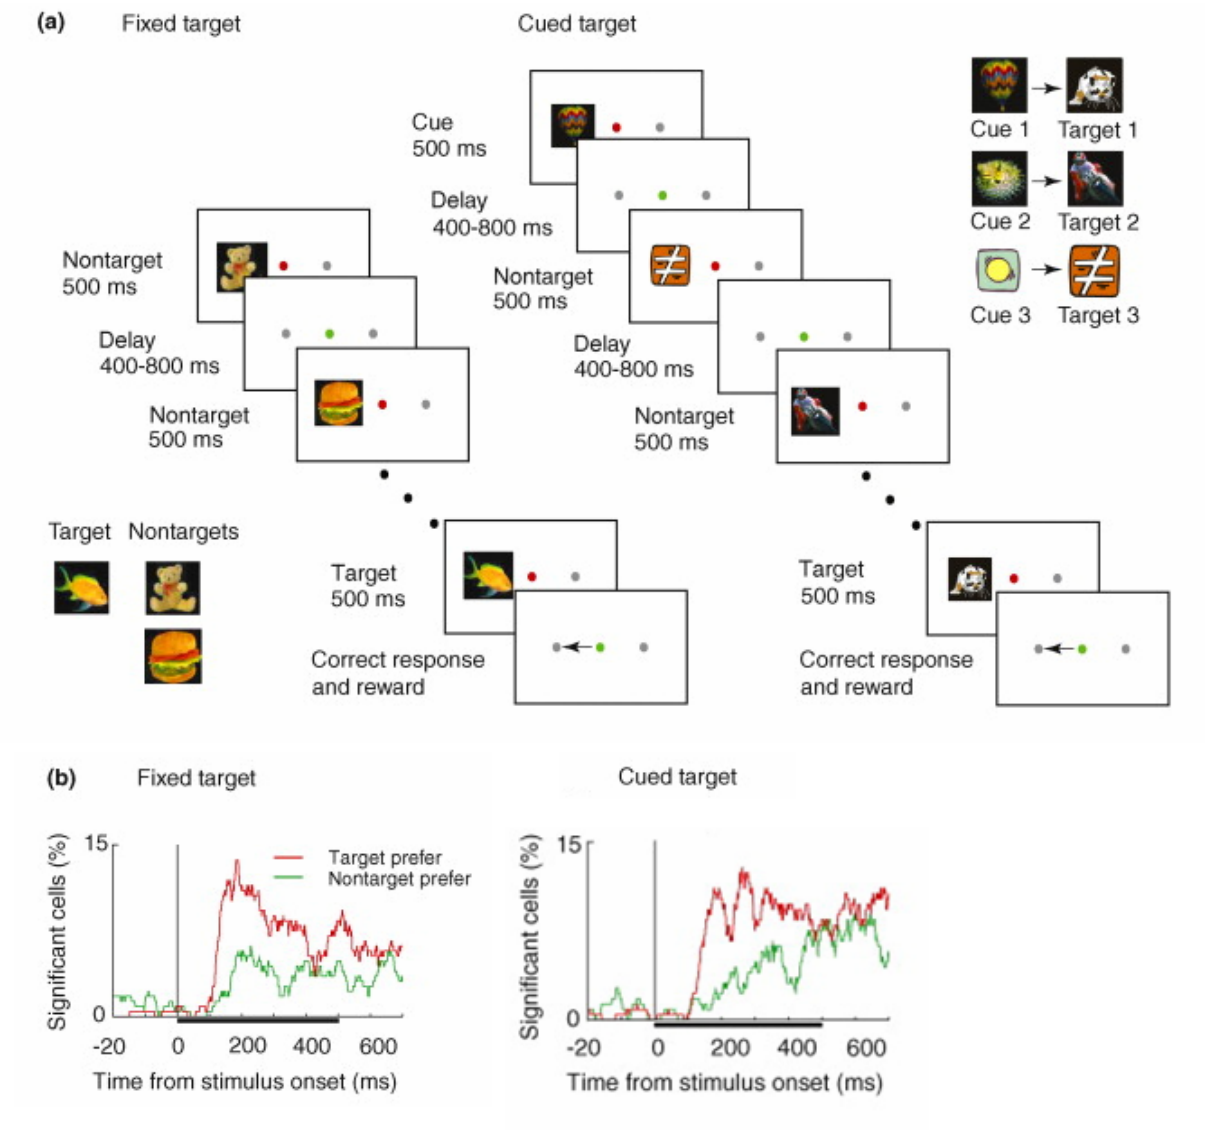

Figure 1.4. (a) Scheme of the task used by Kusunoki et al. (2009). Monkeys viewed a stream of stimuli, presented to one side of the fixation spot, and monitored for the presence

of the target picture, and at its occurrence waited for its offset and then made a saccade to the side it occurred. The targets were fixed across trials in the fixed target condition. In cued target condition, the target on any trial was determined by the nature of the cue at its beginning. Three cue-target combinations were used. (b) Percentages of cells in lateral prefrontal cortex that discriminated targets from non-targets (red – stronger response to targets; green – stronger response to non-targets). Heavy black line shows time of stimulus presentation; discrimination was tested in sliding windows beginning 200 ms before stimulus onset. Reproduced from Duncan, 2010.

Indeed, an important property of these neurons is the flexibility of their coding. Often, neurons cease responding to an event that has lost its importance and assume responding to the now relevant event. This property was seen in the above mentioned study, wherein minimal response was seen for non-targets that had been targets in earlier trials. The rapidity of this adaptation was amply demonstrated in a study by Rao et al. (1997). In their task (Fig. 1.5a), monkeys were shown a sample target stimulus at fixation, the identity of which they had to remember during the following delay ('what delay'). The screen succeeding the delay showed two objects at the periphery, one of which was the earlier shown target. Now the monkeys had to remember the location where the target appeared in the succeeding delay ('where delay'). Finally they were shown a screen with four locations and had to make a saccade at the place where the target had appeared. Thus in this task, the monkeys had to first code for the identity of the target and maintain it during the 'what delay', then represent the spatial location of the target during the 'where delay'. Again roughly half of the neurons recorded coded first for the identity of the target and then its location (Fig. 1.5b).

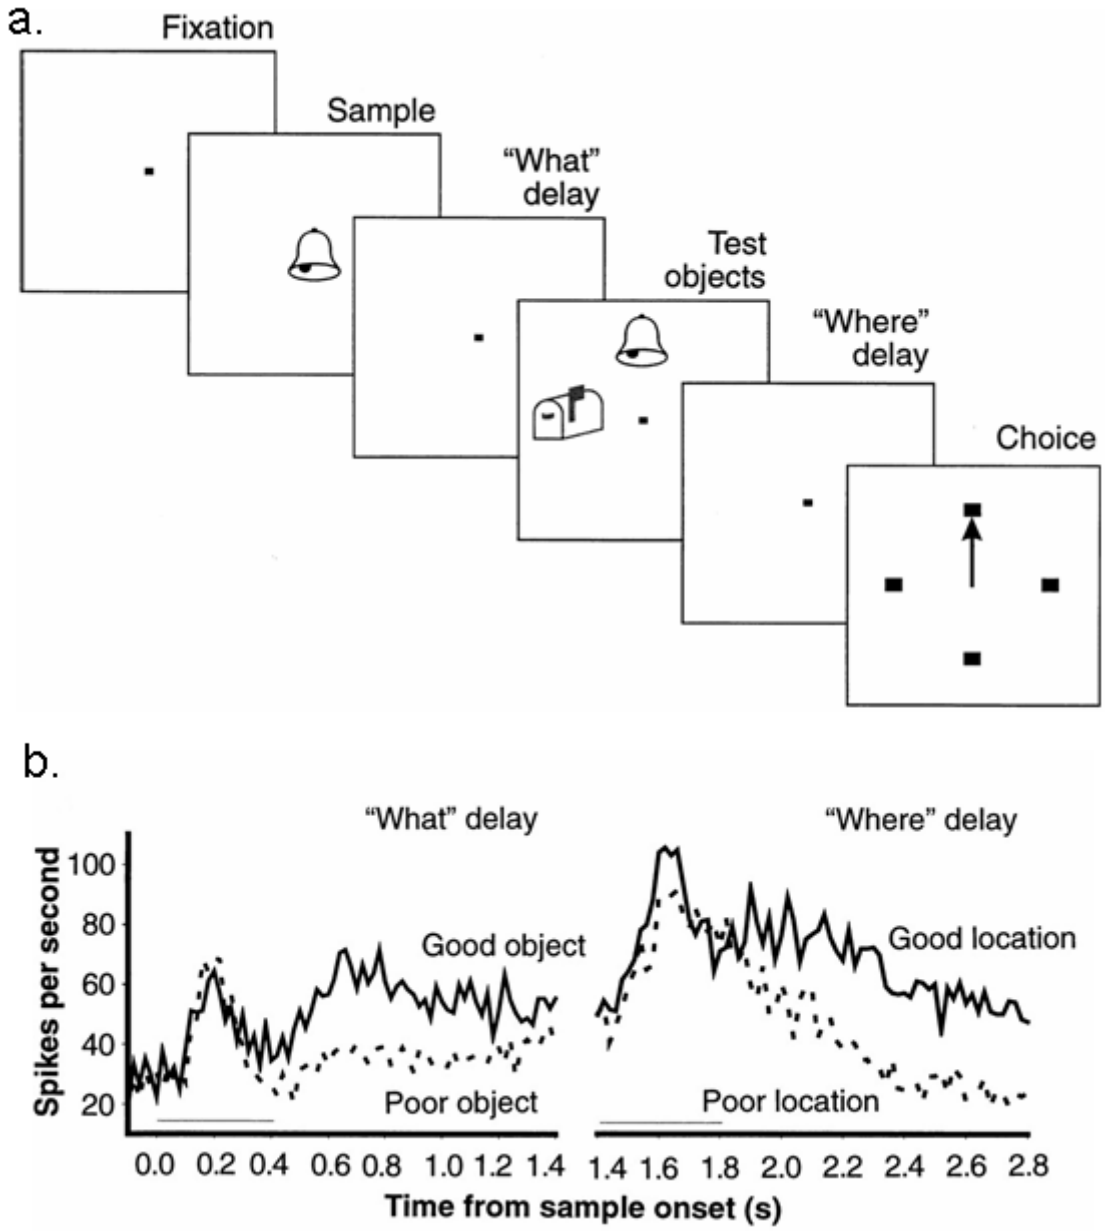

Figure 1.5. (a) A sample trial from Rao et al., (1997); trials began with a sample stimulus presented at the fixation spot, the identity of which was to maintained during the 'what' delay, which was followed by the peripheral presentation of two stimuli, one of which was the sample stimulus. This time the position of the sample stimulus was to be remembered across the 'where' delay, after which monkeys were probed for the position of the sample stimulus, to which they made a saccade. (b) Sample prefrontal neuron that codes for the identity of the object during 'what' delay and the location during 'where'; the two grey bars represent the presentation of sample (left) and test objects (right), respectively . Reproduced from Rao et al., (1997).

A number of human fMRI studies also provide similar evidence (Hon et al., 2006; Thompson and Duncan, 2009; Serences and Boynton, 2007, Haynes et al., 2007). Hampshire et al (2006) designed a task that required subjects to monitor a sequence of visual stimuli and covertly detect the presence of a visual target that had been cued at the beginning of the trial. The stimulus sequence could terminate randomly with a probe asking if the previous stimulus in the stream was the target to which subjects made a yes/no response, otherwise, the detection of visual targets was covert. The actual identity of the target changed across trials. Thus, the target stimuli were made relevant in the context of that trial by the cue but they had the same decision-making and task load requirements as the non-target stimuli. Across many fronto-parietal regions a selective response was seen to the targets, proving that fronto-parietal regions have a predilection for currently relevant events. Interestingly, some prefrontal regions also showed a higher response to other stimuli belonging to the same category as the target. For example, if the target was a particular face, such regions additionally showed greater response to all faces.

Other studies have used multivariate pattern analysis to show a broad representation of attended and task-relevant cognitive content across the fronto-parietal regions (Dux et al., 2009; Eger et al., 2009; Esterman et al., 2009; Greenberg et al., 2010; Jenkins and Ranganath, 2010;). The conceptual basis of this methodology is that if a category of cognitive events (e.g. relevant rule) is represented in a region then the different exemplars of that event (e.g. different rules) will be represented by different neuronal population codes. Since it is very likely that the distribution of such population codes would be different across the voxels composing the region, the pattern of fMRI activity elicited by the exemplars of the represented cognitive event can be expected to be different. So the discriminability of the patterns of activity elicited by the exemplars of the cognitive event can be considered as a measure of the representation of the cognitive event.

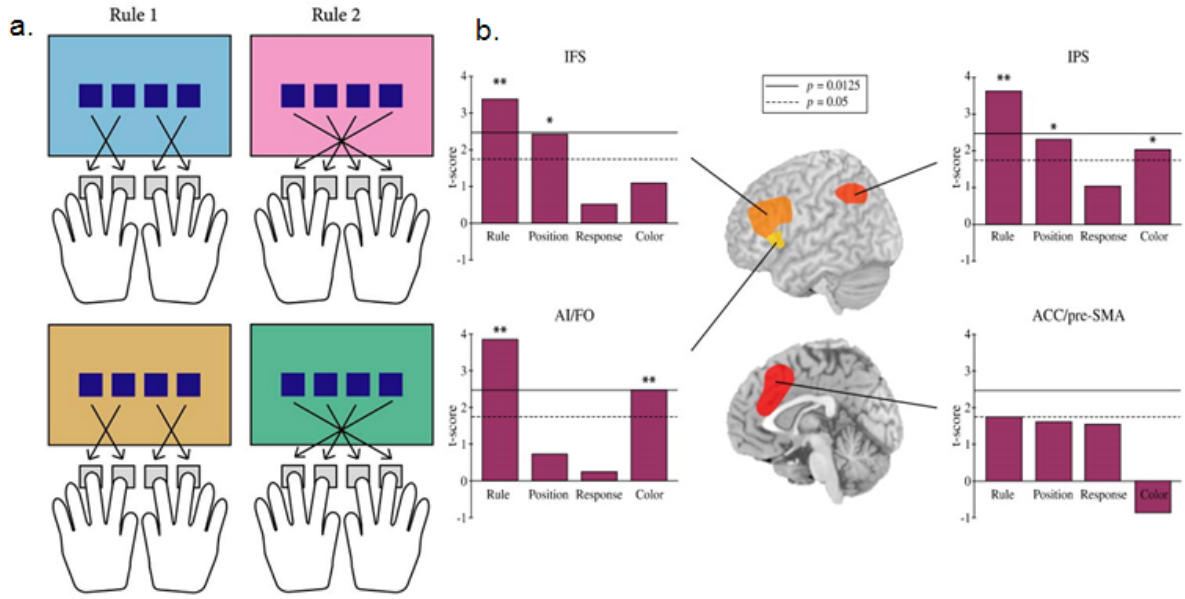

Figure 1.6. (a) Task used by Woolgar et al., (2011); subjects responded to one of the four positions of the stimulus by a key press. Two mapping rules were used, and were cued by the colour of the background. Two colours mapped to each rule. (b) t-score on a one-sample t-test of classification accuracy of patterns of elicited activity against chance. A score of 0 indicates chance level classification and positive values indicate above chance classification. Dotted line: significance threshold p = 0.05; solid line: significance threshold corresponding to Bonferroni-corrected p-value of 0.0125. Reproduced from Woolgar et al., 2011.

In one such study (Woolgar et al., 2011), participants made simple key press responses to stimuli presented at four different positions in a horizontal row on a computer screen. Two rules could be used to map the stimulus position to the response (Fig. 1.6a), which could be one of the four alternative key presses made with index and middle fingers of the two hands. The rule relevant on a trial was cued by the background colour. Two colours were mapped to each rule. This design allowed the authors to look into the discriminability of the pattern of activity elicited by the different colours, response fingers, stimulus positions and rules. The results showed that all features were discriminable in fronto-parietal regions (Fig. 1.6b). Evidence was strongest for the coding of rule followed by colour and stimulus position; responses were least discriminable and their discriminability was significant

only in left inferior frontal sulcus (IFS). Similar results were also obtained by Stiers et al. (2010).

Other studies have demonstrated representation of more abstract task-related information such as the outcome of a decision (Haynes et al., 2007; Kuhn et al., 2010; Bode et al., 2011). Haynes et al (2007) asked subjects to freely decide which of two operations (addition or subtraction) to perform on a forthcoming display of numbers. They were to hold this intention for a variable gap, after which they saw the numbers and performed the intended operation and thus indicated which of the two operations they had decided for. The authors showed that the two decisions were discriminable in a number of prefrontal regions before the actual execution. Hence prefrontal regions represent not only concrete task rules but also abstract task-related decisions.

The representations in fronto-parietal regions are not only limited to those that are explicitly related to a task at hand. Rather, any conscious and attended mental representation seems to be represented in the fronto-parietal regions. Lumer and Rees (1999) found that switches in between bistable percepts (e.g. in binocular rivalry) in the absence of any task-related requirements of decision, control or response, caused activity in the fronto-parietal regions.

A recent study made a strong case that fronto-parietal regions represent any conscious mental representation and activity in them is elicited whenever such representation is changed and a new one is created, irrespective of whether the change has been initiated by the self or exogenously. Knapen et al. (2011) related the magnitude of elicited fronto-parietal activity to the temporal duration of transitions in between the two percepts in bistable viewing conditions. They looked at two forms of bistable percepts. In bistable apparent motion, subjects were shown stimulus with equal motion energy in two opposite directions resulting in the wavering of perception in between the two stable motion percepts. The other

bistable percept they looked at was binocular rivalry, in which the two eyes were given dissimilar inputs that cannot be fused together, and so subjects' percepts swayed alternatively between the two percepts projected to the eyes. In both, binocular rivalry and bistable apparent motion, subjects experience mixture states, when one percept is gradually shifting over onto the other and observers experience a mix of both perceptual interpretations. The duration of this transition period can be variable. They found fronto-parietal activity at the transitions between the percepts, the magnitude of which was dependent on the duration of the transition periods.

Further, they compared the activity elicited by such endogenous shifts between bistable percepts to that elicited when the subjects were exogenously replayed two monostable stimuli corresponding to the two perceptual interpretations they had of the bistable stimuli. In such conditions the cause of the transition of the percept was exogenous and its duration was manipulated by the experimenters. When the duration of transition was controlled for, no difference was found in the magnitude of elicited fronto-parietal activity in the two (exogenous and endogenous) conditions. This makes a strong case that the activity in the fronto-parietal regions was not the cause of transition between the percepts in the bistable conditions, but was the effect of the transitions.

A number of studies have shown greater fronto-parietal activity when a stimulus is perceived consciously compared to subliminal processing (reviewed in Dehaene et al., 2006). However, in most such studies consciousness of a stimulus is usually associated with better decisions made on the stimulus manifested in better performance on the associated task. Such studies do not make it clear if the associated increase in fronto-parietal activity is the result of better task performance on the stimulus or is specifically linked to the subjective conscious perception of it. To circumvent this confound, Lau and Passingham (2006) used metacontrast masking

to create experimental conditions in which the objective task performance on the stimulus was the same but the subjective awareness of that stimulus was not. They found that activity in prefrontal regions was associated specifically with conscious awareness even when the objective task performance on the stimulus had been controlled for.

These regions not only represent the conscious perceptual representations but also the endogenously created contents of imagery. Stokes et al. (2009) asked subjects to visually imagine an 'X' or an 'O', and found that the resultant patterns of activity could be significantly discriminated in anterior insula. Indeed some vegetative state patients can flexibly and reliably elicit distinct patterns of fronto-parietal activity when asked to create an image of common situations like playing tennis or moving around their house (Owen et al., 2006; Monti et al., 2010).

While attention on and consciousness of the mental representation certainly is a feature associated with fronto-parietal representations, some studies suggest that behaviourally relevant information may be represented in these regions in the absence of conscious awareness. Lau and Passingham (2007) cued subjects to do one of two tasks on a displayed word – judge whether it is bisyllabic or whether it referred to something concrete. The cue (diamond or square shape) signalling which task to be done was displayed before the display of the word. On some trials an additional priming cue was displayed very briefly prior to the main cue display. Authors tried to achieve meta-contrast masking of the prime cue by the main cue, resulting in varying levels of perception of the former. On some trials subjects reported consciously seeing the first transient cue; on others it was entirely subliminal. The priming cue could be congruent or incongruent with the actual cue. Interestingly, greater prefrontal activity was elicited on the trials in which the main cue was preceded by incongruent but subliminal prime, suggesting that activity in this region could be triggered by an unconscious event.

However, this and other such studies (Boy et al., 2010; Gaal et al., 2010; Reuss et al., 2011) showing unconscious events triggering fronto-parietal activity do not conclusively prove if this activity is related to an unconscious event, since it is very possible that the greater activity reflects the conscious feeling of difficulty faced by the subjects due to the activation of task irrelevant representations. Soon et al. (2008) showed more conclusive evidence of the presence of information in the fronto-parietal regions to which subjects did not have a conscious access. They asked subjects to randomly choose to make one of two actions (left or right index finger button press), and then proceed on to making the decided response. To allow subjects to convey the time at which they became aware of their decision, a stream of letters, updated every 500 ms, was displayed on the screen. Subjects conveyed the time of their awareness of decision by taking note of the current letter on the screen. The authors discovered that the identity of the future decisions could be predicted from the pattern of activity in the anterior and medial prefrontal and medial parietal regions, seconds before the subjects reported becoming aware of them.

The findings reviewed suggest that fronto-parietal activity is elicited by – and the pattern of elicited activity contains information about – any behaviourally relevant event. However, as discussed in the next section, behaviourally relevant events are situated in the larger context of a purposive mental (or task) episode, and play crucial roles in the organisation of such an episode. It is possible that the fronto-parietal activity seen with behaviourally relevant events is the result of the concomitant changes in the structure of the mental episode.

Consider, for example, the case of visual target detection. Subjects are required to search for a stimulus; the search ends when the target is detected. Can the neural activity found with the detection of the target be attributed solely to the event of target detection? Note that target detection also corresponds to the

boundary between two distinct mental episodes (that of visual search and rest). It is very much possible that the fronto-parietal activity seen, corresponds to such a change in the mental episode rather than attentional detection. Similarly, in many of the findings reviewed above, the behaviourally relevant event is associated with salient roles in the organisation of the task episode. Indeed a criterion for behavioural relevance would be that the event brings about a change in the current mental/behavioural state (or episode) or has clear implications for succeeding mental processes (or how the subsequent mental events are to be organised.

1.2 Mental Organisation

Much of our behaviour and mental life is organised. Consider a sequence of actions for preparing a hypothetical breakfast:

pour water → add coffee → pour cream → toast bread → spread butter

Arguably, in real behaviour such a sequence of actions can only exist in the context of a hierarchical set of organising representations (Lashley 1951, Miller et al., 1960, Schank & Abelson, 1977, Reed et al. 1995), which ensure that the correct action is chosen at each step. Abstract representations defining temporally extended goals organise nested lower-level representations that have temporally more confined roles, define more concrete goals and help select more defined aspects of the cognitive schema; such a hierarchy eventually converges into the selection of the representation of definite actions. This results in a hierarchical parcellation of the mental episode. A consequence of this is the temporal organisation of mind and behaviour, since representations at different levels organise mental episodes of varying temporal abstraction and duration. In other words various representations carve out discrete mental episodes.

An organised morning may be parsed into episodes of toiletries, getting dressed, preparing and eating breakfast etc., based on the different goal

representations organising these episodes. Each of such episodes may be further chunked into smaller episodes. Preparing breakfast, in the earlier example, may be parsed into subtasks of 'preparing coffee' and 'preparing buttered toast' (Fig. 1.7). Due to such an organisation, task-related representations occur in the context of higher over-arching representations. For example, in the case of hypothetical breakfast preparation, the representation of lower level actions (e.g. 'add coffee' or 'toast bread') occurs in the context provided by the higher level representations ('prepare coffee' or 'prepare toast').

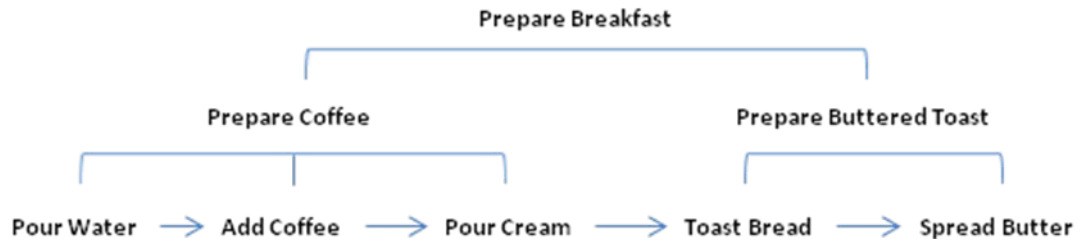

Figure 1.7. Hypothetical organisation of an episode of preparing breakfast**.**

The role of hierarchical representation was demonstrated in an experiment by Schneider and Logan (2007), wherein they asked subjects to memorise two task sequences (called *alpha* and *beta*). These sequences were built of various combinations of two task units – making a judgement about the *size* of the stimulus word (task A), or a judgement on its meaning *semantic* (task B), e.g. alpha = AABB; beta = ABBA. Subjects then performed trials on which they were randomly cued to perform a task at one of the serial positions in a sequence (e.g. 'alpha 2' meant the second task in the alpha list which is size). They found several effects that were consistent with a hierarchical representation of these two lists. One such effect was a large sequence repetition benefit, i.e. subjects were faster if they were cued on the same sequence consecutively than when they were cued on different sequences on consecutive trials. This was the case even when the actual task done in the former condition was different (e.g. beta 1 followed by beta 2), while that done in the latter

condition was the same (alpha 2 followed by beta 1). In fact task repetition benefits only occurred when the tasks were cued from the same list.

In a different study (Schneider and Logan, 2006), the same authors had subjects memorise such task sequences and then execute them repeatedly. For example, if the subjects were executing a sequence ABBA, they performed the tasks cued by the elements of this sequence sequentially on trials 1 to 4, then repeated the whole sequence on trials 5 to 8 and so forth. They found that across all kinds of sequences executed, the RT on trials wherein they executed the first element (i.e. trials 1, 5, 9, 13, etc.) of the sequence was the highest (Fig. 1.8). Two features of this increased RT are noteworthy. In case of sequence ABBA, the first position involves a task repeat, i.e. A to A; whereas positions 2 and 4 involve a task switch, i.e. A to B. If the organisation was not hierarchical, RTs corresponding to the sequence positions 2 and 4 would have been higher than 1, but in the actual RT profile, RT at position 1 was always the highest.

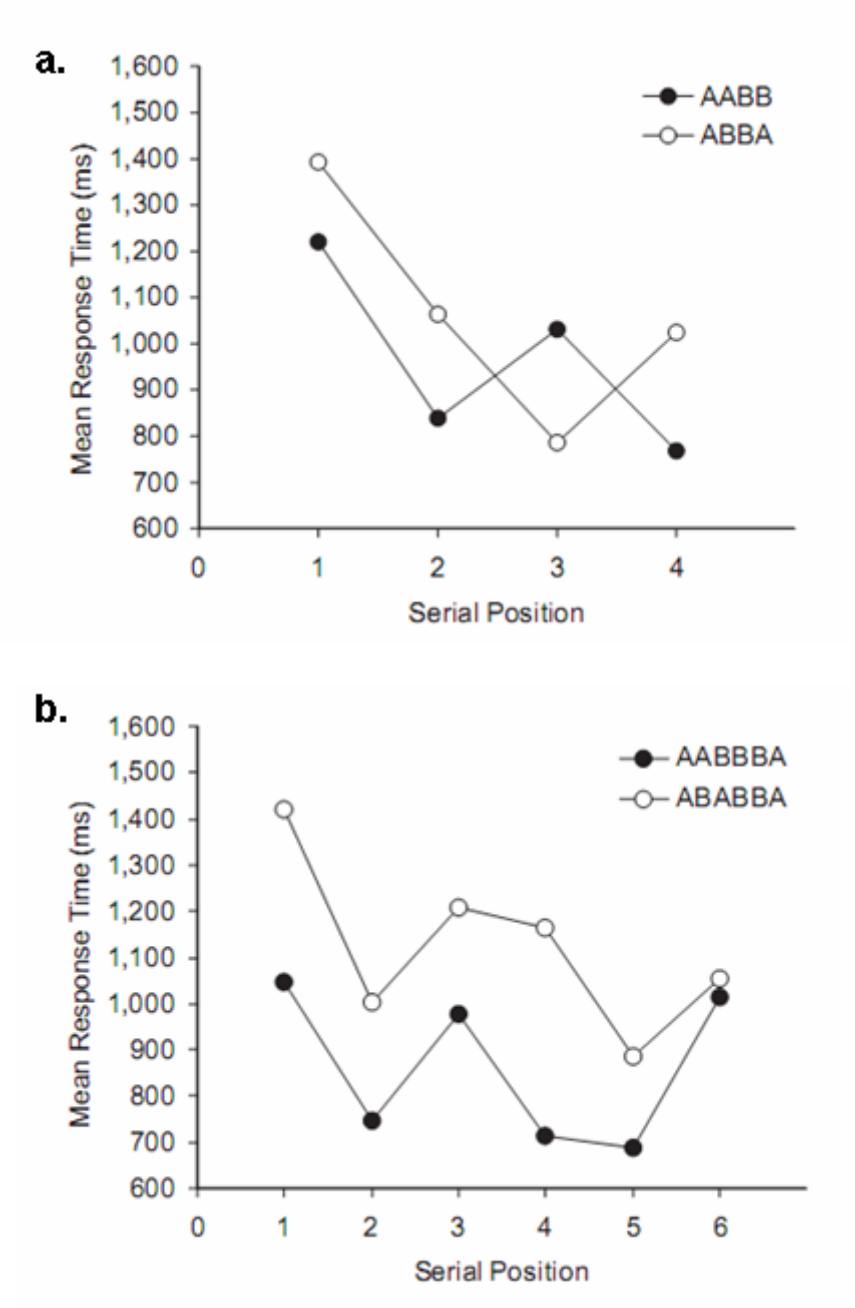

Figure 1.8. Task elements at the beginning of task sequences have the longest reaction time, which depends upon the complexity of the sequence (see text for details). Reproduced from Schneider and Logan (2006).

Secondly, this increase in RT at the first sequence position was not affected by the task at this position being a repetition or a switch with respect to the preceding task. For example, if the sequence was AABB, the first position in the sequence would involve a switch with respect to the preceding task at the fourth position in

the sequence i.e. from B to A : AABB-AABB; whereas if it was ABBA, the first position would involve a task repeat i.e. from A to A: ABBA-ABBA. Nonetheless, the first RT of sequence ABBA was higher than AABB (Fig. 1.8a), plausibly because the former is a more complicated sequence since it involves two task switches, while the latter involves just one. This notion was proven by the findings that the magnitude of the first RT was indeed dependent upon the number of switches within the sequence; hence first RT for a sequence like AABBBA was smaller than that for AABABB (Fig. 1.8b). While these results certainly reinforce the inference of a hierarchical organisation of the resulting task episode, they also suggest that the beginning of the task episode involves instantiation of representations related to the entire length of the episode (and hence the dependence of the first RT on the complexity of the episode).

As has been reviewed in the previous sections, the neural representation of task-related events has been worked upon. However most of these studies have largely ignored the fact that task-related events occur in the context of organised mental episodes. Take the above example of preparing breakfast. An action like 'pour water' can be looked upon as a task event on its own, or this can be looked upon as an event that occurs in the larger context of preparing breakfast. It is, therefore, not known how the larger context of an organised mental episode in which such task related events occur affects the representation of these events. How does the fact that there are underlying representations ('prepare breakfast', 'prepare coffee', 'prepare toast') that have organised the context of lower level sequence of actions (pour water → … → spread butter), affect the representation of the various actions composing the sequence?

An important consequence of the organisation of behaviour into discrete episodes is that specific mental/behavioural events mark the beginning and the end of episodes. The beginning would mark the creation of the episode specific cognitive

focus and the assembly of representations creating this focus, while at the end of episodes, these representations would be dismantled. In the earlier example, 'pour cream' marks the end of the episode of 'preparing coffee', while 'spread butter' completes the preparation of 'buttered toast'. As is evident from Fig. 1.7, in such a scheme, 'spread butter' additionally completes the whole episode of 'prepare breakfast'. This brings forth an important characteristic of organised behaviour. Because the parcellation of a behavioural episode is hierarchical and nested, some events would complete a smaller and hierarchically lower episode (e.g. *subtask*), while others would complete a hierarchically higher and temporally longer episode (e.g. *task*).

1.2.1 Neurocognitive Dynamics at Mental Episode Boundaries

It was seen earlier that fronto-parietal regions flexibly represent all kinds of task-relevant events. In a sequential task, the successive mental states can be expected to be represented in these regions. Old representations can be expected to dissolve away while new ones would be successively assembled. This issue was investigated by Sigala et al. (2008) using an experimental task described earlier (Fig. 1.4a, Cued target task), wherein monkeys watched a series of pictures presented on one of the two sides (left/right) of the fixation spot, and waited covertly for the target specified by the cue at the beginning of the trial. At the appearance of the target, they waited for its offset and then made a saccade to the side on which target had appeared. Three possible cues were associated with three different targets.

The authors compared the activity profiles of 324 neurons recorded from lateral PFC during eighteen possible events - three cues presented to right and left (i.e. six cue types); delay between two stimuli for the trials beginning with the six cue types; and the six targets associated with the six cue types. Mean firing rate of each neuron was obtained in response to these eighteen events; this was normalised for each neuron by dividing the rate of response to an event by mean activity across all

event types. Thus, a pattern of activity was obtained for each event across the neural population sample in the form of a vector. These eighteen vectors were then correlated.

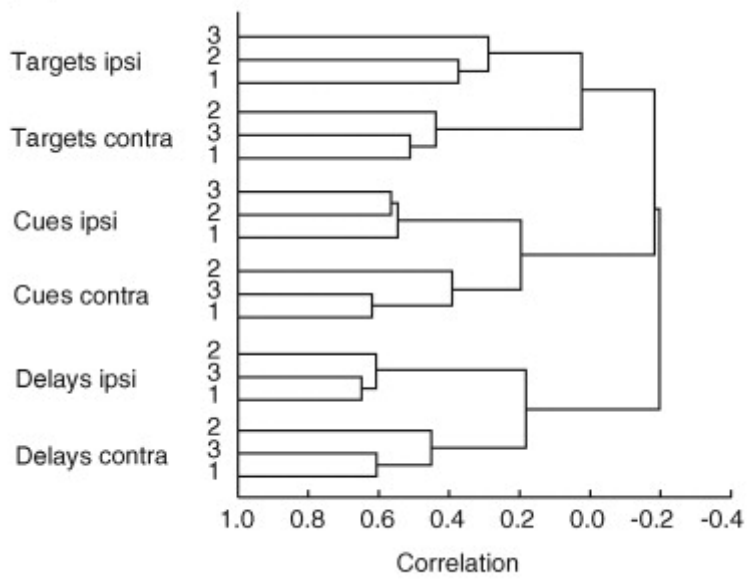

Figure 1.9. Correlation patterns in Sigala et al., (2008). Note that the correlation between the different types of events at the same phase (e.g. between the three cues presented in the same hemifield) are high, while correlation across phases (e.g. between the cues and delay) is minimal. Reproduced from Sigala et al., 2008.

The results (Fig. 1.9) showed high correlations for activity patterns from the same task phase, particularly in the same hemifield, and negligible or even negative correlations for different task phases. This shows that the different steps of a sequential task are represented by different patterns of activity and hence form different neural representations. The structure of neural representation seems to care less about the content of representation at a particular step of the sequence – seen in the high correlations in the activity pattern across the different cue – target identities but at the same phase, e.g. the three targets in the same hemifield had highly correlated patterns. However, for the different phases of the task, the representation was essentially orthogonal, i.e. activity patterns for targets, cues and delays were minimally correlated with each other.

What would happen at the boundaries of hierarchically higher task episodes? Plausibly, such boundaries represent greater revisions in mental representations. In the case of Fig. 1.7, finishing the act of 'pour water' brings the subject to the next action 'add coffee', while the higher level representations of 'prepare coffee' and 'prepare breakfast' under which this action took place are not affected. In contrast finishing the act of 'pour cream' not only changes the lower level action representation but also changes the higher level representation of 'prepare coffee'. At the extreme, the event completing the whole task of 'preparing breakfast' changes the representations at all levels of hierarchy into which the task episode was organised.

Arguably, brain regions representing such organising representations (like 'prepare coffee' or 'prepare breakfast') can be expected to show a change in activity when these are revised at the end of mental episodes. Further, the activity can be expected to be greater when the revision is more intense, e.g. at the end of the overall episode ('prepare breakfast') versus completion of a sub-episode ('prepare coffee'). Note that earlier studies focussing on task-related representations have usually focussed on task-related stimulus, action, rule or other related behavioural events like reward. To my knowledge none of these have explicitly focussed on the implicit representations organising the task episode. As has been mentioned earlier, such representations provide the organising framework that controls the lower level behavioural events like action and related decisions.

1.2.2 Task Phases

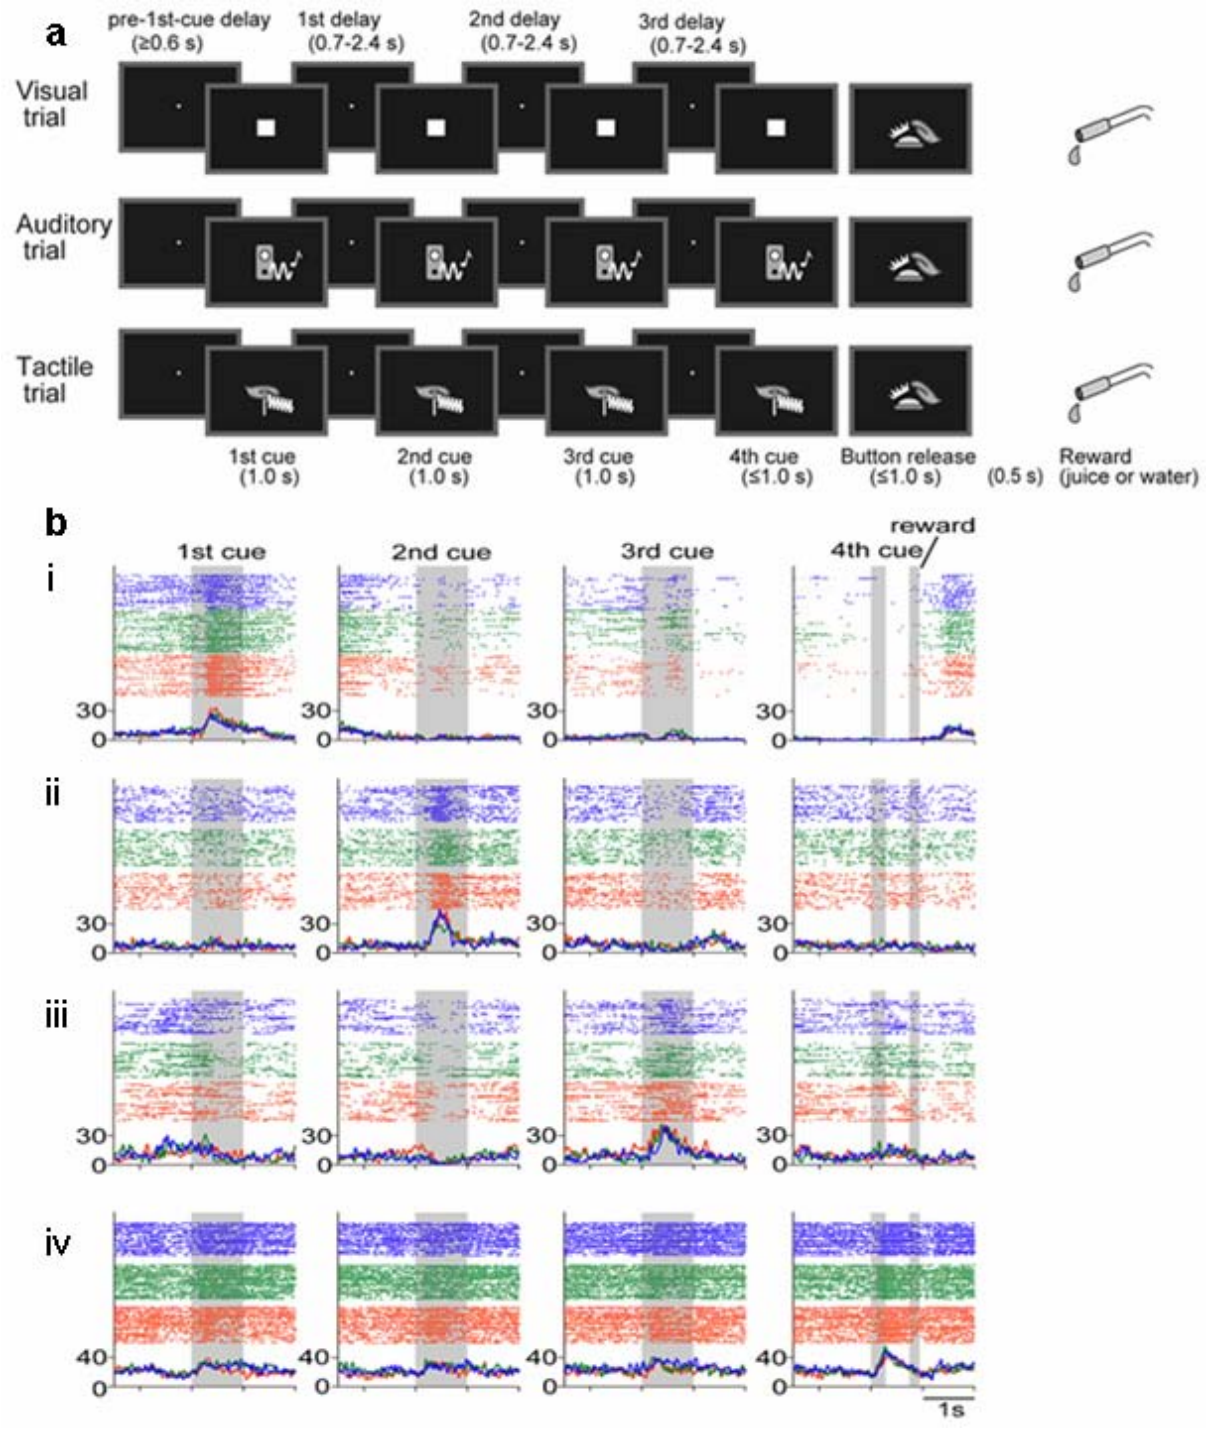

Figure 1.10. (a) Task used by Saga et al., (2011); Monkeys were presented with four sequential cues interrupted with blank delays. After the fourth cue, monkeys released a lever to get a reward. (b) Examples of neurons specific for cues belonging to particular phases. Reproduced from Saga et al., 2011.

Another consequence of organisation of mental/behavioural events into discrete episodes is that the representations at the various ordinal positions assume distinctness because of their rank position. Had the behaviour been a hierarchy-less flat chain of events, rank position would have no meaning and any step would be equivalent to any other. A number of experiments show that fronto-parietal neurons code for the rank position of the behavioural events (Hasegawa et al, 2004; Bardyyeva and Olson, 2010; Campos et al., 2010; Saga et al., 2011). In an experiment by Saga et al (2011), monkeys were to wait for four sequential cues (separated by intervening delays) before releasing a lever and gaining a reward (Fig. 1.10a). The four cues on any trial were identical and could be presented in one of three sensory modalities (visual, auditory and tactile). The sequential cues, therefore, can be regarded as sequential phases that differ only in terms of their position in the sequence. They found that more than 60% of the lateral prefrontal neurons, sensitive to task events, coded information about the phase of the trial. Such neurons, however, varied in the actual detail of this coding, for example, some were specific to the cue at a particular sequential position (Fig. 1.10b) while others coded for specific delay periods between the cues. Another group of neurons coded for a combination of sequential cues and delays.

While representation of the sequential steps of a multi-step task is qualitatively distinct (Sigala et al., 2008; Saga et al., 2011), does the representation of the various steps differ quantitatively as well? Phenomenologically, the various steps do seem to be different. The final steps appear to be more salient and valued than earlier ones. Subjective frustration experienced is greater when subjects commit a mistake closer to the completion of the task (Louro et al., 2007). Likewise, neural activity gets higher as subjects get closer to the reward in a multistep task (Rowe et al., 2008; Sohn and Lee, 2006; Mizuhiki et al., 2007; Shidara et al., 2005, Platt and Glimcher 1999). However, these studies do not make it clear if the effect is related to

the ordinal position of the step with respect to the completion or is due to the proximity and imminence of reward.

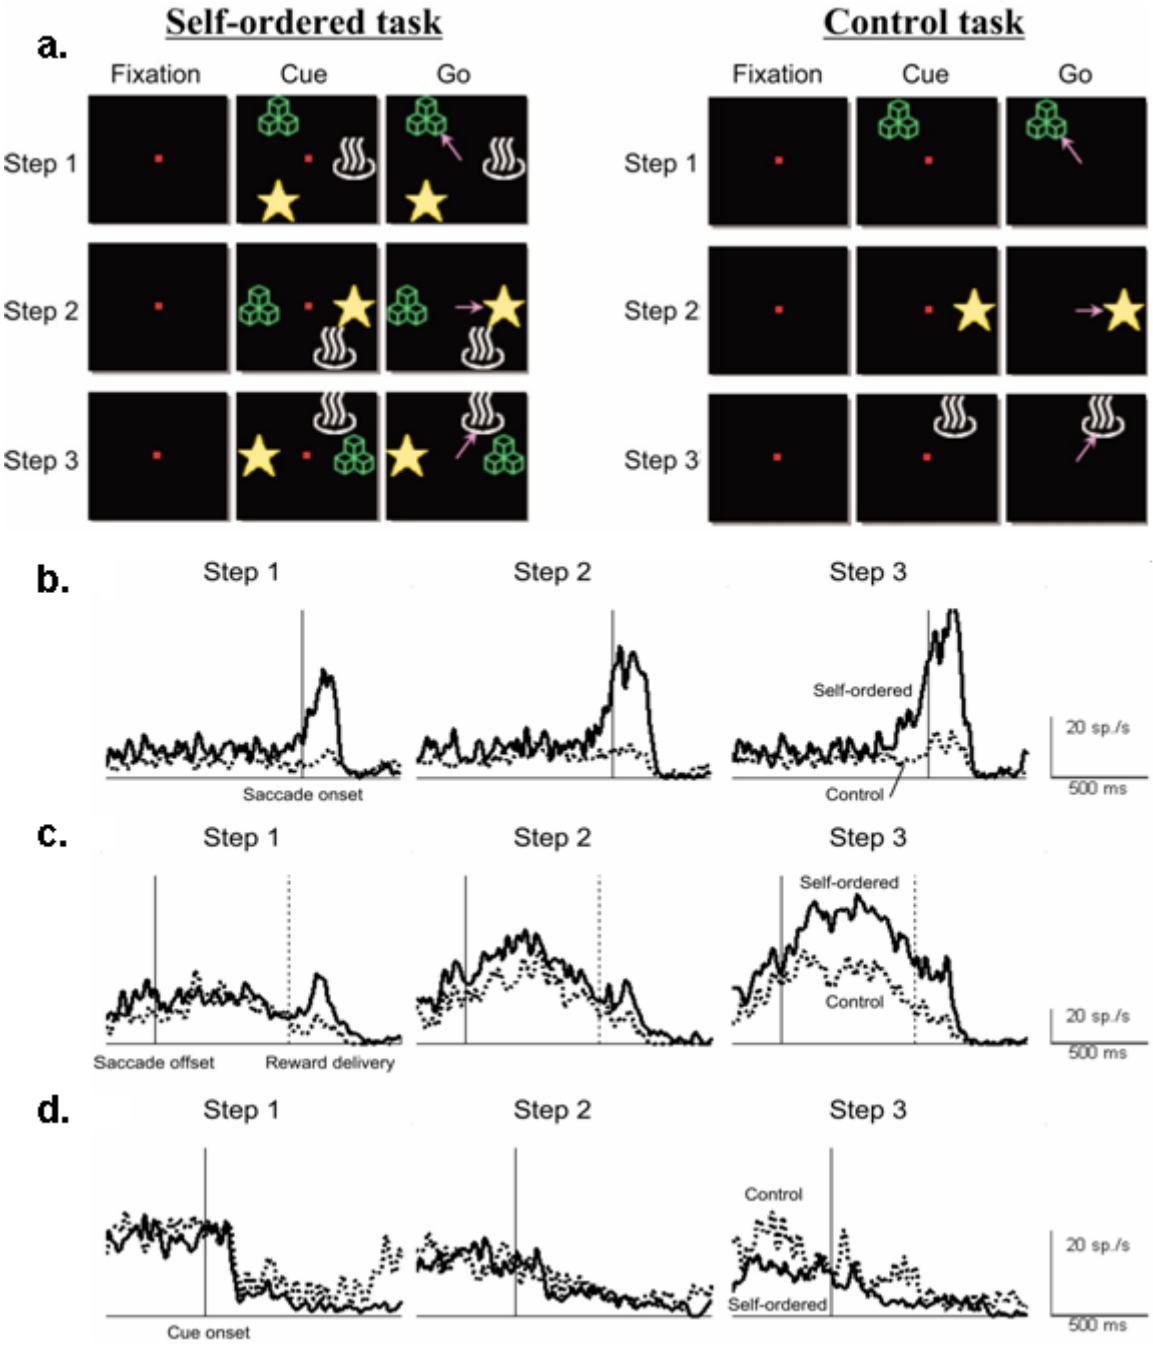

Figure 1.11. (a) The two tasks used by Hasegawa et al. (2004). In the self ordered condition, monkeys performed three steps. In the first step, they were presented with the three stimuli, and were explicitly signalled to choose one of them by making a saccade. In the next two steps, the same three stimuli were presented but at different locations, and monkeys chose a stimulus that had as yet not been chosen by them. In the control task, they were presented with one stimulus at each step to which they made a saccade. In both tasks,

monkeys were rewarded after each step. (b) to (d) Samples of neuronal activity profiles showing effect of task phase at the different epochs of the step. Reproduced from Hasegawa et al., 2004

A study by Hasegawa et al., (2004) showed a similar increase in activity across the sequential steps of a task even when the reward expectancy had been controlled for. Monkeys performed two distinct tasks (Fig. 1.11). In the self-ordered condition, the monkeys fixated at the centre of an array of three pictures, chosen from a set of 6 pictures. On the first step they made a saccade to one of these three pictures. On the 2nd and 3rd steps (*steps 2* and *3*), the same 3 objects reappeared at different spatial positions, and the monkeys had to make saccades to objects that had not been chosen on previous steps. Monkeys received rewards after every step, which increased with the position of the step in the ratio of 1:2:4. To control for the effect of reward, the monkeys did a control task in which they made simple visually guided saccades in three similar steps without the requirements of identification or memory. The reward schedule was identical across the two kinds of tasks.

These two tasks differ in two critical aspects. The three steps of the self-ordered condition have a means-end relation and so form the three phases of the same task. On the contrary, in the control condition the individual steps were independent of each other, and so are less strongly cohered as belonging to the same task. Secondly, the working memory load increased across the three steps of the self ordered condition since the monkeys had to remember the pictures they had already chosen; there was no working memory load in the control condition.

The results (Fig. 1.11, b to d) showed that activity in half of the recorded neurons was modulated with the step position in the self-ordered condition (dark lines in Fig 1.11, b to d). This effect was weak or absent in the control condition (dotted lines), which excluded the possibility that reward was the main cause of modulation. While the working memory load did increase across the steps of the self-ordered condition, many aspects of their results suggest that working memory

load was not the cause behind the modulation of activity. Many neurons showed step modulation of activity in the post saccadic epoch (Fig. 1.11b). Note that in step 3, monkeys can relinquish their memory load after the saccade has been made and hence the post saccadic epoch of step 3 need not have any cognitive load at all. However, as is evident in Fig. 1.11b, neurons do show a modulation of activity across the post saccadic phase of the three steps. Secondly, since the working memory load is constant across the various epochs of the step, the neural effect of load should be constant across all epochs of the step. However, most neurons showed the effect of the step position exclusively at certain epochs of the step – the neurons exemplified in Fig 1.11b, show the effect exclusively at the time of the saccade, while those in Fig 1.11d, showed a decrease in activity across successive steps and this decrease was limited to their activity before the onset of the cue.

1.3 Précis

In this thesis several effects of task organisation on the representation of behavioural events are investigated. In chapter 2, two fMRI experiments are described that investigate the fronto-parietal activity elicited by the completion of task episodes at different levels of hierarchy (subtasks versus tasks). The experiments show that the completion of defined task episodes results in activity in a number of fronto-parietal regions, but especially in anterior prefrontal and anterior cingulate cortices. It is discussed that this activity cannot be attributed to the recall of new task rules, termination of cognitive activity, etc.; rather, they are solely related to change in the organising representations underlying the task episode.

The third chapter looks at the representation of the different phases of a task episode. The fronto-parietal activities elicited by early and late phases of the trial were investigated. The design further manipulated the cognitive load of the two phases independently. This allowed for a direct comparison of the effect of phase vis-à-vis the effect of cognitive load. The results showed that fronto-parietal regions that

increase their activity in response to cognitive load, also increase their activity for the later phases of the task episode.

The fourth chapter investigates if the characteristics of the higher-level representations have a causal role in determining the structure of the organisation of the task episode. Results show this to be true. They also confirm the results of earlier chapters in a different task framework.

The fifth chapter looks at the pattern of activity across the voxels of a region elicited by the different steps of a sequential task, to investigate if the discriminability of these patterns, and by extension, the content of the information in the region, depends upon the sequential position of the step. The results show that the representation of varied kinds of information about the step is affected by their ordinal position, all of them being higher for steps closer to the completion.

Finally, the sixth chapter discusses the various neural possibilities underlying these effects.

CHAPTER 2

Completion of Hierarchical Task Episodes

2.1 Introduction

A central feature of purposive behaviour is parcellation of a main task episode (e.g. preparing breakfast) into smaller subtasks (preparing coffee and buttered toast), whose achievement in their proper sequence culminates in achievement of the main purpose. The whole sequence of behavioural events (e.g. pour water → add coffee → add cream → toast bread → spread butter) may be organised into a nested hierarchy of tasks and subtasks, corresponding to the abstractedness of purpose being maintained. The representations underpinning the subtask would be nested within those underpinning the task episode.

Attempts at modelling cognitive operations underlying organised behaviour find such hierarchy of representations to be critical (Cooper and Shallice, 2000). Without higher level representations organising the events at the lower level, the creation of correct sequences tends to be slow, since lot of options need to considered before the correct one is selected, and inflexible, as minor changes can upset the process (Russell and Norvig, 1995). Further such systems find it difficult to follow long term goal or to follow multiple goals (Gat, 1998).

In hierarchically organised episodes, discrete behavioural events would correspond to the creation, revision and dismantling of representations at different levels of hierarchy. In the current example, the event ‘add cream’ completes the subtask of preparing coffee and hence marks the revision of subtask level organising representations, while the event ‘spread butter’ completes the whole task of preparing breakfast and changes the representations up till the highest level.

Correspondingly, the different behavioural events making up the sequence would differ in terms of the change in mental representations that takes place at them – with greater changes occurring at the behavioural events marking the end of hierarchically abstract and temporally extensive episodes.

A key role in purposive behaviour is played by prefrontal cortex (Luria, 1966; Duncan, 1986). In the behaving monkey, cells of lateral and medial prefrontal cortex code many kinds of information relevant to a current cognitive operation, including stimuli, responses, rules, working memory contents etc. (Duncan, 2001; Miller and Cohen, 2001; Procyk et al., 2000). In functional magnetic resonance imaging (fMRI) studies of the human brain, using either adaptation or multivoxel pattern analysis to examine the detailed content of task representations, extensive coding of task-relevant information is seen across multiple regions of frontal and parietal cortex (e.g. Li et al., 2007; Thompson and Duncan, 2009; Woolgar et al., 2011). It is frequently proposed that adaptive fronto-parietal representations are a source of cognitive control, biasing processing in other brain regions (Desimone and Duncan, 1995; Miller and Cohen, 2001; Rigotti et al., 2010). In such a framework, completion of each task step signals that a previous control representation should be dissolved and the next put into place (Sigala et al., 2008). In a hierarchical task that requires control at numerous levels, task steps crossing across hierarchical episodes would result in revision of representations at multiple levels .

We know little of how fronto-parietal control activity is influenced by hierarchical task-subtask structure. While several studies have linked hierarchical task control to the organisation of prefrontal cortex, especially along the antero-posterior axis, these studies have not been explicitly concerned with task-subtask structure (Koechlin, 2003; Badre, 2007, Christoff, 2006). For example, Badre (2007) showed more posterior frontal activity for simple one-level decisions (e.g. red > left key) than for two-level decisions in which one stimulus feature indicated how the

other should be processed. Similarly, Koechlin (2003) showed anterior frontal activity when the context of a whole task block determines how a stimulus should be interpreted. Though control in such tasks is certainly hierarchical, this sense of hierarchy is different from a comparison of task and subtask completion in complex, multi–step behaviour.

To investigate this issue, we devised a sequential target detection task with a hierarchical structure, such that some target detections represent the achievement of subgoals, others the achievement of a higher level goal. We reasoned that, when a subtask is completed, its specific content loses relevance, but higher level representations pertaining to the overarching task episode must remain in behavioural control. By comparison, completion of a more substantial behavioural episode requires a more substantial revision of task related representations, and perhaps producing a stronger or more extensive pattern of fronto-parietal activity.

2.2 Materials and Methods

2.2.1 Tasks

Experiment 1

Two experiments were conducted. In the main experiment (Experiment 1), participants monitored a series of letters presented at a rate of 1/1.3s on a computer screen (Fig. 2.1). The task was to detect four target letters in turn. A three - letter word (e.g. 'CAT') presented at the start of the trial indicated the first three targets (e.g. Fig. 2.1, T1 'C', T2 'A', T3 'T'). The fourth target was always the letter X. At the end of the 52 s letter stream, the participants were probed to indicate whether X had appeared (true on 50% of trials called complete trials, remaining trials referred to as incomplete trials).

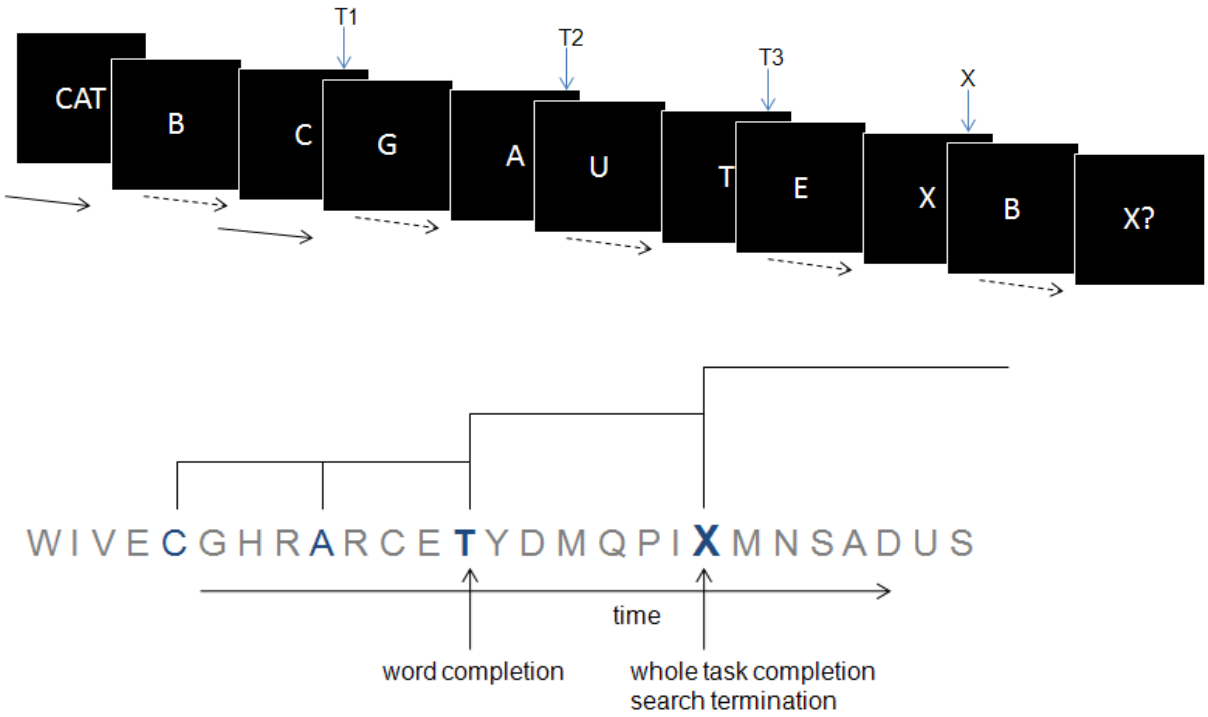

Figure 2.1. Structure of a typical trial (Experiment 1). Trials began with a 3-letter cue word. The three letters of this word were to be covertly detected, in the correct order, in the ensuing letter stream; after all three had been detected, search for the letter 'X' began. The complete sequence of four target letters appeared in only half of the trials. The letter stream ended with a probe asking whether the letter 'X' had appeared in the correct sequential position, i.e. following the three letters of the target word. Thus the first two targets (T1 and T2) changed the representations at the lowest level (component letters of the first target word; level 1); the third target, T3, completed the representations at the next highest level (complete target word; level 2), while the fourth target, X, completed the whole goal of the task, and changed the representation at the highest level (level 3).

At the start of each trial, the instruction screen specifying the first three targets (Fig. 2.1, 'CAT') was presented for 3500 ms. The letter sequence began after a jittered gap of 1000 – 5000 ms, each letter being presented for 800 ms with a gap of 500 ms in between consecutive letters. The letter stream consisted of a total of 40 letter presentations and lasted 52 s. Participants began by monitoring for T1 (Fig. 2.1, 'C'), at the detection of which search for T2 started, and so on. Targets were only

relevant once the appropriate preceding targets had already been detected; for example, while Xs might appear at any point in the sequence, only an X appearing after sequential detection of T1, T2 and T3 was relevant. In this way participants were obliged to search for the specified four targets in turn, with a positive response at trial end only if all four targets had appeared in the correct order. Within each trial, the inter-target interval varied randomly between 1.3 and 32 s. Responses were made when the letter stream terminated with a probe ('X?' displayed for 2000 ms) asking whether the letter 'X' had appeared at the relevant position i.e. after the detection of T1-T3. Responses were made on a button box positioned under the participant's right hand (index finger for 'yes', middle finger for 'no'). A variable inter-trial interval of 2000 to 7000 ms preceded onset of the cue for the next trial. (See Appendix B.2 for exact instructions).

All stimuli were centred on the screen, visible from the participant's position in the scanner via a mirror mounted within the head coil. Letters subtended a visual angle of 2° vertically. The experiment was controlled by a program written in Visual Basic. Subjects learnt the task in a ten minute pre-scan practice session and proceeded to a scanning session of an hour, which was divided into three separate scanning runs, each consisting of 20 trials.

Task-subtask structure was emphasised in the initial instructions that mentioned "two tasks are to be done on each trial; search the letters of the cued word, and then search for X". This was reinforced by the distinction of the sources that guided these two tasks - a cue-word to be kept in working memory as it changed from trial to trial, and the letter 'X' that was constant across trials. According to such a structure, T1 and T2 completed subtasks at the lowest level (component letters of the initial target word; level 1), T3 completed a task at the next higher level (target word completed; level 2), while 'X' completed the entire task episode (level 3). Participants were informed that correct 'yes' responses on complete trials increased

the current score by +1, whereas correct 'no' response on an incomplete trials did nothing to the score; an incorrect response on trial decreased the current score by -1.

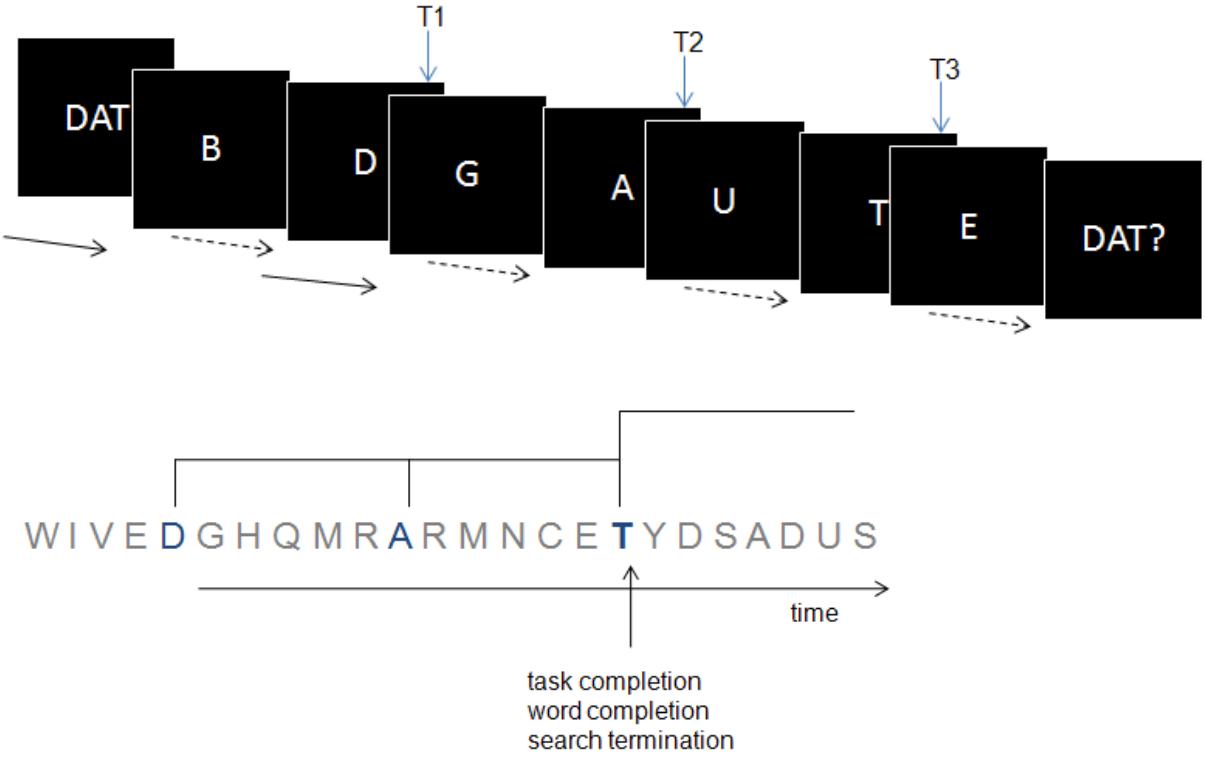

Figure 2.2. Structure of a typical trial in Experiment 2, similar to Experiment 1 but with only three targets T1-T3.

Experiment 2

In Experiment 1, when each of targets T1 to T3 was detected, the previous search was to be abandoned and a new one began. In contrast, detection of the final target (X) was associated with termination of the search and an increase in the current score, which could be interpreted as a form of reward. To examine the significance of these features, in a follow-up experiment (Experiment 2) the task was slightly modified - T3 detection completed the task, i.e. there was no requirement to

search for a final 'X' (Fig. 2.2). Visual search now terminated at T3 rather than X; further an explicit reward (of 10p) was given at the completion of the task. Comparison between experiments allowed a direct examination of the importance of these factors. At the end of each trial participants were probed to indicate whether all the three target letters had appeared in the correct order (Fig. 2.2, 'DAT?'); as before this was the case in only half the trials, with 25%, 18% and 6% of the trials having two, one and no target letters respectively. In Experiment 2 each letter stream was composed of just 40 letters, lasting a total of 40 s. The scoring system was same as before, but with an explicit addition of monetary reward.

An inherent problem in experimental designs focussing on sequential events is that the latter events will necessarily occur after the earlier ones and hence there is a possibility of such event regressors being correlated. This, however, was obviated in the design of the current experiments by having a large temporal jitter between the different target events (1.3 to 32 s). In the second experiment an additional feature was added. Here the number of target events appearing in a trial varied from none to all, thus allowing for greater dissociation between them as the earlier target events were not always followed by the later target events. Finally, a distinct set of FIR regressors modelled the entire duration of the letter stream (see below) and ensured that effects related to the temporal positions within the letter stream were regressed out and did not affect the estimates of target related activity. Lastly, the effect of expectancy was obviated by ensuring that the probability of any instance of the letter stream being a target was constant across the entire sequence of the letter stream.

2.2.2 Participants

18 participants (10 female, mean age = 22.5 ± 3.6 years) in the first experiment and 21 (15 females, mean age 24.5 ± 4.1years) in the second experiment were recruited from the MRC-CBU volunteer panel. Participants were right handed

and had normal or corrected vision. Informed consent was taken and the participants were reimbursed for their time. The study had the approval of Hertfordshire Local Research Ethics Committee.

2.2.3 Acquisition

fMRI data were acquired using a Siemens 3T TimTrio scanner with a 12 channel head coil. A sequential descending T2*-weighted echo planar imaging (EPI) acquisition sequence was used with the following parameters: Acquisition time 2000 ms; echo time 30 ms; 32 oblique slices with slice thickness of 3 mm and a 0.75 mm interslice gap; inplane resolution 3.0x3.0 mm; matrix 64x64; field of view 192mm; flip angle 78 deg. T1-weighted MPRAGE structural images were also acquired for all participants (slice thickness 1.0 mm, resolution 1.0x1.0x1.5 mm, field of view 256 mm, 160 slices).

2.2.4 Analysis

The fMRI data were analysed using SPM5 (Wellcome Department of Imaging Neuroscience, London, England; www.fil.ion.ucl.ac.uk). Prior to statistical analysis, all EPI volumes were slice-time corrected using the first slice as a reference, and then realigned into a standard orientation using the first volume as a reference. These realigned were then normalised into the Montreal Neurological Institute (MNI) space, and spatially smoothed using an 8 mm full-width half-maximum (FWHM) Gaussian kernel. During the normalisation stage, voxels were resampled to a size of 3 × 3 × 3 mm. The time course of each voxel was high pass filtered with a cutoff period of 90 s.

Statistical analysis was carried out using a general linear model. To capture activity related to target detection, a 16 sec epoch starting from each target onset was modelled using an FIR basis set of eight 2 s long boxcar regressors. In this way the response to target detections could be modelled without making any

assumptions about the shape of the BOLD response. To capture and regress out the possible effects of position of target events in the letter stream, the entire letter stream was modelled using another FIR set of 25 2 s long boxcar regressors (in experiment 2, the duration of letter stream required only 20 such regressors, since the letter stream were shorter in this experiment). Additionally, the cue and probe were modelled using epoch regressors, of width equal to the duration of respective events, convolved with a basis function representing the canonical hemodynamic response. Movement parameters and block means were included as covariates of no interest. Parameter estimates for each regressor were calculated from the least squares fit of the model to the data, and estimates for individual participants were entered into a random effects group analyses.

Whole-brain comparisons were performed using paired t-tests on the relevant contrast values from each participant's first-level analysis. Unless otherwise specified, all results are reported at a threshold of $p < .01$, corrected for multiple comparisons using the false discovery rate (FDR). Coordinates for peak activation are reported using an MNI template.

To capture fronto-parietal regions widely engaged in cognitive control, ten regions of interest (ROIs) were created as spheres of 10 mm radius at coordinates (Table 2.1) that have been shown to be consistently active in varied tasks (Duncan, 2006; Dosenbach et al., 2006). The ROIs (in MNI space) were bilateral inferior frontal sulcus (IFS; central coordinate ±41 23 29), bilateral intra-parietal sulcus (IPS; ±37 -56 41), bilateral anterior insula extending into frontal operculum (AI/FO; ±35 18 3), anterior cingulate (ACC; 0 31 24), and pre-supplementary motor area (pre-SMA; 0 18 50), all taken from Duncan (2006); along with bilateral anterior prefrontal cortex (APFC; 27 50 23 and -28 51 15) taken from Dosenbach (2006). ROIs were constructed using the MarsBar toolbox for SPM (http://marsbar.sourceforge.net; Brett, Johnsrude & Owen, 2002). Estimated data were averaged across voxels within each

ROI using the MarsBar toolbox and the mean values were exported for analysis using SPSS.

ROI	Coordinates
IFS	±41 23 29
IPS	±37 -56 41
AI/FO	±35 18 3
APFC R L	27 50 23 -28 51 15
ACC	0 31 24
pre-SMA	0 18 50

Table 2.1 List of ROIs with their coordinates in MNI space. Note that for midline regions (ACC and pre-SMA), a single ROI centred on the midline was used.

2.3 Results

2.3.1 Experiment 1

Behaviour

Average response time to the probe was 723 ± 15 ms, and accuracies for responses to the probe exceeded 97% for most subjects (mean 97.1 ± 1.4), suggesting that they were indeed covertly detecting target letters.

Imaging

ROI analyses

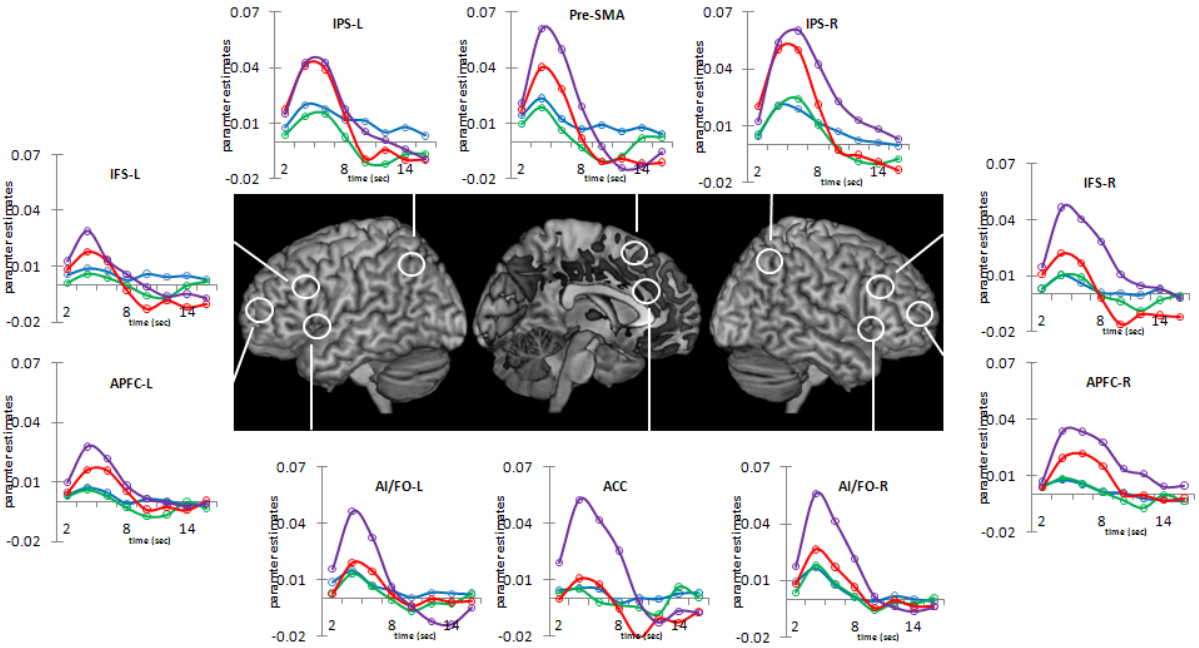

Figure 2.3. (Exp 1) Time course of activity in the different ROIs, constructed from the estimates of eight FIR regressors linked to each target event (blue: T1; green: T2; red: T3; purple: X).

The FIR model provided estimates of activity in eight successive 2-second long windows starting from each target onset, which were plotted to construct an estimate of the time course of activity following each target detection. Fig. 2.3 shows these plotted for the ROIs; the difference across the various target detections is evident. Note that the time courses for T1 and T2 are nearly identical in many ROIs. For further analysis, an index of phasic BOLD response was calculated by subtracting the estimate of first FIR bin from the average of the second and third. Mean values of this index for each ROI appear in Fig. 2.4. It is evident that X, which completes the overall goal (level 3) shows the highest activity in all ROIs. In contrast, T1 and T2 (level 1 subgoals) show the least; indeed in ACC, pre-SMA and APFC, as is evident from the error bars depicting one standard error, their activation index does not differ from zero.

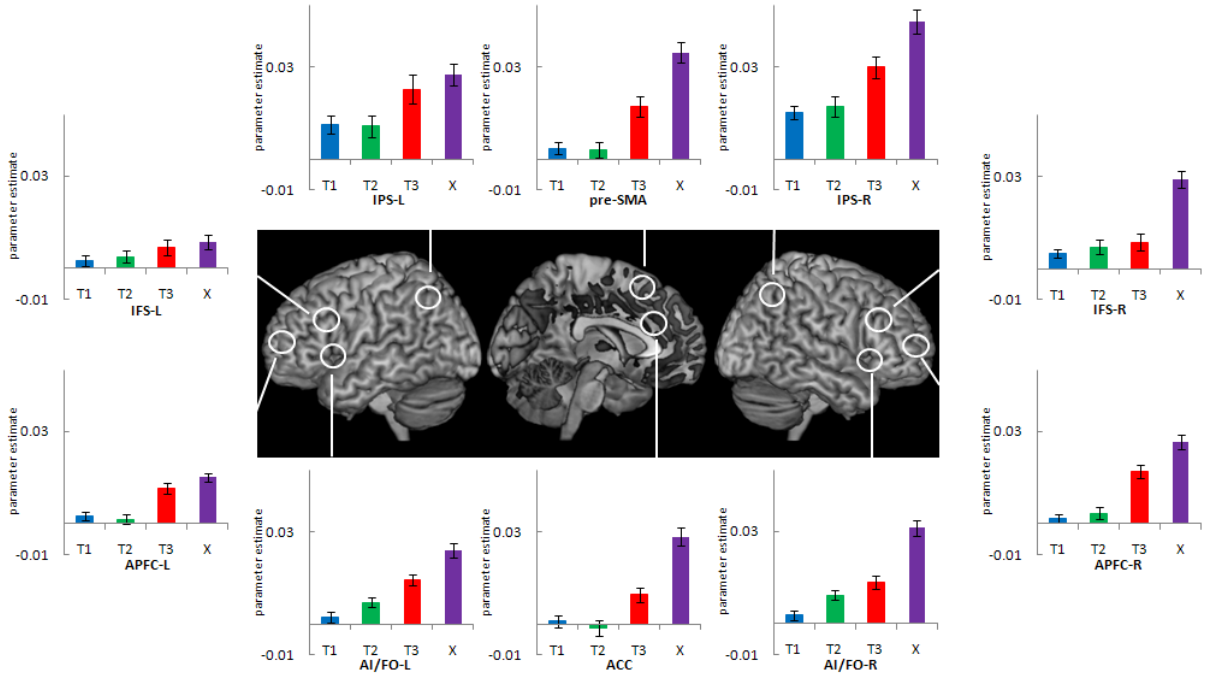

Figure 2.4. (Exp 1) Comparison of phasic activity in response to various target events. An index of phasic activity was derived for each target event by subtracting the estimate of the first FIR regressor from the average of the second and third. The error bars represent one standard error of the mean. All ROIs showed a main effect of target position on the magnitude of this index (Table 2.2a).

(a)

ROI	Target position $F_{(3,51)}$	Hemisphere $F_{(1,17)}$	Target x Hemi $F_{(3,51)}$
IFS	**4.9***	**11.4***	**14.2***
IPS	**6.3***	**9.9***	2.7
AI/FO	**7.1***	2.6	1.4
APFC	**9.1***	**4.9***	**3.4***
ACC	**9.0***	-	-
Pre-SMA	**11.2***	-	-

(b)

ROI	Laterality	T1 & T2 vs T3	T1 & T2 vs X	T3 vs X
IFS	L	1	1.2	0.3
	R	0.67	**4.9***	**3.8***
IPS	L	1.9	**2.6***	0.6
	R	**2.5***	**4.4***	1.8
AI/FO	L	2	**2.9***	1.6
	R	1.2	**3.7 ***	**3.2***
APFC	L	**2.8***	**3.7***	0.8
	R	**3.3***	**4.8***	1.7
ACC		**2.6***	**4.4***	**3.1***
Pre-SMA		2.2	**4.8***	**3.0***

Table 2.2. [a] Effect of the position of target detection, laterality of the region and their interaction in different ROIs. Values in bold were significant at $p < 0.05$. [b] $t_{(17)}$ values from the pairwise comparisons of different target detections at various ROIs. Values in bold were significant at a Holm-Bonferonni corrected threshold of 0.05 for multiple comparisons.

To confirm that activity across the various target type, a repeated measures ANOVA was done with target type (T1, T2, T3, X) and ROIs as factors. As would be expected, there was a significant main effect of target type ($F_{3, 51} = 13.5$, $p < 0.001$). The main effect of ROIs was significant as well ($F_{9, 153} = 9.4$, $p < 0.001$), and so was an interaction between ROIs and target type ($F_{27,459} = 2.0$, $p < 0.01$), showing that the differential response to the target types varied across the ROIs.

For more detailed investigation, activity indices for each ROI were examined by ANOVA, with factors target type and hemisphere (for lateral ROIs only). All ROIs showed a main effect of target type (Table 2.2a). Amongst the bilateral ROIs, IFS, APFC and IPS showed a main effect of laterality reflecting greater activity on the right. In IFS and APFC significant interaction of target type with hemisphere showed that the differential activity was greater in the right hemisphere.

To amplify the significant effect of target type, pairwise comparisons compared T3 (level 2 subgoal) with the average of T1 and T2 (level 1). APFC, ACC and right IPS showed significantly higher activity to T3 (Table 2.2b). Amongst other target events, X was significantly higher than T1 and T2 in all ROIs except left IFS. In all ROIs, the tendency was for greater activity to X than T3, however, this was significant only in right IFS, right AI/FO, ACC and pre-SMA. There was no significant difference between T1 and T2 in any ROI.

Whole Brain Analyses

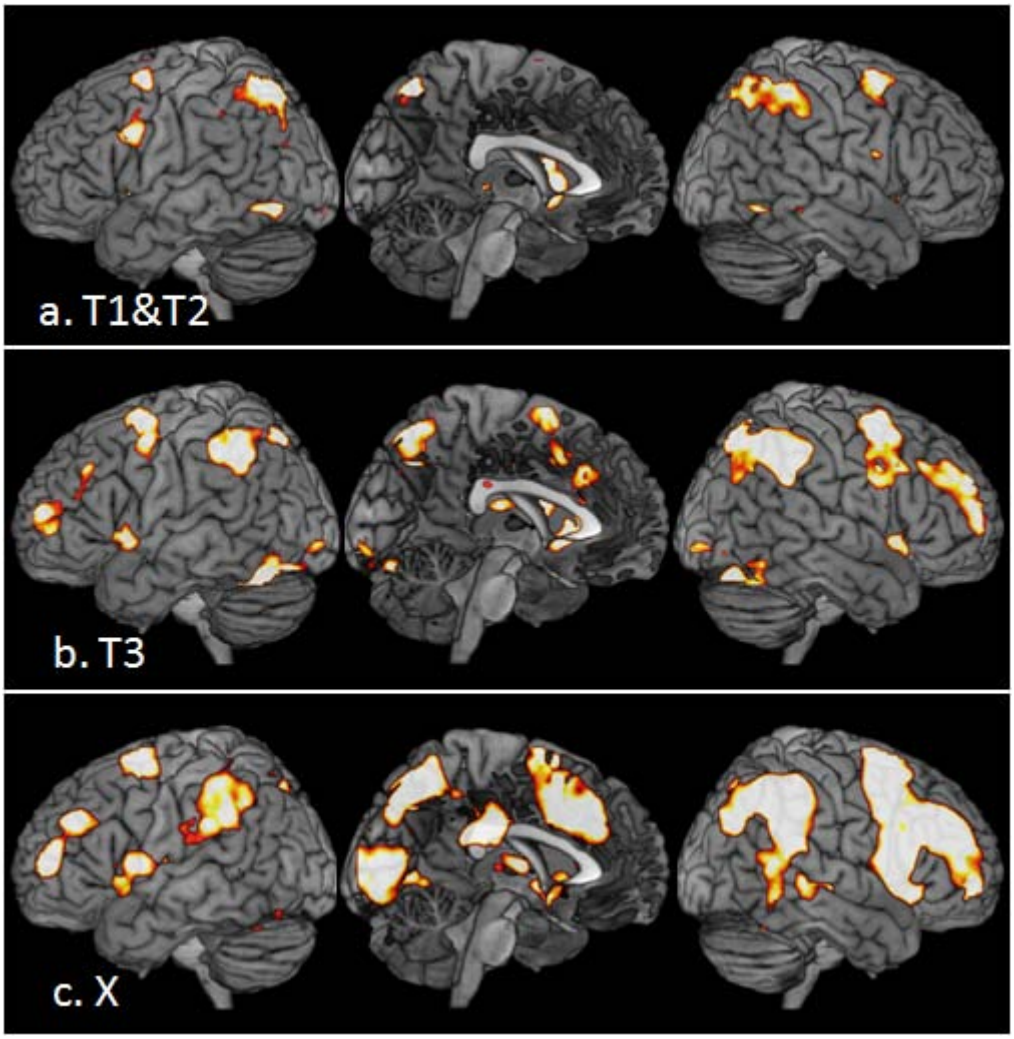

Figure 2.5. (Exp 1) Regions where the activation indices for the various targets were significantly greater than zero. As in the ROIs, X and T3 elicited activity in much wider network of regions than T1 and T2.

In the first series of contrasts, the regions where the activation index for the different target detections differed significantly from zero were investigated. As in the ROIs, the involvement of other regions varied with the position of the targets. T1 and T2 (Fig. 2.5a) did not elicit much activity beyond few foci in posterior frontal (dorsal premotor and intrafrontal junction) and IPS. T3 (Fig. 2.5b) elicited much more widespread and intense activity spreading onto anterior and medial prefrontal

regions, anterior insula, precuneus; activity was most widespread for X (Fig. 2.5c) and additionally involved interparietal lobule, temporo-parietal junction and cuneus.

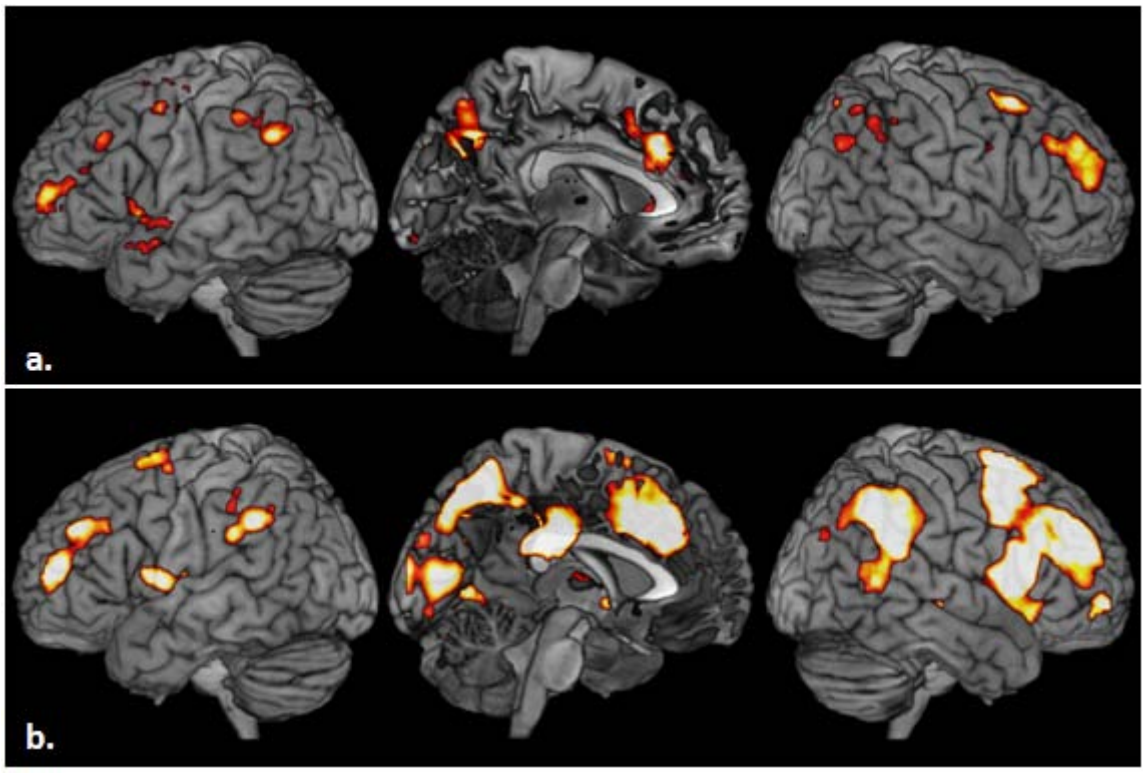

Figure 2.6. (Exp 1) [a] Whole brain render showing areas that had higher activation index for 'T3' compared to 'T1' and 'T2' ; [b], areas with greater activity for 'X' than 'T1' and 'T2'.

To check if these differences were significant, two contrasts compared the activity during T3 and X with that during the earlier target detections. Consonant with the ROI results, the first contrast (Fig. 2.6a) showed that T3 elicited greater activity in APFC, ACC and parts of pre-SMA along with lateral parietal and precuneus (for list of coordinates see Appendix, table A2.1). Results of the second contrast showed greater activity for X in widespread frontal, parietal and occipital regions with a dominance on the right (table A2.2). In the right prefrontal cortex it involved inferior frontal sulcus, frontal operculum and the anterior insula, extending anteriorly through the middle frontal gyrus up to anterior prefrontal region. In the

left prefrontal cortex only the anterior prefrontal regions and rolandic operculum were involved. Medially, stronger activity for X compared to T1/T2 was seen in the anterior cingulate and the pre-SMA on both sides. Posteriorly, significant activity was found in the intra-parietal sulcus extending into the temporo-parietal junctions, precuneus, cuneus and parts of calcarine and lingual regions.

To summarise activity was strongest and most widespread in response to the completion of the final goal (level 3). Selected regions especially APFC and ACC showed greater activity for change in representations at level 2 compared to that at level 1.

Default Mode Network

Apart from the MD regions, there are other fronto-parietal regions that show modulation of BOLD signal during task phase, however these regions that are part of what is commonly called the default mode network show a decrease in BOLD signal during task periods, when MD regions show increased BOLD signal. It is commonly assumed that this decrease in DMN regions is an undifferentiated response to cognitive load during the task periods compared to periods of relative cognitive inactivity (Fox et al., 2005; Buckner et al., 2008). However, some recent evidence suggests that these regions do have a task dependent modulation of activity. For example, Mayer et al., (2009) showed that qualitatively different regions of the DMN responded to visual attention load versus working memory load.

A recent study gives evidence of possible role of DMN regions in the accessing hierarchically organised information (Rogers et al., 2010). Subjects were trained to do discriminate between image pairs using transitive inference. Some image pairs were derived from an implicit image sequence (e.g. A>B>C>D>E), whereas others were independent (e.g. F>G, H>J). The authors found greater functional connectivity between DMN regions (like posterior cingulate) and task related (multiple demand)

fronto-parietal regions when subjects were making judgements on pairs belonged to the implicit sequence than when the image pairs belonged to independent sets.

It is not known whether activity in the DMN fronto-parietal ROIs is modulated by the organisation of the task episode. Since the four target detections in the current task are matched in their attentional and working memory demands, a variance in their modulation of activity across the four targets would suggest that these regions are indeed sensitive to the organisation of the task episode.

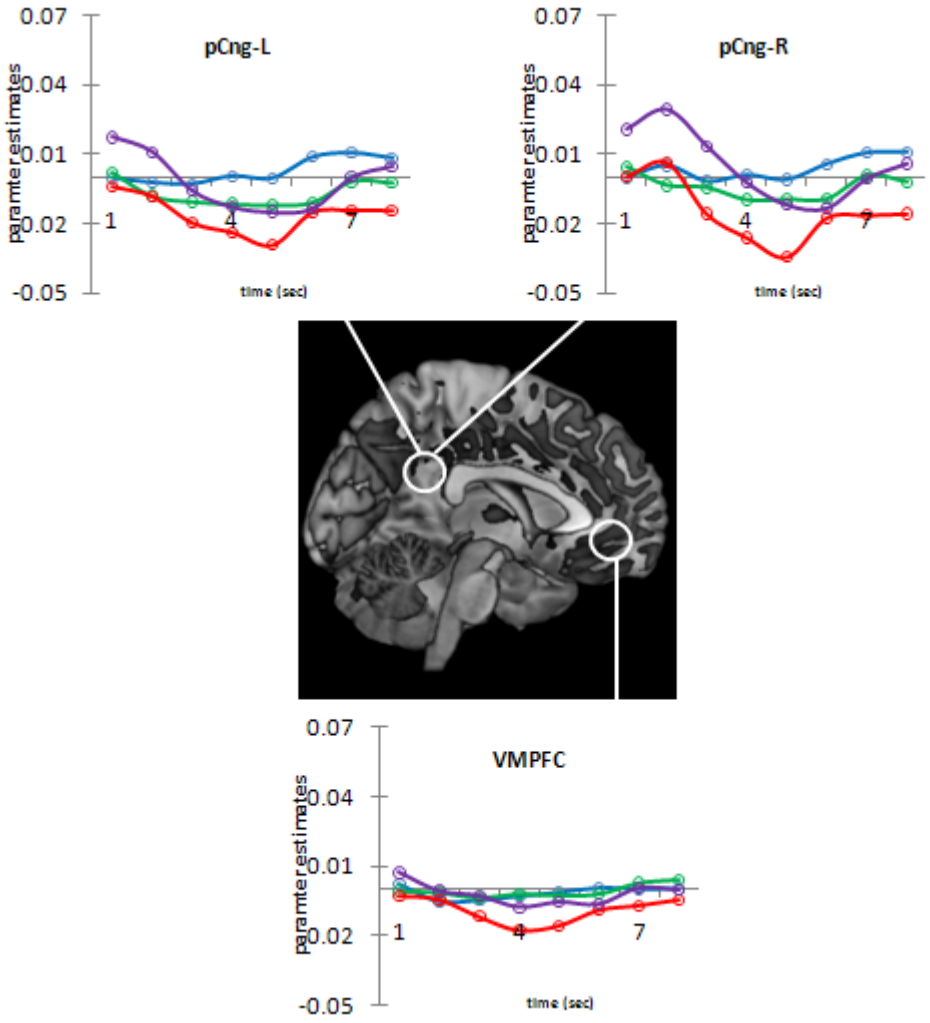

Figure 2.7. (Exp 1) Estimated time course of activity in the Default Mode ROIs (blue: T1; green: T2; red: T3; purple: X). Note that the change in activity is greater for T3 and X.

Three ROIs were constructed based on Dosenbach (2007) – ventromedial prefrontal cortex (1, 31, -2), and bilateral posterior cingulate (10, 56, 16; -11, 77, 13). The time course of activity in these ROIs for the different target events (Fig. 2.7) show that while first and second target detection barely cause any decrease in activity, the third and the final target detections caused mark decrease in BOLD signal. To compare the phasic change in activity in response to the various target detections, the first five FIR bins were compared for each ROI across the four target detection events to see if the effect of time interacted with that of the target type.

All three DMN ROIs showed a significant interaction between time bin and target type, suggesting that modulation of activity in these regions was different for different target types (vmPFC: $F_{(12,\ 204)} = 2.07$ $p<0.05$; right posterior cingulate: $F_{(12,\ 204)} = 4.03$, $p<0.001$, left posterior cingulate: $F_{(12,\ 204)} = 2.74$, $p<0.01$). The effect of task episode organisation is not just limited to fronto-parietal regions that show an increase in activity during task episodes. The result further suggests that the modulation of activity in the DMN regions during task episodes is not undifferentiated inhibition rather they also seem to carry information about the structure of organisation of the episode.

2.3.2 Experiment 2

In Experiment 2, only the 3 letters of the target word (T1 to T3) were to be found on each trial. By associating T3 with task termination and reward – which in Experiment 1 were linked to the final target 'X' – we planned to investigate the importance of these factors. For brain regions where activity is sensitive only to task/subtask structure, the contrast of T3 vs. T1/T2 should show similar results in the two experiments. Additional influence of task termination and reward should be

shown by additional T3 activity in Experiment 2, corresponding to the strong activity seen, in Experiment 1, for the final 'X'.

Behaviour

Average response time to the probe was 770 ± 18 ms, and accuracies for responses to the probe exceeded 95% for most subjects (mean 96.2 ± 2.7).

Imaging

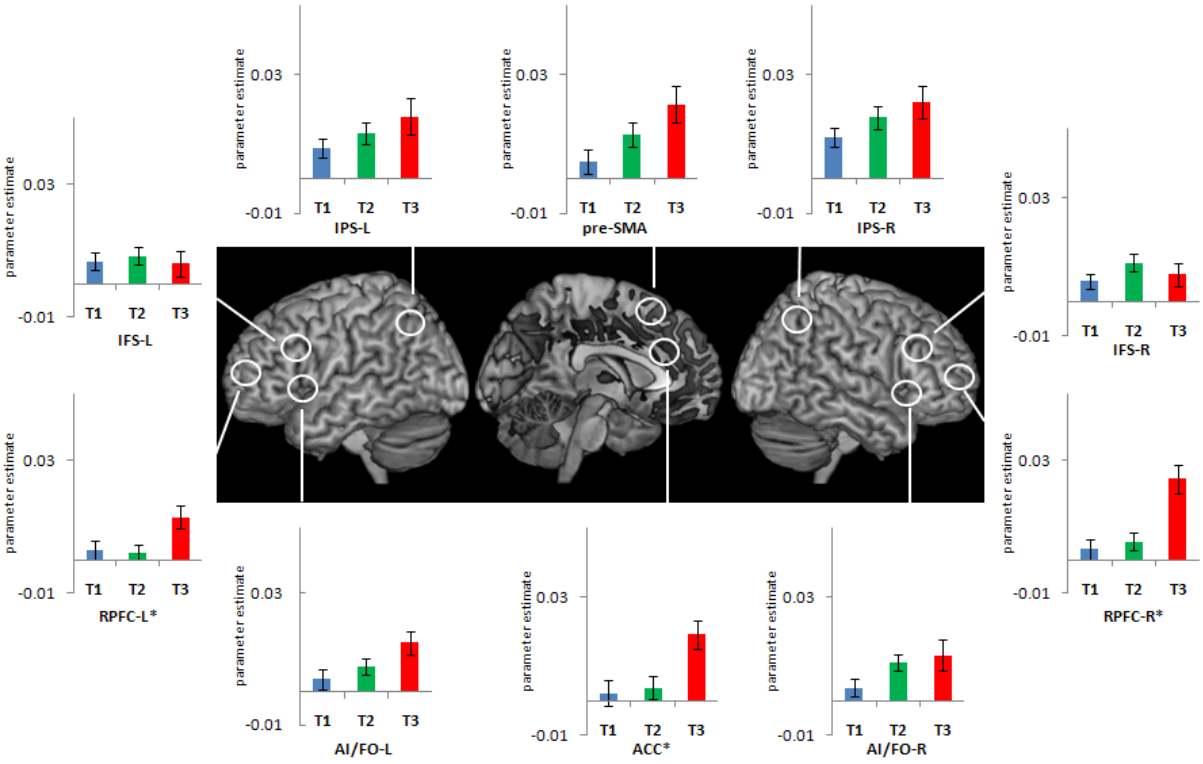

Figure 2.8. (Exp 2) Comparison of phasic activity in response to the three target events in experiment 2. Only APFC and ACC showed effect of target position on the magnitude of activity.

Fig. 2.8 shows plots of index of activation for the three target events from the various MD ROIs. A repeated measures ANOVA with target type (T1, T2, T3) and hemisphere (for lateral ROIs only) as factors showed the significant effect of target type only in APFC and ACC, where responses to T1 and T2 were not significant (one sample $t_{(20)} < 1.5$, $p>0.12$)

To directly compare the two experiments, we carried out an ANOVA with factors target type (T1, T2, T3), hemisphere (for lateral ROIs only) and experiment. No ROI showed a significant interaction, either between experiment and target type ($p >0.3$ in all comparisons) or between experiment, target type and hemisphere ($p > 0.08$ in all comparisons). The results suggest that associating T3 with task termination and reward has little effect on the fronto-parietal activity. Instead it is the hierarchical position of the event in the plan of the task that determines activity, in particular the strong response to T3 in the APFC and ACC.

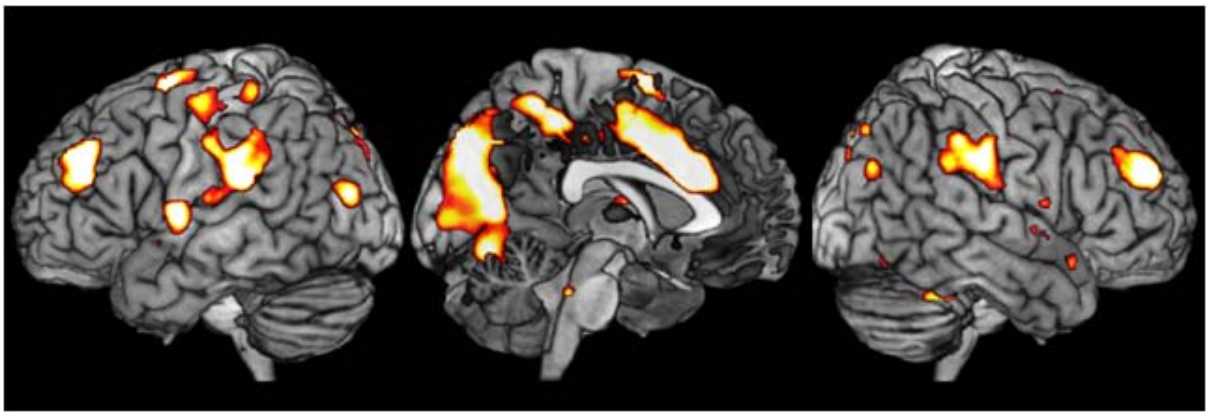

Figure 2.9. (Exp 2) The contrast of T3 against T1 & T2 showed significant activity in anterior prefrontal, pre-SMA, ACC along with regions near TPJ, precuneus and cuneus.

Lastly, the activity during T3 was contrasted with that during the earlier target detections. The results in fronto-parietal regions were very similar to those in the contrast of T3 and earlier targets in the previous experiment, although the current contrast additionally involved precuneus, cuneus and TPJ. Recall that the latter set of

regions were seen in the contrast of earlier targets with X. This makes it likely that these regions were specifically active at the end of the overall task episode.

2.4 Discussion

These experiments show that in various regions of the fronto-parietal cortices, activities generated by task events vary with the hierarchical status of the episode they complete. Events completing the lowest level subtasks (level 1), led to the weakest fronto-parietal activity. At the next higher level (level 2), stronger activity was seen especially in anterior and medial prefrontal regions. Completion of the top-level task or the whole trial (level 3) resulted in greatest and most widespread activity. These results show that the fronto-parietal activity elicited by a task event depends upon the role played by it in the overall organization of the behaviour.

Why should completion of task episode elicit fronto-parietal activity? A task episode is characterized by the organization of multiple mental processes to enable task-relevant processing in various cognitive domains. While the exact nature of mental structures that organize such an episode is a matter of ongoing investigation, task descriptions, representations of purpose/goal (Miller and Cohen, 2001; Duncan, 2006), and task rules (Duncan et al., 2008; Sakai, 2008; Dumontheil et al., 2010) are considered important components that could act as the source of control and bias representations in diverse brain systems to achieve relevant neural processing. At the end of a task episode, these structures would be changed or dismantled depending upon the nature of the future behavioural episode. A plausible interpretation of our result is that the restructuring of these mental representations at the end of the corresponding episode is the elicitor of fronto-parietal activity.

Purposive behaviour is further characterized by a hierarchical organization in which the overall task episode is further organized into smaller nested episodes, each one having its own cognitive focus and organized to achieve a smaller goal

(Lashley, 1951). This organisation is achieved through a similar nested hierarchy of mental representations (Miller et al., 1960; Schank and Abelson, 1977; Barsalou 1988; Reed et al. 1995; Schneider and Logan, 2007), wherein the lower level representations organize smaller episodes of behaviour which in turn are organized by higher level representations into larger and more complex behavioural episodes. As lower level nested episodes are completed, the lower level representations lose their relevance, while the higher level representations organizing larger episodes remain in control; on the other hand completion of the overall goal changes the representations at all levels and brings about a greater change in the mental organisation. It can hence be expected that the activity elicited by the changes in the mental representations organizing the behaviour would be more intense/widespread at the completion of higher level episodes. Our results above show that this was indeed the case in fronto-parietal regions.

Apart from a change in the underlying representations, completion of purposive episodes additionally corresponds to the termination of the ongoing task (and so a change in the task set) and to a feeling of reward corresponding to the achievement of the goal. These, however, appeared to be less significant. T3 of Experiment 2 was similar to 'X' of Experiment 1 in terms of these factors, since visual search terminated at both, and, plausibly, both resulted in a feeling of a successful completion of the trial. T3 of Experiment 2 additionally had a monetary reward attached to it. However, in terms of the level of episode completed, this T3 was identical to the T3 of experiment 1 (Fig. 2.1 & 2.2). As would be expected if the hierarchy of episode completed was the major determinant of fronto-parietal activity, T3 in the two experiments elicited very similar patterns of fronto-parietal activity (Fig. 2.4 & 2.8). Hierarchical changes in the representations organizing the behaviour, therefore, seem to be the most plausible explanation of our results.

Earlier investigations into changes in task related representations have approached it with the concept of task set (Sakai, 2008), understood as a neurocognitive structure linking the stimulus to the relevant response as per the task rules and have shown that restructuring the task set in response to changes in task rules corresponds to activity in the fronto-parietal regions. Our concept of the organising representations subsumes this concept of task set and also includes other representations like task and goal descriptions. End of task episodes will always revise the representations corresponding to the task and goal descriptions even if the relevant rules linking stimulus to response remains the same, as was the case in our experiments. As discussed above changes in the task set that occurred at the termination of the visual search were less important than changes in task description.

Our results also showed a functional distinction amongst the fronto-parietal areas. Regions like IPS, IFS and pre-SMA were active during all target events, whereas APFC and ACC were active only at the completion of task episodes, suggesting that while any change in the underlying representations elicited activity in the former areas, only changes in higher level representations activated the latter. APFC and ACC have been shown to be active at the beginning of the trial blocks (e.g. Dosenbach et al., 2006; Sakai and Passingham, 2006; Haynes et al.,, 2007) and to have sustained activity for the duration of task episodes (e.g. Koechlin et al., 2003; Braver et al., 2003; Dosenbach et al., 2006; Marklund et al., 2007). In the light of our findings, it is plausible that this pattern of activity reflects the role of these regions in the creation and maintenance, respectively, of the higher level task representations that organize the whole episode.

Compared to our finding that anterior prefrontal activity can be elicited in very simple mental conditions without any serious requirement of control or working memory organisation and manipulation, a number of other studies have found such

activity in situations that require maintenance of information across periods of unrelated tasks (Koechlin et al., 1999; Pisapia et al., 2007; Reynolds et al., 2009; Gilbert, 2011). Typically, such studies compares trials in which more than one task are executed such that information about one of the task has to be maintained while a different task is being executed and which is followed by the task that required the maintained information, to trials in which the same tasks are executed sequentially but without there being a requirement to maintain any information across a different task period.

In a task by Koechlin et al., (1999) subjects did four kinds of trials (Fig. 2.10). In the control condition, they judged whether two successively presented letters were also in immediate succession in the word 'tablet' (only upper-case letters were presented). In the delay conditions, subjects had to ignore lower-case letters which were used to occasionally delay the response required by the control condition. In the dual task condition, subjects had to respond as in the control condition for both upper and lower case letter series, but at every first letter showing a case change they judged if the letter was 'T' (or 't'). Finally in the branched condition, they had to respond to the upper case as in the delay condition and to the lower case as in the dual task condition. This meant that specifically in the branch conditions, subjects had to remember the identity of the last upper case letter presented before the case switch, while they were responding to the lower case letters. Anterior prefrontal activity was found associated with the branching task condition.

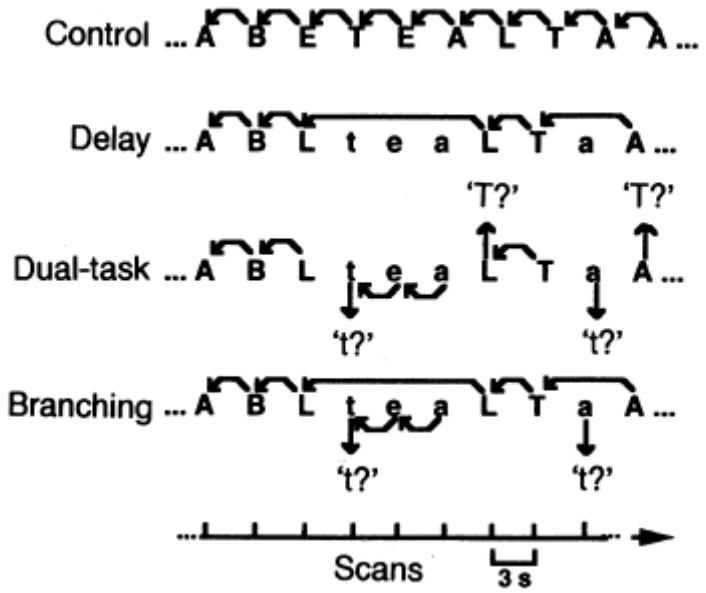

Figure 2.10. Design of task used by Koechlin et al., (1999). Four kinds of trials were used. Control condition: subjects had to decide whether two successively presented letters were also in immediate succession in the word `tablet' (only upper-case letters were presented). Delay condition: subjects had to ignore lower-case letters which were used to occasionally delay the response required by the control condition. Dual task condition: subjects had to respond as in the control condition for both upper- and lower-case letter series with one exception. Subjects had to decide whether every first letter indicating a case change was the letter T (or t). Branching condition: subjects had to respond to upper-case letters exactly as in the delay condition and to lower-case letters exactly as in the dual-task condition. Reproduced from Koechlin et al., (1999)

Conceptually similar task design was used by Pisapia et al, (2007), Reynolds et al., (2009) and Gilbert (2011), in which the condition resulting in the anterior prefrontal activity involved a task episode during which some not immediately relevant information was to be maintained, which was used to execute a task after the current episode was over.

It can be argued that the requirement to maintain information across an irrelevant episode binds what would have been two unconnected task episodes into a common task. Binding the separate task episodes into a common task would create a task plan that has higher levels of organisation than were present in the separate

task episodes. In the dual task condition of Koechlin et al., (1999), the lower case and higher case episodes were disconnected to each other. However, in the branching task condition, the requirement to maintain the upper case information across the lower case episode means that the lower case episode be represented as part of the same higher level mental episode as the upper case one, resulting in a more hierarchical mental plan.

Another support for this thesis comes from studies on reasoning (Smith et al., 2007). Compared to the trials that require separate inferences of the relation amongst two separate pairs of objects, trials that additionally require a comparison of these relations elicited greater activity in APFC. Again, making separate inferences on two separate pairs creates two separate mental episodes, while the requirement to compare the two inferences could result in a common organisation of these hitherto separate episodes.

The current results show that a hierarchical mental plan without any requirement of maintaining and integrating information across branched and unrelated task episodes can result in anterior prefrontal activity.

Other fronto-parietal regions – IFS, IPS and pre-SMA, in contrast, are active with any cognitive challenge (Duncan, 2010). Regions like IFS (DLPFC in many studies) show transient activity in situations that show sustained activity in APFC (Braver et al., 2003; Dosenbach et al., 2006; Reynolds et al., 2009). Interestingly, change in task representations within an episode of behaviour, for example, switching between the different rules active in the block, frequently result in fronto-parietal activity which in lateral PFC tends to be limited to the IFS (e.g. Dove et al., 2000; Braver et al., 2003; Brass et al., 2005; Savine and Braver, 2010; Kim et al., 2011).

Additionally our results also suggest a functional distinction along the antero-posterior axis of the lateral and medial prefrontal cortex. Consonant with our finding

that the behaviour of IFS and pre-SMA were similar, while that of APFC paralleled ACC, Taren et al. (2011) found that the posterior regions of medial PFC had maximal resting state functional connectivity with the posterior parts of lateral PFC, whereas anterior regions of the two were maximally connected with each other.

Earlier studies have tried to link different aspects of hierarchical behaviour with different regions along the antero-posterior axis of the lateral PFC. Badre and D'Esposito (2007) noted that tasks with increased levels of decisions were accompanied with activity in more anterior regions of lateral PFC. Compared to single level decisions where one feature (e.g. colour) determined the response, two level decisions, where one feature (e.g. shape) determined how the other feature (e.g. colour) would determine the response, was associated with more anterior locus of prefrontal activity. APFC activity was found in tasks that required four level decisions (i.e. when the context of the trial block determined the relevant dimensions which in-turn determined the relevant features that determined how the colour would map to the response). Koechlin et al. (2003) found that increasing the temporal abstraction of the context that determined a response led to more anterior locus of prefrontal activity, for e.g. compared to the trials when the trial information was sufficient to allow response selection, trials that required the consideration of the context of the block for response selection to occur elicited activity at more anterior loci in the lateral PFC.

Clearly in both of these studies, as noted earlier for studies showing activity in anterior prefrontal regions, the kinds of trials that result in anterior activity require a more complex behavioural plan for execution. For example, in Badre and D'Esposito (2007), compared to trials with single level decision wherein the colour directly allows for the selection of response, trials with two-level decisions have got an additional level of hierarchy - subjects first determine what colour to response mapping is valid in the light of the shape (subgoal), and then determine the relevant

response based on the colour (completion of goal). Similarly, in Koechlin et al. (2003), the trials wherein the context of the block determined how the stimuli will map onto the response, clearly required an additional level of representation (context of the trial block) compared to the trials where the response could be determined by the stimulus itself.

A difference between earlier studies and our study is worth noting. Earlier studies had found APFC activity in trials that had very high cognitive demands compared to trials that had activity limited to posterior prefrontal regions. For e.g. four nested levels of decision processes (Badre and D'Esposito, 2007) or information integration across three/four levels (Koechlin and Summerfield, 2007). In our experiments simple events marking the end of trial, without concomitant complex decision process or complex information integration, produced activity in these regions (cf. Hon et al., 2006).

Neurophysiological evidence is also supportive of our findings. Neurons from around the principal sulcus have been found to be active during all kinds of task events (reviewed in Duncan, 2010) and seem to code for all stages of the task episode (Sigala et al., 2008). In contrast, in the only neurophysiological study of the anterior prefrontal cortex, Tsujimoto et al., (2010) found its neurons to be active only at the end of the task episode.

An important suggestion of the current results is that mental organisation is itself an elicitor of fronto-parietal activity. It is obvious in the light of above results that to answer a question like – what neural regions underpin visual target detection? – requires consideration of the role of the act of target detection in the mental organisation underpinning the experimental design. An experimental trial is itself an organised mental episode, this fact needs to be kept in mind while making conclusions about the neural correlates of cognitive events.

CHAPTER 3

Phase of Task Episode and Fronto-Parietal Activity

3.1 *Introduction*

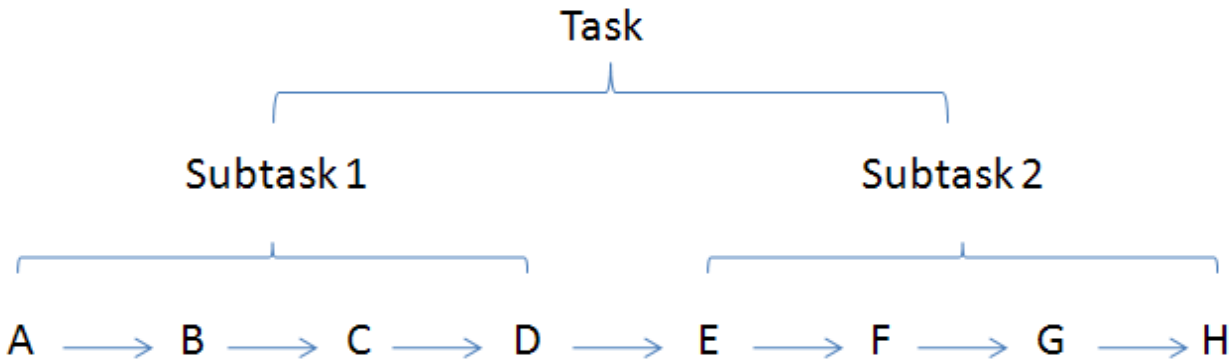

Figure 3.1. Scheme of a typical plan underlying a goal directed behavioural episode. The sequence of cognitive/task events are organised into discrete subtasks, which are further grouped into a higher level task.

Consider a behavioural episode consisting of a sequence of cognitive events ('A' to 'h'), which is organised as a single task episode comprising of two subtasks (Fig. 3.1). In this behavioural episode, apart from their intrinsic identities (perception, decision making, response selection, task switch and so forth), the cognitive events also have unique roles due to their position in the plan of the behaviour. These roles would be the hierarchy of goal they achieve ('D' completes a subtask, while 'H' completes a task and the whole plan) and the phase of the plan they represent (events in 'subtask1' represent the early phase of the plan whereas those of 'subtask2' represent the final phase since they culminate in the completion of the plan). These roles of the cognitive events come from the peculiarities of the

organisation of the mental episode, which, amongst others, would depend upon the goals to be achieved, the trajectory planned to achieve the goal, the organisation of the semantic knowledge of the world, the conceptual structure of the actions to be done, habits and the environmental/social constraints of the subject.

How does the context of the plan affect the neural activity elicited by a cognitive event within it? Does the role that the cognitive event has in the plan of the behaviour affect its representation in the fronto-parietal regions? In the previous chapter, it was shown that compared to the target detections within a subtask, those completing the respective subtask and task led to increasingly greater fronto-parietal activity, and the target detection that completed the whole plan resulted in maximal activity. This chapter will explore whether the phase of the plan that the cognitive events represent affects their elicited neural activity. This should be the case, if the fronto-parietal activity representation of an event depends upon the larger context. During the early phases of the episode, representations about the upcoming phases or steps have to be additionally maintained, while during the final phase only the representations relevant to that the current step need to be maintained. It is possible that this might affect the fronto-parietal activity elicited during the different phases of the task episode.

Our experimental design required the subjects to do two blocks of tasks on every trial (fig. 3.2). One block consisted of arithmetic tasks and the other of word tasks, their order in the trials was balanced, such that the first and the second task blocks composing the trial consisted of both kinds of tasks and were identical when averaged across the experimental session. To ensure that the subjects conceive the two blocks as part of the same mental plan, information relevant to both blocks was cued simultaneously at the beginning of the trial and the subjects executed the first block using the information relevant to that block, while separately maintaining the information relevant to the second block, and subsequently executed the second

block. Due to this feature of the design, while completing the first block, the subjects still had to maintain the relevant information for the second block and only at the completion of the second block could they rest and wait for the next trial to begin. Thus, the plan of a trial can be considered to be biphasic –the completion of the first task block completes the initial phase, and the completion of the second completes the final phase. Since subjects always maintained information relevant to the second block while executing the first, when averaged across the length of the experimental session, the cognitive load of the first block was higher than that of the second. Hence from this aspect of the design it could be expected that brain areas sensitive to cognitive load would be more active during the first block, whereas brain areas that prefer the final phase of a mental plan would be more active during the second.

The main aim was to see the effect of the phase of the mental plan represented by the task on its elicited fronto-parietal activity. In current understanding, the most reliable determinant of task-related fronto-parietal activity is its associated cognitive load (Miller and Cohen, 2001), so to get a relative picture of the effect of phase, the level of difficulty of the two task blocks was additionally manipulated independently between two levels – 'easy' and 'hard'. The execution of the harder blocks was more complex and required maintenance of greater amount of information from the cue displayed at the start of the trial.

The experimental design allowed for two distinct kinds of cognitive load manipulations during the first task block (Fig. 3.2). In the trials in which this was the hard block, greater information load maintained from the cue was relevant to it (hence, in the context of the first block, referred to as the *current* load); however, when the second block was the hard block, the extra information load maintained during the first block pertained not to the current block, but the future block (hence referred to as the *prospective* load).

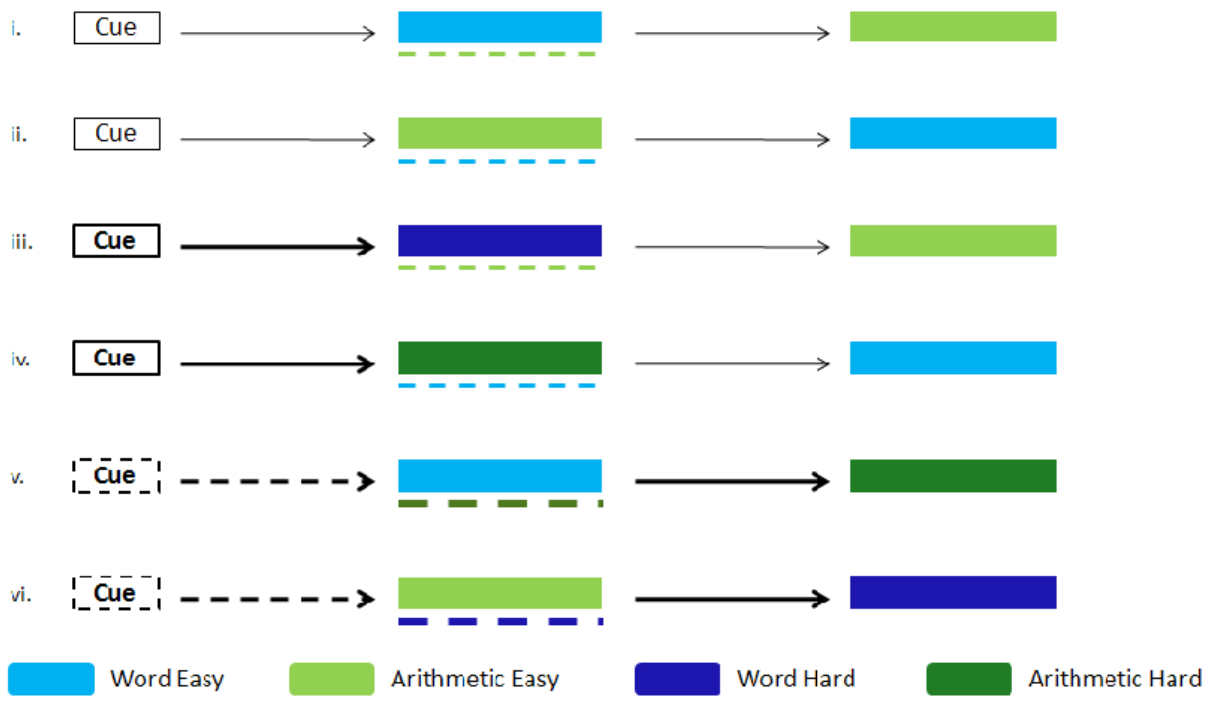

Figure 3.2. Scheme of the different kinds of trials. Thin lines refer to periods with low working memory load, thick lines to high working memory load relevant to the imminent or the current task and dashed lines to high working memory load relevant to the future task.

It is plausible that these two kinds of information are represented separately. Greater current load could be represented as part of the current task set (Sakai, 2008), whereas prospective load might require a separate and distinct representation. This view gets support from some studies (Koechlin et al., 1999, 2007; Pisapia et al., 2007; Reynolds et al., 2009) that show that blocks requiring maintenance of prospective information were associated with activity in the anterior prefrontal cortex, compared to blocks with current load only, which were associated with activity in the dorsolateral prefrontal cortex. Hence, it was predicted that the modulation of prospective load would be associated with a different pattern of activity than that elicited by the modulation of the current load. Further, since prospective load was limited to the first block while the second block had no such requirement, regions involved in maintaining such prospective information would

show greater activity during the first compared to the second block. On the contrary, regions that prefer the final phase of the task plan would elicit greater activity during the second block compared to the first.

It was found that increased current cognitive load resulted in an increase in activity in the fronto-parietal cortices, specifically in the multiple demand regions (Duncan, 2010); however, increased prospective load did not result in significant activity in any brain region. Importantly, we found that the phase of the task indeed had an effect on the elicited fronto-parietal activity: the second task block, although identical to the first and with lesser load, was associated with much greater activity in widespread regions of the fronto-parietal cortex, including those that had shown an effect of current cognitive load.

3.2 Methods

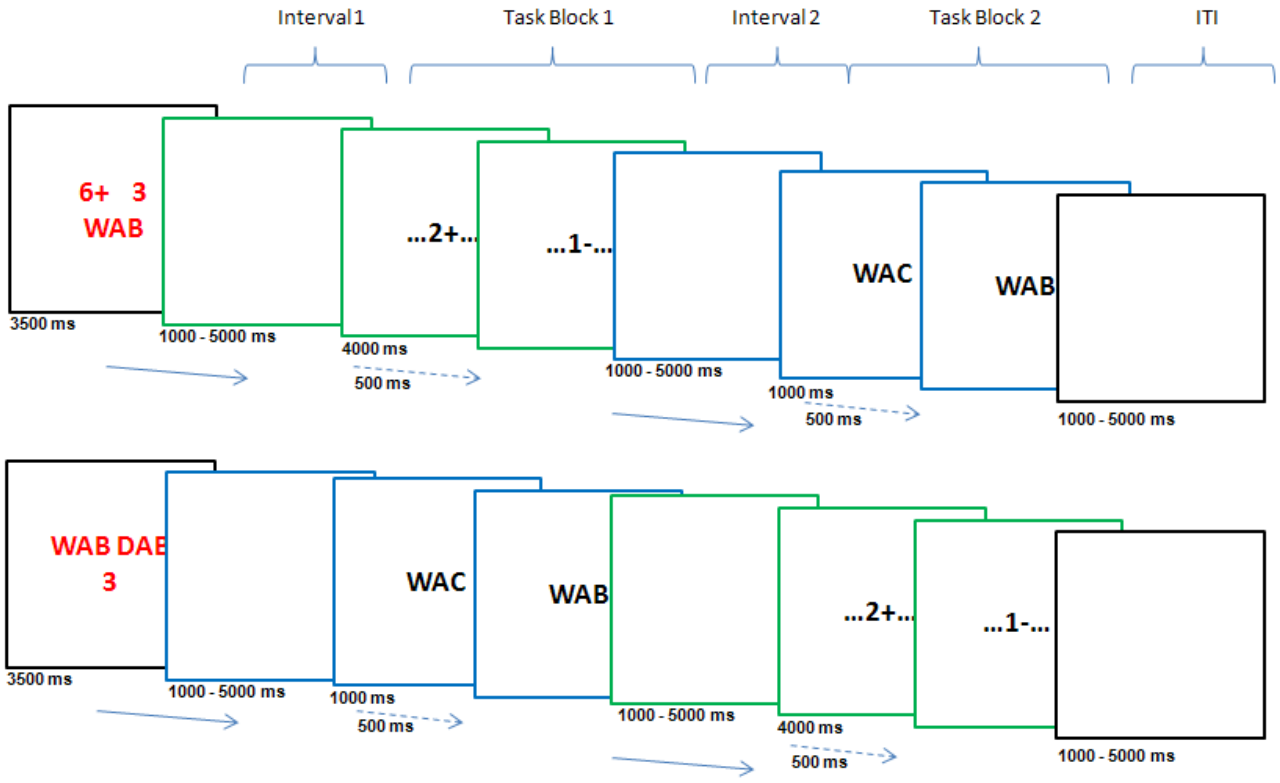

Figure 3.3. Scheme of trials. Trials began with a cue display of numerical information (for the arithmetic block) and pseudoword(s) for the word block. In trials with an easy arithmetic block a single number was displayed, whereas two numbers and an arithmetic operator were

displayed on trials with a hard arithmetic block; one and two word(s) were displayed prior to easy and hard word blocks respectively. The information displayed on top signalled the nature of first task. The cue was followed by the first interval when subjects prepared for the first block. An outlining square the coloured margins of which reflected the upcoming block (blue: word; green: arithmetic) was displayed throughout this interval. The two task blocks consisted of a series of stimulus presentation to which subjects made the relevant response using the relevant cued information (see text). The two task blocks were always different within a trial but identical when averaged across the experimental session.

As schematised in fig. 3.3, trials began with a cue display (on for 3500 ms) of numerical information for the arithmetic task and word(s) for the word task. The numerical information could be a single number in case of easy blocks e.g. “5”; and two numbers and an arithmetic operator, + or -, in the format “x± y” (e.g. “2+ 3”) for difficult blocks. One word was presented before easy blocks e.g. “kaf” and two before the difficult blocks e.g. “daj & gaz”. The information relevant to the first task block was displayed on top, e.g. numerical information above the word(s) meant that the arithmetic task would be the first block followed by the word block.

The cue display was followed by a jittered gap, hence referred to as the first preparatory interval (1000 to 5000 ms) during which an outline square whose coloured margins signalled the nature of the forthcoming task block (green- arithmetic block, blue- word block) was displayed on the screen to reinforce the relevant task set. The jittered temporal duration kept the subjects in a state of readiness throughout the interval. At the end of this interval the stimuli relevant to the first task block appeared in this square. In the arithmetic blocks these were a series of one to three 4000ms displays of a number and an arithmetic operator, preceded and succeeded by blanks (e.g. “...8-...”); in word blocks a series of two- to nine-word stimuli (“baf”) were displayed (1000 ms per display). The inter stimulus interval (ISI) was 500 ms in all task blocks.

During easy arithmetic blocks, subjects mentally put the cued number (“5”) in the second blank (“...8-**...**”), solved the resulting expression (“8-5”) and then answered through the button box positioned under their right hand if the solution was less than 10 (index finger: 'yes', middle finger: 'no'). In difficult blocks, when the cued information maintained consisted of two numbers and an arithmetic operator (“2+ 3”), the subjects placed the number and the operator that were displayed on the left side in the cue, in the first blank of the stimuli (“2+8-...”), and the number that was displayed on the right side in the cue in the second blank (“2+8-3”), and, as in easy blocks, mentally solved the resulting expression and responded if the solution was less than 10; for e.g. if the cue was “2+ 3” and the displayed stimuli “...8-...”, the resulting expression becomes “2+8-3”.The arithmetic operations always involved numbers less than 20. The stimuli in these blocks were displayed for 4000 ms and a valid response could be made during this period. Number of sequential stimuli presented in an arithmetic block varied randomly from one to three with an ISI of 500 ms, giving the duration of such blocks a jitter between 4 to 13 seconds. In the word task blocks, subjects responded if the displayed word was the same as the cued word in easy blocks or was one of the cued words in difficult blocks (index finger: yes; middle finger no). Such blocks consisted of two to nine sequential stimuli, each displayed for 1000 ms with an ISI of 500 ms, giving their duration a jitter of 2.5 to 13 seconds.

The first task block was followed by the second preparatory interval (which was identical to the first interval) during which a blank margin was displayed, the colour of which signalled the nature of the second task block; within a trial the nature of the task in the two blocks was always different. This interval lasted for 1000 to 5000 ms during which the subjects prepared for the second block. The structure of the ensuing second block was same as the first, except that it ended with a black margined square at which subjects could relax and wait for the next trial to begin (inter trial interval: 1000 to 5000 ms).

In fifteen percent of the trials the first task block did not appear. In such cases the first preparatory interval ended with a change in colour of the margin which marked the beginning of the second preparatory interval. At this point the subjects switched to preparing for the second task block. In another fifteen percent of trials the second task block did not appear. In such cases the second preparatory interval ended with the margins of the blank square turning black signalling that the trial had ended. (See Appendix B.3 for exact instructions).

The design of trials described above thus created two distinct phases in each trial wherein subjects prepared for and then executed two distinct tasks. Averaged across the experimental session, the first and the second task blocks were identical except for their position in the trial. The second task block can be considered to represent the final phase of the trial since it culminated in its completion. The overall experimental design of two kinds of task blocks (arithmetic and word) with two levels of difficulty (easy and hard) at two phases (initial and final) of a trial potentially created eight kinds of trials (depicted here by the identity of their first and second task blocks) –

I. Easy Arithmetic → Easy Word

II. Easy Word → Easy Arithmetic

III. Hard Arithmetic → Easy Word

IV. Hard Word → Easy Arithmetic

V. Easy Arithmetic → Hard Word

VI. Easy Word → Hard Arithmetic

VII. Hard Arithmetic → Hard Word

VIII. Hard Word→ Hard Arithmetic

The last two kinds of trials were not used since behavioural pilots showed too many errors in such trials due to excessive working memory load. Our experiment therefore had six kinds of trials (I to VI).

The requirement that the information relevant to both task blocks be encoded at the cue display meant that the subjects had to maintain information relevant to the second task block while executing the first, resulting in greater cognitive load in the initial phase of the trial compared to the final phase. This cognitive load, hence called prospective load, depended upon the level of difficulty of the second block since the amount of information to be maintained differed for easy and difficult blocks. Accordingly the first blocks preceding easy second blocks had lower prospective load compared to the first blocks preceding difficult second blocks.

The first task blocks were therefore be classified as easy (in trial types I and II), hard (in trial types III and IV) and easy with greater prospective load (in trial types V and VI). The 1st intervals preceding these task blocks were likewise classified as having low or high information load relevant to the first task block (i.e. low or high current load) and low or high information relevant to the second task block (low or high prospective load). The second task blocks and their preceding intervals could only be classified as easy (I to IV) and hard (V and VI), since they had no prospective load.

All stimuli (visual angle ½ deg x 1 deg) were centred on the screen, visible from the participant's position in the scanner via a mirror mounted within the head coil. The experiment was controlled by a program written in Visual Basic. Subjects did a 20 min pre-scan practise session. The scan lasted for an hour and was divided into three separate 20-minutes sessions, each having 55 trials.

3.2.1 Participants

15 participants (6 females, mean age = 23.5 ± 3.6 years) were recruited from the MRC-CBU volunteer panel. Participants were right handed and had normal or corrected vision. Informed consent was taken and the participants were reimbursed for their time. The study had the approval of the Hertfordshire Local Research Ethics Committee.

3.2.2 Acquisition

fMRI data were acquired using a Siemens 3T TimTrio scanner with a 12 channel head coil. A sequential descending T2*-weighed echo planar imaging (EPI) acquisition sequence was used with the following parameters: Acquisition time 2000 ms; echo time 30 ms; 32 oblique slices with slice thickness of 3 mm and a 0.75 mm interslice gap; in-plane resolution 3.0x3.0 mm; matrix 64x64; field of view 192mm and a flip angle of 78 deg. T1-weighted MPRAGE structural images were also acquired for all participants (slice thickness 1.0 mm, resolution 1.0x1.0x1.5 mm, field of view 256 mm, 160 slices).

3.2.3 Analysis

The fMRI data were analysed using SPM5 (Wellcome Department of Imaging Neuroscience, London, England; www.fil.ion.ucl.ac.uk). Prior to statistical analysis, all EPI volumes were slice-time corrected using the first slice as a reference, and then realigned into a standard orientation using the first volume as a reference. These realigned images were then normalised into the Montreal Neurological Institute (MNI) space, and spatially smoothed using an 8 mm full-width half-maximum (FWHM) Gaussian kernel. During the normalisation stage, voxels were resampled to a size of 3 × 3 × 3 mm. The time course of each voxel was high-pass filtered with a cutoff period of 90 s.

The various stages of the trial – cue, first preparatory interval, first task block, second preparatory interval and the second task block – were modelled using epoch regressors, of width equal to their duration. Further, the different categories of these stages based on their associated cognitive load (see above) were separately modelled. Event regressors of no duration modelled the completion of the two task blocks. All regressors were convolved with a basis function representing the canonical hemodynamic response. Movement parameters and block means were included as covariates of no interest. Parameter estimates for each regressor were calculated from the least squares fit of the model to the data, and estimates for individual participants were entered into a random effects group analyses. The results were rendered on to a whole brain template displayed at a false discovery rate corrected threshold of $p<0.01$ (unless specified otherwise).

To capture fronto-parietal regions widely engaged in cognitive control, ten regions of interest (ROIs; Table 1, Figure 2) were created as 10 mm diameter spheres at coordinates that have been shown to be consistently active in varied tasks (Duncan, 2006; Dosenbach, 2006). The ROIs (in MNI space) were bilateral inferior frontal sulcus (IFS; central coordinate ±41 23 29), bilateral intra-parietal sulcus (IPS; ±37 56 41), bilateral anterior insula extending into frontal operculum (AI/FO; ±35 18 3), anterior cingulate (ACC; 0 31 24), and pre-supplementary motor area (pre-SMA; 0 18 50), all taken from Duncan (2006); along with bilateral anterior prefrontal cortex (APFC; 27 50 23 and -28 51 15) taken from Dosenbach (2006). ROIs were constructed using the MarsBar toolbox for SPM (http://marsbar.sourceforge.net; Brett, Johnsrude & Owen, 2002). Estimated data were averaged across voxels within each ROI using the MarsBar toolbox and the mean values were exported for analysis using SPSS.

3.3 Results

In the first part I discuss the results pertaining to the effects of cognitive load. In light of these results the second part examines the effect of the phase of the trial.

3.3.1 Cognitive Load

Behaviour

Overall performance was fairly accurate with error rates being less than 3% of the total responses for most subjects (mean 2.8 ± 0.8). Trials with any error were excluded from further analysis. For further analysis the behavioural parameters on the arithmetic and word blocks were separated.

Arithmetic blocks

Arithmetic		1	2	3
E1	RT (ms)	1541 (69)	1149 (90)	1166 (93)
	Error rate (%)	2.2 (1)	2.6 (1.2)	1.1 (0.6)
E2	RT(ms)	1418 (78)	1127 (80)	1059 (61)
	Error rate (%)	2.4 (0.9)	2.2 (0.7)	1.3 (0.6)
H1	RT(ms)	2312 (101)	1873 (109)	1871 (161)
	Error rate (%)	7.6 (2.1)	4.8 (1.9)	4.8 (2.6)
H2	RT(ms)	2303 (87)	1912 (74)	1697 (153)
	Error rate (%)	8.7 (2.3)	5.1 (2.1)	4.7 (2.7)
E1 + P	RT(ms)	1613 (93)	1167 (69)	1104 (67)
	Error rate (%)	4.1 (1.6)	3.4 (1.5)	0 (0)

Table 3.1. Reaction times and error rates at the various positions within the arithmetic blocks, values in parenthesis represent standard error. E1 - Easy block during phase 1; E2 -

Easy block during phase 2; H1 – Difficult block during phase 1; H2 – Difficult block during phase 2; E1 + P – Easy block with high prospective load

Table 3.1 shows the mean RTs and error rates at different positions in the easy and hard arithmetic blocks from the two phases, and in the easy blocks with high prospective load in the first phase. Performance during the different arithmetic blocks was compared to examine the effect of task difficulty (performance on I & II vs. performance on III & VI) and position of the block in the trial (performance on I & III vs. performance in II & VI). Repeated measures ANOVA with level of current task difficulty (easy or hard), task phase (1st or 2nd), and position of the stimulus within the block (1-3) as factors showed a significant effect of the level of difficulty of the block on RT ($F_{(1, 12)} = 131.72$, $p<0.001$), but did not show any effect of the phase of the block ($F_{(1, 12)} = 2.7$, $p=0.12$). Position of the stimuli in the block also affected the RT ($F_{(2, 24)} =48.4$, $p<0.001$). Pairwise comparisons showed that this was driven mainly by the first RT being higher than the rest (mean difference > 378 ms; Holm-Bonferroni corrected $p<0.001$ in all cases) as would be expected from a blocked set of responses (Allport and Wylie, 2000). The second and third RTs did not differ (mean difference = 67.1 ms; 95% CI: -86.9 to 221.1). A similar ANOVA for errors, showed more errors in difficult blocks ($F_{(1,13)} = 10.62$, $p<0.01$), but did not show any significant difference in error rates across the two phases ($F_{(1,13)} = 0.11$, $p=0.74$); error rates varied marginally with the position in the block ($F_{(2,26)} = 3.05$, $p = 0.05$), again reflecting higher error rates at the first position.

A separate ANOVA compared the two easy arithmetic blocks that differed in their prospective load trial type I vs. trial type V. Reaction times (1294 ± 75 vs. 1285 ± 71; $F_{(1, 11)} = 0.65$, $p=0.8$) and error rates (1.3 ± 0.3 vs. 1.2 ± 0.3) were not significantly different between the two.

Word blocks

Word		1	2	3	4	5	6	7	8
E1	RT(ms)	683 (13)	538 (16)	500 (12)	499 (14)	491 (16)	510 (20)	521 (22)	508 (22)
	Error rate (%)	0.0	1.2 (0.3)	1.0 (0.2)	1.6 (0.3)	1.2 (0.4)	1.3 (0.4)	1.7 (0.6)	1.6 (0.8)
E2	RT(ms)	666 (12)	544 (15)	506 (16)	496 (20)	481 (16)	502 (20)	519 (17)	509 (28)
	Error rate (%)	2.1 (0.9)	1.9 (0.3)	1.7 (0.2)	1.4 (0.4)	1.8 (0.7)	1.4 (0.5)	1.8 (0.7)	1.5 (0.7)
H1	RT(ms)	735 (14)	600 (13)	568 (12)	571 (14)	582 (13)	586 (18)	564 (18)	574 (14)
	Error rate (%)	1.8 (0.8)	2.0 (0.5)	1.3 (0.4)	1.9 (0.6)	2.7 (1.1)	2.0 (0.6)	1.2 (0.3)	0.0 (0)
H2	RT(ms)	723 (13)	604 (18)	580 (14)	574 (14)	586 (16)	569 (18)	546 (15)	560 (29)
	Error rate (%)	2.0 (0.8)	1.3 (0.4)	1.6 (0.7)	1.2 (0.3)	1.4 (0.4)	1.6 (0.6)	1.2 (0.5)	0.0 (0)
E1 + P	RT(ms)	711 (18)	552 (19)	524 (17)	520 (16)	524 (18)	523 (19)	533 (21)	535 (30)
	Error rate (%)	1.6 (0.7)	1.6 (0.4)	1.7 (0.6)	1.2 (0.7)	1.8 (0.5)	1.6 (0.8)	1.8 (0.7)	1.9 (0.6)

Table 3.2. Reaction times and error rates at the various positions within the word blocks, values in parenthesis represent standard error. E1 - Easy block during phase 1; E2 - Easy block during phase 2; H1 – Difficult block during phase 1; H2 – Difficult block during phase 2; E1 + P – Easy block with high prospective load

Table 3.2 tabulates the reaction times and the error rates in easy and difficult word blocks, across the two phases. Performance in word blocks was compared for effect of difficulty (I & II vs. IV & V) and position in the trial (I & V vs. II & IV). Repeated measures ANOVA with the level of current task difficulty, phase and the position of the stimulus in the block as factors showed a significant effect of difficulty on RT ($F_{(1, 12)} = 65.8$, $p<0.001$), but no effect of the phase of the trial ($F_{(1, 12)} = 0.6$, $p=0.7$). RTs varied significantly with the position of the stimulus within the block ($F_{(7, 84)} =58.9$, $p<0.001$). Similar to the case in arithmetic blocks, pairwise comparisons showed the first RT to be the highest (mean difference > 140 ms, Holm-Bonferroni corrected $p<0.001$ in all cases). None of the other comparisons reached significance (mean difference <50.1 ms, Holm-Bonferroni corrected $p>0.05$ in all cases). Error rates were not significantly affected by the difficulty ($F_{(1, 10)} = 0.01$, $p=0.9$) or phase ($F_{(1, 10)} = 0.001$, $p=0.87$) or by position within block ($F_{(7,84)}=0.6$, $p =0.4$).

Performance on the word block from trial type II was compared with the word block from trial type VI to look at the effect of greater prospective load. While mean RTs during the task blocks with greater prospective load was higher (558 ± 16 vs. 537 ± 14), the difference failed to reach significance ($F_{(1, 11)} = 3.62$, $p=0.08$). Error rates (2.1 ± 0.8 vs. 1.9 ± 0.9) did not differ significantly between the two blocks either, ($F_{(1, 12)} = 0.46$, $p=0.45$).

These data show that behavioural indices were affected by the level of difficulty of the current task but were not different across the two phases of the trial. The manipulation of prospective load did not have a significant effect on the behaviour.

Imaging

Each trial consisted of discrete stages (fig. 3.2): cue → interval1 → taskblock1 → interval2 → taskblock2. The effect of cognitive load was studied at each stage.

Cue

The cues showed information relevant to the arithmetic and the word blocks to be maintained in working memory until the relevant block was complete. The amount of information to be internalised at the cue display varied depending upon the difficulty of the task blocks of that trial. The cues in the trials with two 'easy' task blocks (trial types I and II), called 'easy' cues, showed a single number relevant to the easy arithmetic block and a single word relevant to the easy word block. In other trials, 'hard' cues consisted either of - two numbers with an arithmetical operation and one word, or of two words and one number.

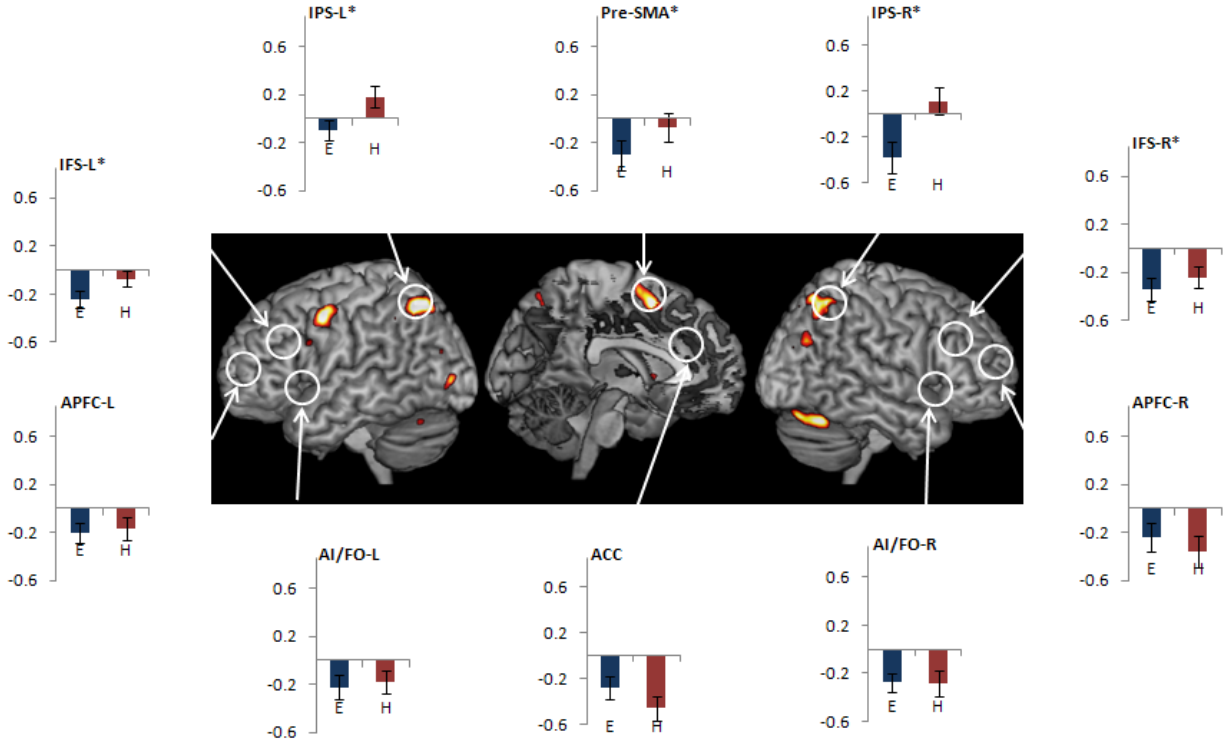

Figure 3.4. The brain render shows areas with greater activity during the hard compared to that during easy cues. Significant activity was present in IPS, left premotor and pre-SMA (marked by asterisk). The plots show estimates of activity in the MD regions during these cues. (E: easy; H: hard)

The brain activity in response to easy versus hard cues was compared. The whole brain render in fig. 3.4, shows regions more active during hard cues.

Significantly greater activity was present in left premotor cortex, left IPS and left pre-SMA.

In the ROI analysis, MD regions were investigated for the effect of cue difficulty. A repeated measures ANOVA with difficulty (low and high) and laterality (for the bilateral ROIs only) as factors found significant effect of difficulty in IFS ($F_{(1, 12)} = 8.5$, $p<0.01$), IPS ($F_{(1, 12)} = 15.5$, $p<0.001$) and pre-SMA ($F_{(1, 12)} = 7.5$, $p<0.01$). The effect of laterality (right > left) reached significance only in IPS ($F_{(1, 12)} = 5.4$, $p<0.01$).

First interval

During the gap between the cue and the first block of the task, two kinds of information were maintained – those pertaining to the first block i.e. the imminent task and those pertaining to the second block i.e. the prospective task. Working memory maintenance, in general, has been shown to be associated with fronto-parietal activity (Courtney et al., 1998). It has also been shown that different kinds of working memory (spatial vs. verbal) relate to activity in different regions of the prefrontal cortex (Sakai and Passingham, 2003). In the current design the working memory contents differed by their relevance to the current and prospective tasks. It is plausible that the working memory contents relevant to the imminent task would be incorporated in the current task set (Sakai, 2008) especially since this information formed the template on which further task-related processing was to be done - the cued numbers were to be used for the processing of each of the presented numerical problems and the cued words were the template against which each presented word was to be judged. On the other hand the information pertaining to the prospective task had to be maintained separately against potential interference from the currently relevant information (Koechlin, 1999, 2007). During different kinds of trials the load of these two kinds of information was separately manipulated – in trials III and IV the current load was high, whereas in trials V and VI the prospective load was high. It was expected that the regions representing this information would be

associated with greater activity when the load of the respective information was high, and this could potentially reveal regions that represent current versus prospectively relevant task information.

Estimates of activity in intervals with greater load were compared with those of intervals with lower load in two separate contrasts (III & IV vs. I & II; V & VI vs. I & II). No brain region showed differential activity in any of these comparisons. Fig. 3.5 shows the estimates of activity during the three kinds of intervals in the MD regions. Only IPS showed differential activity, which curiously was greater for the interval with low load (i.e. from trials I and II), and was only marginally significant ($F_{(2, 24)} = 3.5$, $p=0.048$). Estimates of activity for the three kinds of interval were not significantly different in other ROIs ($F_{(2, 24)} < 2.9$, $p>0.08$, for other ROIs).

Although IPS has been mostly shown to have greater activity with increased memory load (Duncan and Owen, 2000), a decrease in activity in this region has been reported for situations that overwhelm working memory capacity (Linden et al., 2003). Although the difficult trials in the current experiment did require maintenance of three to four items, an overload is unlikely given that the performance was uniformly good.

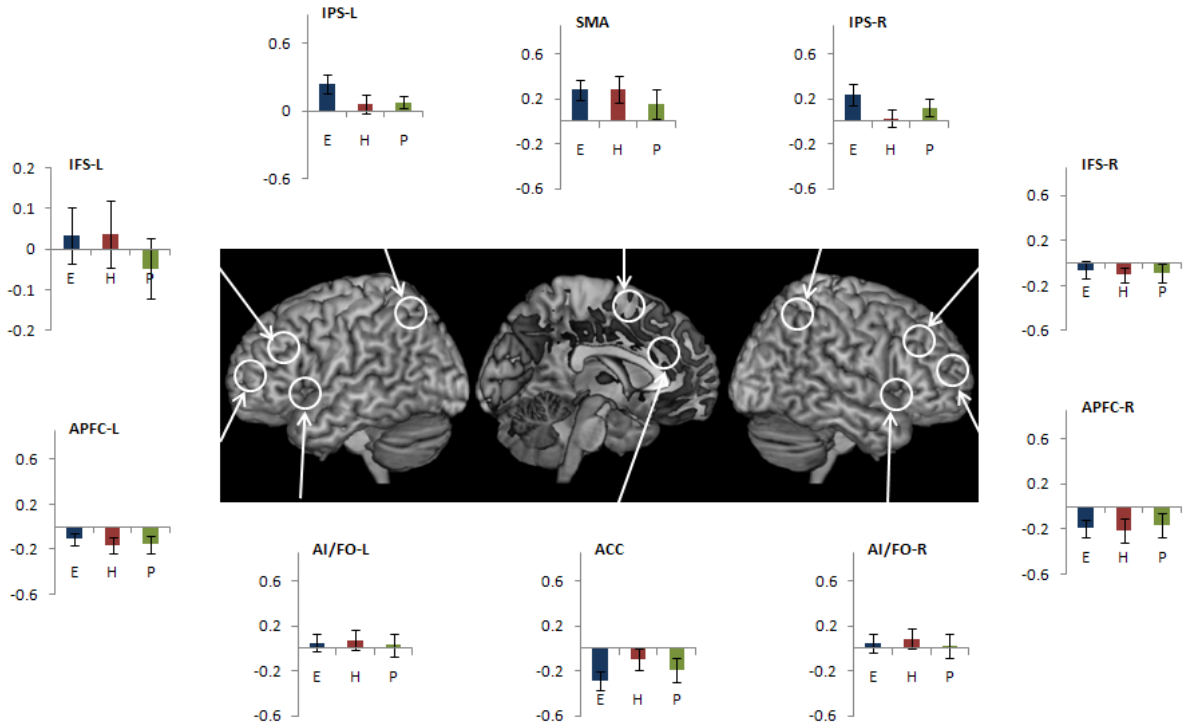

Figure 3.5. None of the brain regions showed differential activity for the different kinds of intervals. Amongst MD regions, IPS showed greater activity for the intervals with low working memory load. (E: Easy; H: High current load; P: High prospective load)

First Task Block

The first task blocks were categorised as 'easy' (I and II), 'hard' (III and IV) and 'easy with high prospective load' (V and VI). The effects of the two kinds of cognitive load (prospective and current) were assessed separately.

Easy and hard first task blocks (I & II and III & IV) differed in the working memory content (one number vs. two numbers & an arithmetic operator, in arithmetic blocks; one word vs. two words, in word blocks) and complexity of computation (one step operation vs. two step operation). Both contributed to the enhanced cognitive load in the hard block.

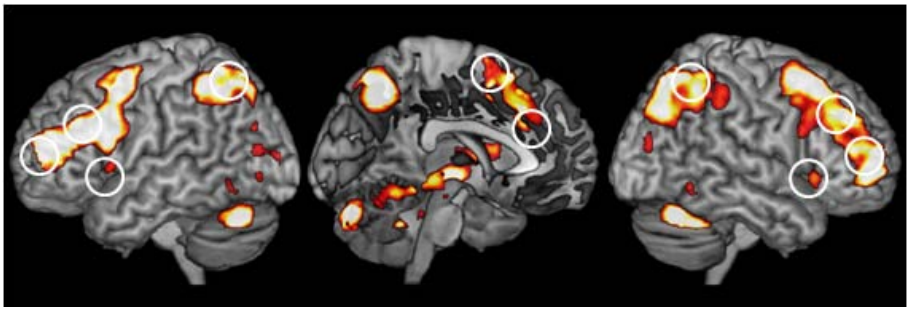

Figure 3.6. Brain render showing areas with greater activity during the hard first task blocks compared to the corresponding easy blocks. These included lateral prefrontal cortex, insula, pre-SMA and ACC, lateral parietal cortex and the precuneus. Lateral prefrontal and parietal activities were centred on the MD regions (circles)

The hard first task blocks were contrasted against the corresponding easy blocks. Greater activity during hard task blocks was found in bilateral fronto-parietal regions (fig. 3.6; peak coordinates in Appendix, table A3.1). Bilaterally in prefrontal cortex, these areas extended antero-posteriorly along the inferior frontal sulcus, from precentral regions to the anterior lateral prefrontal cortex, along with insula; medially, areas between anterior cingulate to the pre-supplementary region were involved. Bilaterally in the parietal cortex, areas with greater activity were centred laterally on the intraparietal sulcus (and extended into superior and inferior parietal lobules and into the middle occipital regions) and medially on the precuneus. Non cortical areas that showed greater activity for difficult blocks included lateral cerebellar cortex, anterior thalamus, and parts of basal ganglia.

In the second contrast the two kinds of easy blocks differing in their prospective load were compared. No brain region showed greater activity for task blocks with greater prospective load even at a liberal threshold of uncorrected $p < 0.001$.

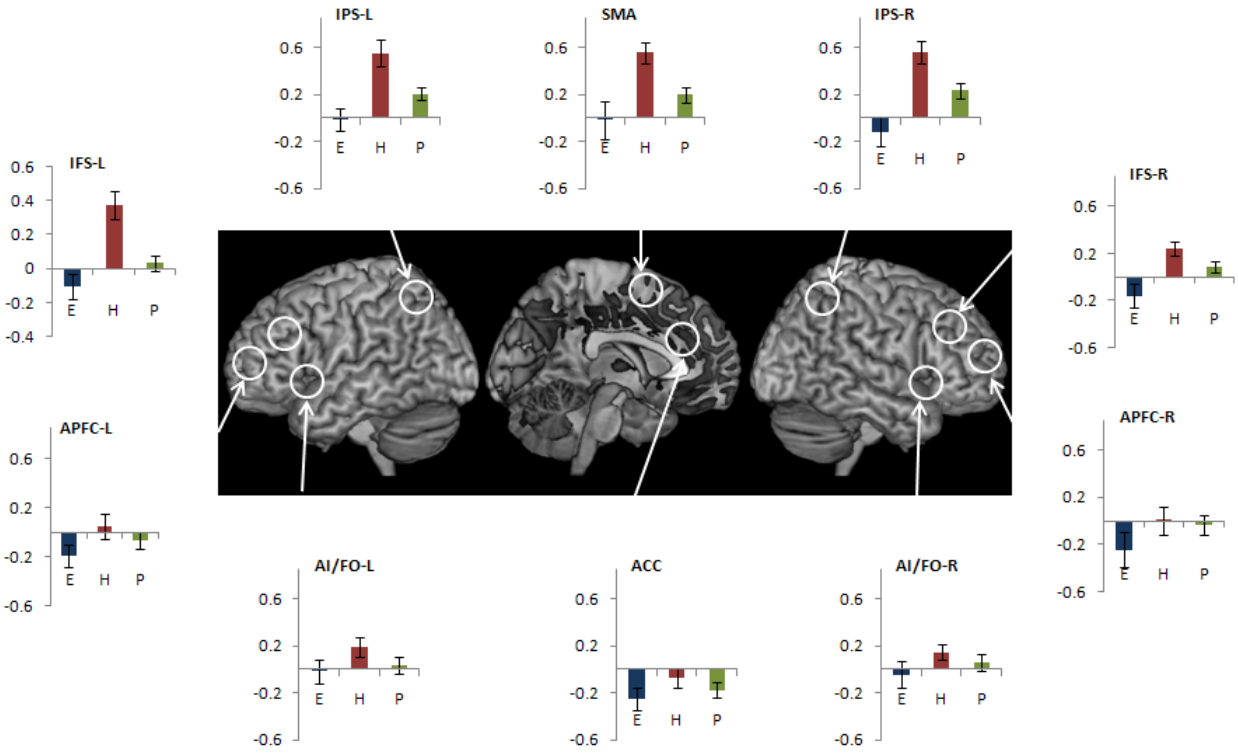

Figure 3.7. Estimates of activity during the three kinds of first task blocks (E: Easy; H: Hard; P: Easy with high prospective load). The estimates of activity of hard task blocks were significantly greater than that of the easy blocks in all ROIs. This effect was bilaterally similar in all ROIs except IFS, where it was greater in the left; Estimates of activity of easy blocks with high prospective load were numerically greater than that of easy blocks, but reached statistical significance only in IFS-R and IPS-R, and close to significance in APFC-L (p=0.06)

The activities in the MD regions during the three kinds of task blocks were then compared (fig. 3.7). A repeated measures ANOVA with load and ROIs as factors showed a main effect for both ($F_{(9,108)} = 8.8$, $p < 0.001$ and $F_{(2,24)} = 14.6$, $p < 0.001$, respectively), along with a significant interaction between them ($F_{(18,216)} = 7.1$, $p < 0.001$), suggesting that the effect of load varied across the different regions. A second model investigated the different ROIs separately to look at the effect of cognitive load (easy, hard, and easy with high prospective load) and laterality (for bilateral ROIs only). All MD regions showed a significant effect of cognitive load ($F_{(2, 24)} > 5.9$, $p < 0.03$, for all ROIs), which was not affected by the laterality in any ROI

except IFS, where the effect was greater on the left ($F_{(2, 24)} = 6.3$, $p < 0.01$). Pairwise comparisons showed that the hard task blocks had significantly higher activity compared to the easy blocks in all ROIs (paired $t(12) > 2.8$, Holm-Bonferroni corrected $p < 0.05$ in all ROIs). Activity for higher prospective load blocks was greater than in the blocks with low prospective load only in right IFS (paired $t(12) = 2.6$, Holm-Bonferroni corrected $p < 0.05$) and right IPS (paired $t(12) = 3.6$, Holm-Bonferroni corrected $p < 0.01$) and nearly so in left APFC (paired $t(12) = 2.3$, Holm-Bonferroni corrected $p = 0.07$).

Second interval

A comparison of the estimates of activity preceding the easy and hard second task blocks was made. These intervals like the case in the first intervals (see above) differed in the amount of information maintained in working memory. Note that this interval was in the second phase of the trial and so did not have any prospective load.

Similar to the case in the interval before the first task block (see above), no brain region showed differential activity across the intervals with different working memory load. The estimates of activity in the MD regions during these intervals were not different either ($F_{(1, 12)} < 1.4$, $p > 0.22$, for all ROIs).

Second Task Block

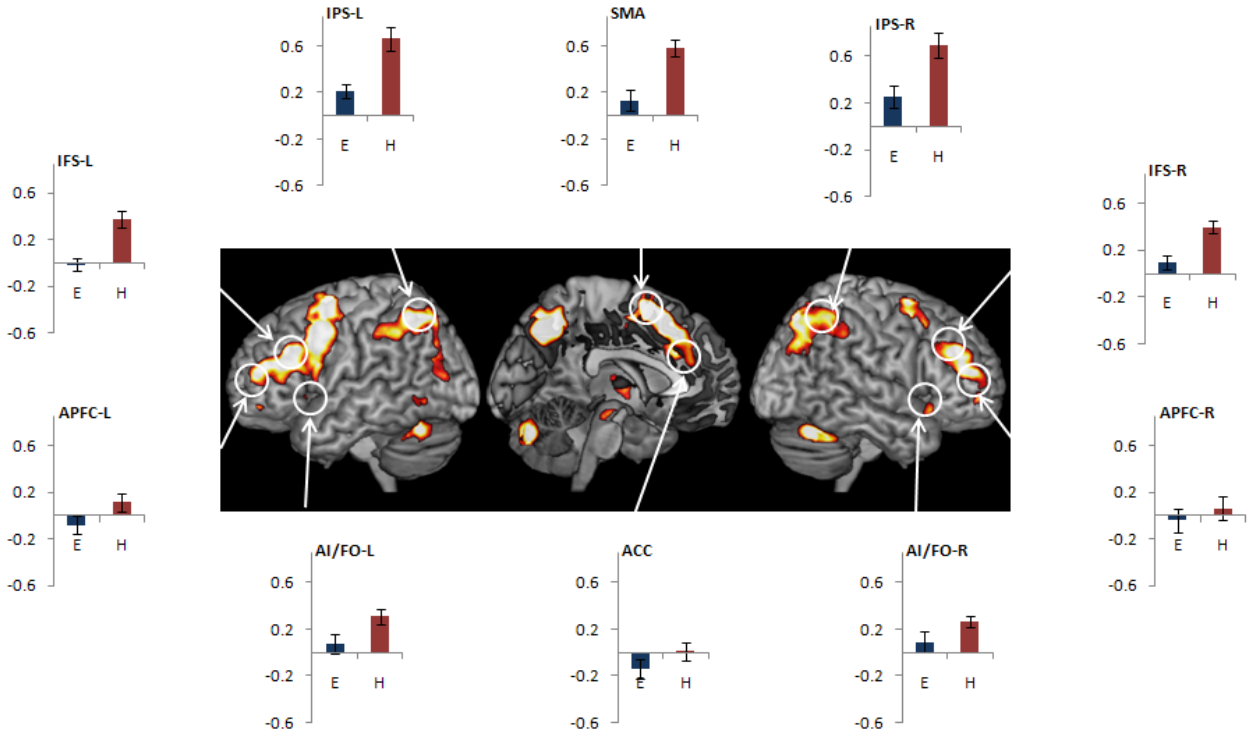

Figure 3.8. The brain render shows areas with greater activity during the hard second task blocks. These included regions along the IFS and IPS, pre-SMA, ACC and precuneus. The plots show the estimates of activity during the easy and hard second task blocks. Estimates for the latter were significantly higher in all ROIs.

Similar to the first task block, the hard second task blocks were contrasted against the easy ones (fig. 3.8). The results were similar to the effect of current task difficulty in the first block and the same set of brain regions showed greater activity for increased cognitive load – bilateral prefrontal regions along the inferior frontal sulci, insula, pre-SMA and ACC; bilateral parietal regions around the IPS, including parts of IPL and angular gyrus, and precuneus; parts of lateral cerebellum, thalamus and basal ganglia. Further, as was the case in the first task block, all of the MD ROIs showed a significant effect of the current task difficulty ($F_{(2, 24)} > 9.5$, $p < 0.01$, for all

ROIs). This effect of difficulty was significantly greater in the left hemisphere for the anterior insula ($F_{(1, 12)}$ =9.8, p<0.01) and anterior prefrontal ROIs ($F_{(1, 12)}$ =7.3, p=0.02).

The results so far show that hard blocks with greater working memory requirements and more difficult task were associated with greater activity in widespread areas of the fronto-parietal cortices (fig. 3.6 & 3.8) including all MD regions (fig. 3.7). Task blocks that differed only in their prospective information load (i.e. the two kinds of easy first blocks that differed only in their associated prospective load) resulted in minimal differential activity that was limited to right IFS and right IPS. Interestingly, the pre-block intervals with similar working memory disparity as the task blocks did not show any difference in associated fronto-parietal activity, which suggests that the activity seen in right IFS and IPS could have been due to the requirement to maintain the information and carry out a different task concurrently.

3.3.2 Phase

Having established the effect of cognitive load of task blocks on their elicited activity, the effect of the position of the block in the plan of the trial was investigated. As described earlier, trials had two discrete phases – the first interval and the first task block constituted the initial phase and the second interval and the second task block were part of the final phase. Across the experiment these two phases were identical to each other, except that the working memory load of the first phase was higher, since subjects additionally maintained information pertaining to the second task block.

Behaviour

It was seen earlier that the behaviour in both kinds of task blocks were not affected significantly by the phase of the trial in which those blocks occurred.

Imaging

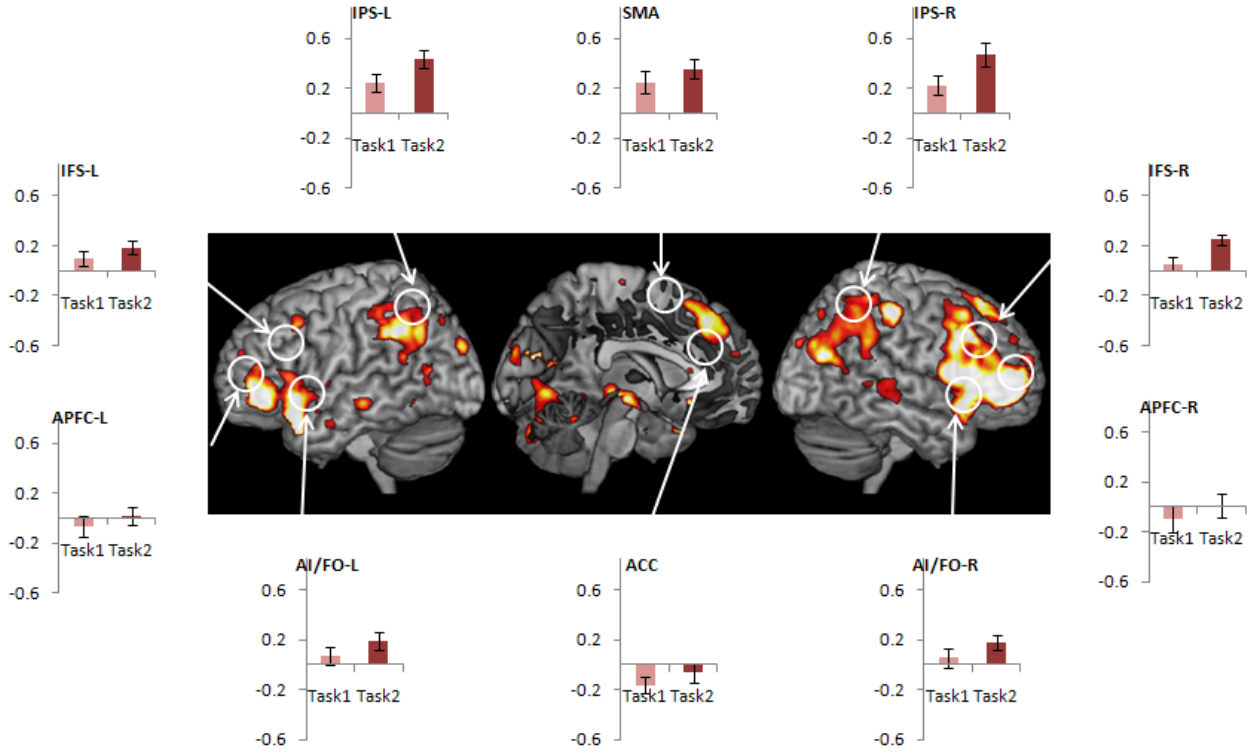

Figure 3.9. Regions showing greater activity for the second task block. Plots show estimates of activity for the first and second task blocks; the estimates for the second task blocks were higher in all ROIs.

The two task blocks were compared to reveal brain regions sensitive to the phase of the trial. No brain region showed greater activity for the first task block compared to the second. This is interesting since in all trials subjects maintained extra prospective information during the first phase that was to be used during the second phase. On the contrary, the reverse contrast showed widespread regions with greater activity for the second block (fig. 3.9; peak coordinates in Appendix,

table A3.2), which included the right prefrontal cortex (right inferior frontal gyrus especially the frontal operculum/anterior insula, lateral anterior prefrontal cortex, premotor regions, posterior regions of superior frontal gyrus extending medially into superior medial regions and parts of dorsal ACC), left prefrontal cortex (inferior frontal gyrus especially frontal operculum/anterior insula extending on to the lateral anterior prefrontal region, posterior portions of middle frontal gyrus and superior medial frontal gyrus), bilateral parietal cortex (inferior parietal lobule, supramarginal and angular gyri) and anterior parts of the superior temporal gyrus.

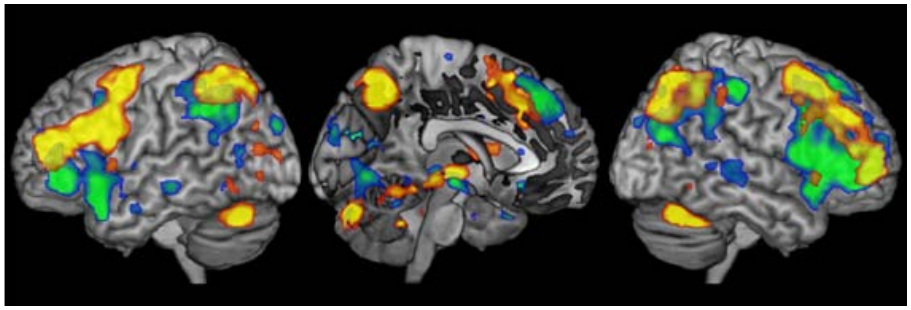

Figure 3.10. The contrasts – Hard vs. Easy blocks (red-yellow), and Second vs. First blocks (blue-green) overlaid together for comparison; in lateral frontal and parietal regions the peak voxels of the former contrast were more dorsal compared to the later. The two overlapped in anterior prefrontal cortex, anterior insula and inferior parietal lobule.

This pattern of activity was qualitatively different from the pattern of activity in response to increased task difficulty (fig. 3.10). In the lateral frontal and parietal cortices, the peak voxels of the former were ventral to the latter. In the prefrontal regions, for example, the peak voxels showing phase related activity were centred on the inferior frontal gyri and parts of frontal operculum extending on to anterior prefrontal regions, whereas the regions showing increased activity with task difficulty were centred on the MD regions - IFS and extending along its axis to the anterior prefrontal regions. Prefrontal activity for both contrasts converged on the anterior prefrontal cortex and anterior insula.

While the pattern of activity showing the effect of phase was different from the MD pattern, the individual MD ROIs were affected by the phase of the trial (plots in fig. 3.9), being significantly greater for the second task block. Repeated measures ANOVA with the phase of the block and laterality of the ROI (for bilateral ROIs) as factors, showed a main effect of the phase in all MD regions ($F_{(1, 12)} > 9.5$, $p<0.01$ for all ROIs). While all ROIs showed the effect of phase, individual ROIs differed in the extent of this effect. A comparison of the effect of phase across different ROIs showed a significant interaction between the two ($F_{(9,108)} = 5.5$, $p<0.001$). Further, the effect was bilaterally symmetrical except in the IFS, where the difference was greater in the right hemisphere ($F_{(1, 12)} = 16.8$, $p<0.001$). These results show that the effect of the phase of the trial on MD regions was similar to that of cognitive load and affected the activity in all MD regions. However, there was one interesting difference. The effect of cognitive load was greater in the left prefrontal ROIs, whereas the effect of the phase was greater in the right IFS.

The activity during the first and the second intervals were compared. Recall that the temporal separation between them was similar to that between the two task blocks. Both the whole brain contrast and the ROI analysis of the MD regions ($F_{(1, 12)} < 1.1$, $p>0.3$) showed null results (fig. 3.11), which suggests that there was no gradual increase in neural activity with the temporal progression of the trial and the difference in activity found between the two task blocks was likely to be due to increase in task-related activity during the final phase of the trial.

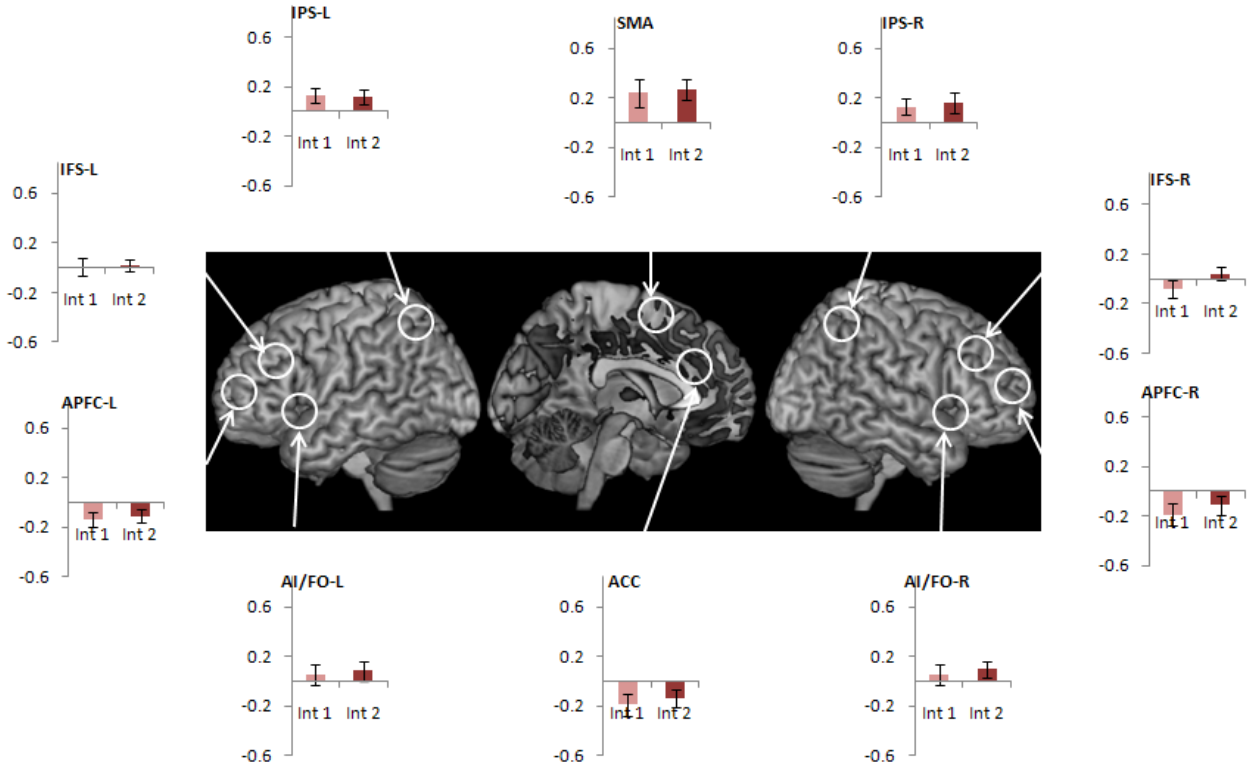

Figure 3.11. Comparison of activity during the first and the second intervals. No brain region showed any difference in the level of activity during these intervals.

End of task blocks

Finally the activities related to the end of the two task blocks were compared. The end of the first task block was related to a task switch (Monsell, 2003), wherein a new task set relevant to the second task block was configured, whereas the end of the second task block initiated a period of rest. However, from the perspective of their role in the plan of the trial, the completion of the first block completed a subtask, whereas the completion of the second block completed the whole plan (ref chapter 2).

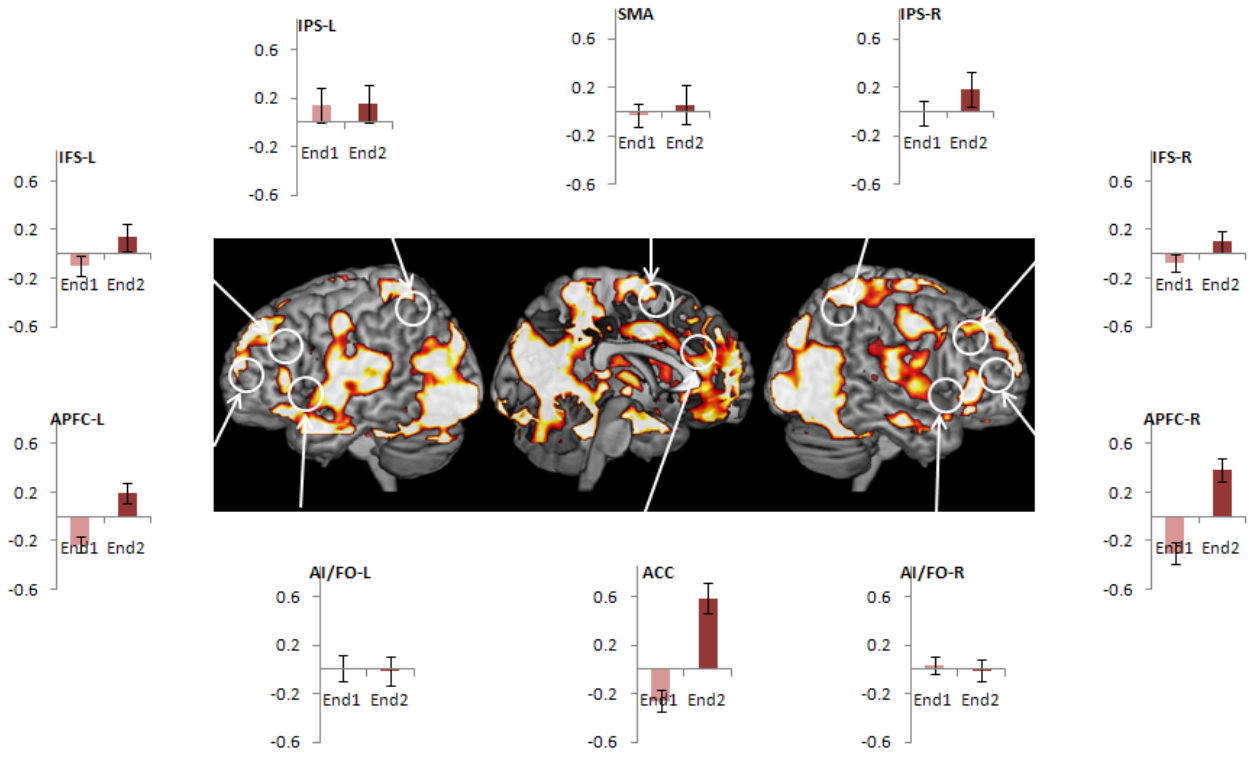

Figure 3.12. Brain regions with greater activity at the completion of the second task block. Plots show the estimates of activity for the ends of the first and second task blocks; the estimates for the second task block were greater in most MD ROIs, but the difference reached statistical significance only in APFC and ACC.

In whole brain analysis, no region showed greater activity for the end of the first task block compared to the second, even at a liberal threshold of uncorrected $p < 0.01$. On the contrary, widespread regions showed greater activity for the completion of the second task block than the first (fig. 3.12). These included bilateral precentral, inferior, and superior frontal gyrii, anterior and middle cingulate cortices, medial and ventromedial prefrontal cortex; bilateral precuneii, posterior cingulate and regions around the temporo-parietal junction in parietal cortex; parts of inferior temporal cortex and middle and superior temporal gyri; and most of the occipital cortex, except for posterior most parts of calcarine and lingual gyri. Apart from these cortical regions – bilateral hippocampus, parahippocampus, amygdala, basal ganglia

and cerebellar regions also showed greater activity bilaterally for the end of the second task block compared to the first.

Some of the regions above have been shown to be more active during periods of rest (medial frontal and parietal cortex, lateral temporal regions, hippocampii and parahippocampii). However, this as an immediate explanation seems less plausible for other regions (fig. 3.12) like inferior frontal gyrii, anterior and medial cingulate, precentral gyrus, early visual and subcortical regions.

Estimates of activity related to the end of the second task block were higher in most ROIs (fig. 3.12), but this difference reached significance only in APFC ($F_{(1, 12)}$ = 16.0, $p<0.001$) and ACC ($F_{(1, 12)}$ = 17.2, $p<0.001$). In APFC this difference was significantly greater on the right ($F_{(1, 12)}$ = 5.4, $p=0.03$). Interestingly this pattern of activity in the MD ROIs for the completion of the trial matches the results of earlier experiments (chapter 2) that showed APFC (right > left) and ACC to be most consistent in showing greater activity for task or plan completion. Indeed compared to their level of activity at the end of the second task block, earlier task stages had resulted in far less activity in these regions (compare figs. 3.9, 3.11, 3.12).

3.4 Discussion

The above results show that the task elements in the final phases of a trial are associated with greater activity in widespread fronto-parietal regions compared to those in the initial phases; this effect was also present in those specific regions of these cortices where activity has been shown to reflect task complexity and cognitive control requirements (Duncan, 2010). Indeed the effect of phase in the MD regions was as widespread as the effect of task difficulty/cognitive load.

3.4.1 Cognitive Load

Hard task blocks showed greater activity in the extended regions of fronto-parietal cortices. In the lateral prefrontal cortex, these regions extended antero-

posteriorly from the premotor to the anterior prefrontal cortex (fig. 3.6 & 3.8). It is interesting to compare this result with a popular view that progressively anterior regions along the antero-posterior axis of the lateral prefrontal cortex are sensitive to hierarchically higher aspects of cognitive control (Badre, 2009). Badre and D'Esposito (2007) showed more posterior frontal activity for simple one-level decisions (e.g. red > left key) than for two-level decisions in which one stimulus feature indicated how the other should be processed, and more anterior frontal activity for higher level decisions, such that anterior prefrontal activity occurred during four-level decisions. Similarly, Koechlin et al., (2003) showed an increasingly anterior locus of prefrontal activity when increasingly greater extents of temporal context determined how a stimulus should be interpreted (for e.g. context of a trial vs. context of the block of trials). In our experiments the hard blocks producing activity in the entire length of the same antero-posterior axis of lateral prefrontal cortex, differed from the easy blocks in working memory load (2 items vs. 4 items) and in the number of computational steps (one vs. two). Since the pre-block intervals having similar working memory disparity did not differ in their elicited activity, it is likely that the difference between hard and easy blocks was related to the complexity of computation performed. Arguably, these blocks did not differ in the levels of decision hierarchy needed to reach the solution (c.f. Badre and D'Esposito, 2007; Christoff et al., 2007) or in the levels of temporal contexts across which the information was to be integrated (c.f. Koechlin et al.,, 2003). This raises the possibility that the antero-posterior shift of prefrontal activity seen in these studies were just related to the greater number of computational steps and were not specific to the hierarchy of the control structures per se.

Intriguingly, we did not find any effect of working memory load alone. Fronto-parietal activity during the intervals preceding the task blocks did not vary with the amount of information maintained. The current results confirm a claim made in the last chapter, that the anterior prefrontal activity is not related directly to the

requirements of maintaining information while carrying out an unrelated task as has been hypothesised by earlier studies (Koechlin et al., 2007; Pisapia et al., 2007; Reynolds et al., 2009). Although the first task blocks had prospective information load, while second blocks had none, no region showed greater activity for the first task block compared to the second. Amongst the first task blocks those with greater prospective load only showed a minimal increase in right IFS and IPS. Further, anterior prefrontal activity was found in hard blocks compared to easy blocks, although they did not differ in their prospective information load.

3.4.2 Phase

The most significant result of the current study was that the phase of the plan in which the task blocks occurred determined their elicited neural activity in the fronto-parietal cortex. Greater activity for the second task block was most prominent in frontal operculum/anterior insula, anterior prefrontal cortex, medial prefrontal (medial superior frontal and dorsal ACC), and IPS extending into IPL. Interestingly, some of these regions viz. APFC, AI/FO, ACC were found by Dosenbach et al. (2006) as core regions showing sustained activity throughout a task block across a number of different experiments. Evidence of distinction between the two hemispheres was also discernible. Left prefrontal ROIs showed greater effect of task difficulty (IFS, AI/FO, APFC), whereas the effect of phase was greater in right IFS.

What is different between the first and the second task blocks? The first possibility is expectation of completion/reward. It is interesting in this regard to compare the above results with the phenomenon of reward expectancy. Tasks requiring subjects to perform a number of trials to achieve a discrete reward show that the reaction time and error rates decrease for trials closer to the reward, with the last trial yielding the fastest and most accurate response (Shidara and Richmond, 1998; Sohn and Lee, 2007). Likewise, the associated neural activity in prefrontal,

parietal, temporal and striatal neurons are seen to increase as trials get closer to the reward (Bowman et al., 1996; Shidara and Richmond, 1998).

Reward expectancy, however, seems unlikely to be the explanation, primarily because there was no reward or achievement to be expected. Further, beyond a general instruction for being accurate and fast, the task instructions did not mention any end to be hoped for; no explicit scoring was done or mentioned. The profile of behavioural performance does not fit this picture either. Reaction time in both phases did not decrease after the second stimulus, whereas studies on reward expectancy show a progressive decline in reaction time till the end. Neither were the RTs different across the two phases. Furthermore, there was no difference in activity in between the two intervals preceding the two blocks, even though the temporal difference between them was similar to the temporal difference in between the two blocks. This also argues against a general increase in activity through the length of the trial and is evidence that the effect of phase was specific to the task blocks.

A second possibility is that the decrease in the cognitive load that happens in between the first and the second block is the actual cause of increase in activity. Counterintuitive as this is in the light of the prevailing view about the relation between fronto-parietal activity and the cognitive load, it certainly is a possibility. The experimental design provides with a means to test this. If some fronto-parietal regions increase their activity only consequent to a decrease in task load that happens between the first and the second task block, it would be expected that an increase in the same kind of task load during the first task block would decrease the resulting fronto-parietal activity. However, contrary to this, increasing the prospective load during the first task block lead to an increase in activity (fig 3.7), thus, showing that decrease in task load in and of itself cannot be an explanation for the increase seen in fronto-parietal activity across the two phases.

A third possibility is that a decrease in the representational load pertaining to the structure of the episode, leads to an increase in activity elicited by the events of the later phase. During the first phase, in addition to the representations related to the current phase, subjects have to additionally maintain representations pertaining to the later phases. While in the later phases, those pertaining to the earlier phases can be dismantled, freeing up the representational space, which can then be utilised by the ongoing task events. This possibility suggests a kind of neurocognitive storage resource, loading which does not lead to an increase in fronto-parietal activity.

This chapter thus adds onto the findings of the previous chapter that the context of the plan of the behaviour affects the fronto-parietal activity related to any task event or task episode within it. While the previous chapter showed the effect of the hierarchical position of the task episode completed on the elicited fronto-parietal activity, the current results show the effect of the phase in which the task events lie. However, the current experiment leaves it unclear if the effect seen is specific to the final phase, or the increase in activity would be a feature of any subsequent phase when compared to the previous phase. This issue is investigated in the next chapter.

CHAPTER 4

Higher Level Representations & Organisation of Task Episodes

4.1 Introduction

Purposive behaviour is organised into a nested hierarchy of subtasks and tasks (Miller et al., 1960). How this organisation is achieved is an open question. It was speculated in earlier chapters that the behavioural hierarchy stems from the hierarchical organisation of relevant representations in the mind such that abstract representations such as task or goal descriptions subsume and organise more concrete representations such as action plans and rules that in turn organise even more concrete representations like those of discrete actions (for a similar account, see Cooper and Shallice, 2000). Other accounts (Botivinick and Plaut, 2004), on the contrary, have tried to show that such hierarchical organisation of relevant representations are epiphenomenal to the output of action control systems which can be instantiated as a simple recurrent network that does not require hierarchically organised set of representations (see also Botivinick, 2009).

One of the ways to test the two accounts is to see if task episodes with identical lower level elements are organised differently when the organisation of their higher level representations is changed. Consider a hypothetical scheme of making tea:

Boil water → Add tea → Brew → Add sugar → Add milk

Will the nature of the higher level representation under which this episode is organised affect the nature of its organisation?

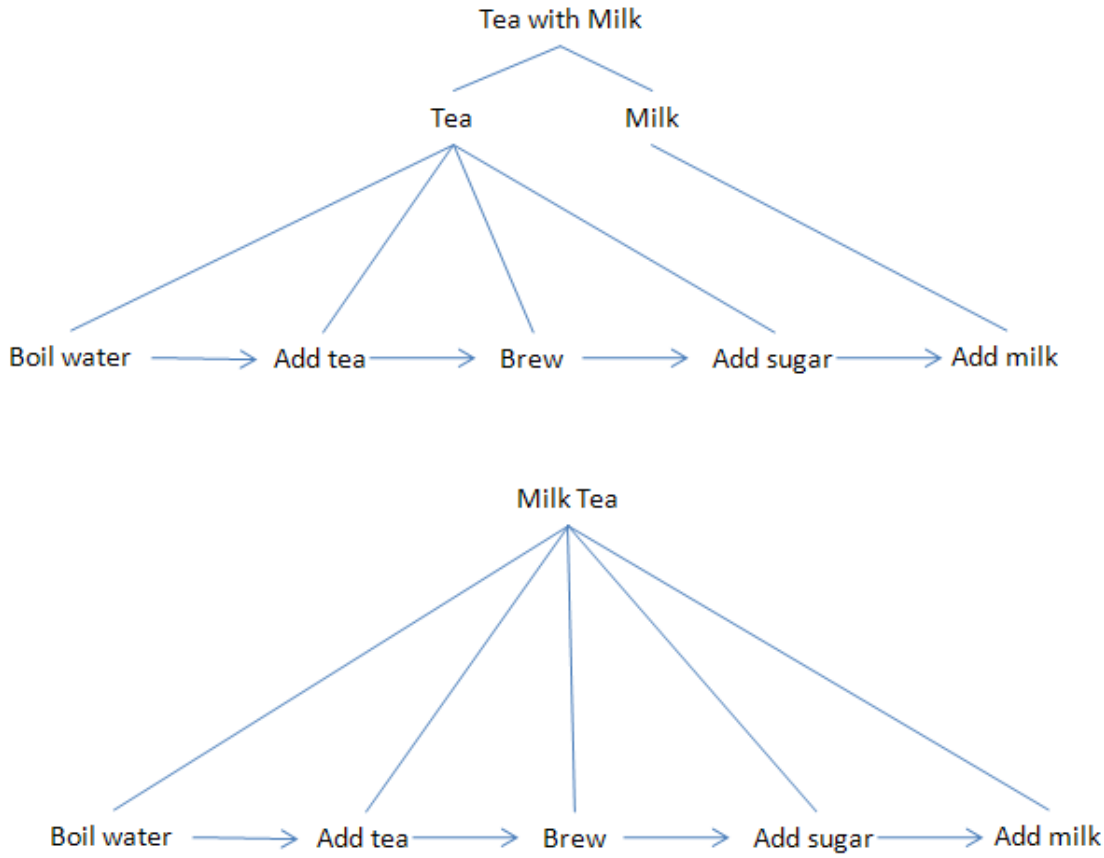

Figure 4.1. Proposed organisation of a hypothetical sequence of actions for preparing tea under two different higher representations, if the nature of latter was causal to the details of organisation (see text).

Consider the representation 'tea with milk' that emphasises the distinction of the last step ('add milk'), versus the representation 'milk tea', which de-emphasises this distinction. Will the lower level sequence of task elements be organised differently in the two cases? Fig. 4.1 depicts the two potentially different ways of organisation if the higher level label affected the organisation. On the other hand, if the properties of higher level label had no causal role, the organisation of the episodes would be identical in the two cases.

Results described in the chapter 2, furnished evidence in support of the former view by showing that the way the task was described, did affect the organisation of the task episode. Although subjects carried out four identical target detections, the first three had been described to them as being part of 'one task', and the fourth as part of the 'other task'. The fronto-parietal activity at the third target detection was higher than the first two, and was similar to the case when the third target detection did actually complete the task i.e. when there were only three targets to detect.

That the nature of task instruction determines its execution is a truism. But the way the same task instructions are organised by the subjects affects their performance. This is especially evident in complex tasks. Duncan et al. (unpublished data) instructed subjects to do a task with multiple contingent rules (chapter 1, Fig. 1.2). Subjects are shown a stream of pairs of stimuli (letters or numbers) presented simultaneously on the two sides of the centre of a computer screen. The last three pairs of stimuli of the stream were preceded by arrow heads pointing to left or right. The subjects are to follow a series of instructions in a particular sequence – attend to left (or right), read aloud the letters on the attended side, add pairs of numbers but ignore the asterisks, after the appearance of '>' read aloud the letters on the right, if '<' appears, read aloud the letters on the left side. It had been shown by the authors earlier (Duncan et al., 2008) that the last part of the task (taking note of the arrow head and attending to the requisite side) tends to be frequently neglected. In the current study, the authors found that the way subjects were asked to recall the instructions, affected their performance. Subjects that were free to recall in any way they chose showed poorer performance than subjects that were explicitly instructed to organise the recall into a series of bullet points. Note that the instructions were the same in both groups, the difference lay in the way they were made to organise the instructions. Plausibly, the effect of different kinds of recall was the different

pattern of organisation of task instructions (higher level representations), which affected the assembly of task control structures.

In the current task, the subjects were biased to organise the higher level representations into specific ways, then the organisation of the resulting task episode was judged from their behavioural profile and from the elicited neural effects of task organisation described in earlier chapters. Each trial in the current experiment, consisted of a sequence of eight steps (called subtrials), the rule relevant to them was cued to the subjects as an eight letter string composed of letters 'A' or 'B' (e.g. AABBBBAA). On the subsequent eight subtrials, responses were made according to the rule represented by the identity of the letter at respective position in the memorised string. 'A': categorise the letter as vowel/consonant; 'B': categorise the number as even/odd. For example, if the string was AABBBBAA, rule A was relevant on the first two subtrials, followed by rule B on the next four, followed by rule A again in the last two subtrials.

To bias the organisation of the representation of this sequence, subjects underwent a 20 minute pre-scan practise session during which they performed trials with four subtrials guided by four letter strings e.g. AABB. It was hence expected that in the main scanning session the subjects would be biased towards representing the eight letter string as two chunks of four letters (e.g. AABBBBAA as AABB and BBAA).

Two kinds of strings were used in the scanning sessions. In one (trial type 1), the first and the second group of four letters were different e.g. AABBBBAA; they were the same in trial type 2, e.g. AABBAABB. It was hoped that these two kinds of strings would be represented differently. The sequence in trial type 1 required that all eight elements be kept in working memory from which individual letters at the relevant position were to be accessed. However, in the case of trial type 2, there was no need to remember all eight elements, and it was plausibly more parsimonious to

remember just four elements and then repeat them. If the properties of higher level representations causally determine how the behavioural episode is organised, it can be expected that these two trial types will be differently organised.

In trial type 1, since the entire sequence had to be represented as one unit from which individual elements were to be accessed, it was expected that the entire trial would be organised strongly as one episode albeit with some evidence of further organisation in two subtask episodes due to the bias created by the pre-scan session (Fig. 4.2a). In trial type 2, on the other hand, there was no need for the entire sequence to be represented, rather just representing the repeating four elements was enough, so this trial type was expected to be strongly chunked as two sub-episodes (Fig. 4.2b). These expectations made specific predictions about the profile of behaviour and that of elicited fronto-parietal activity in the two trial types.

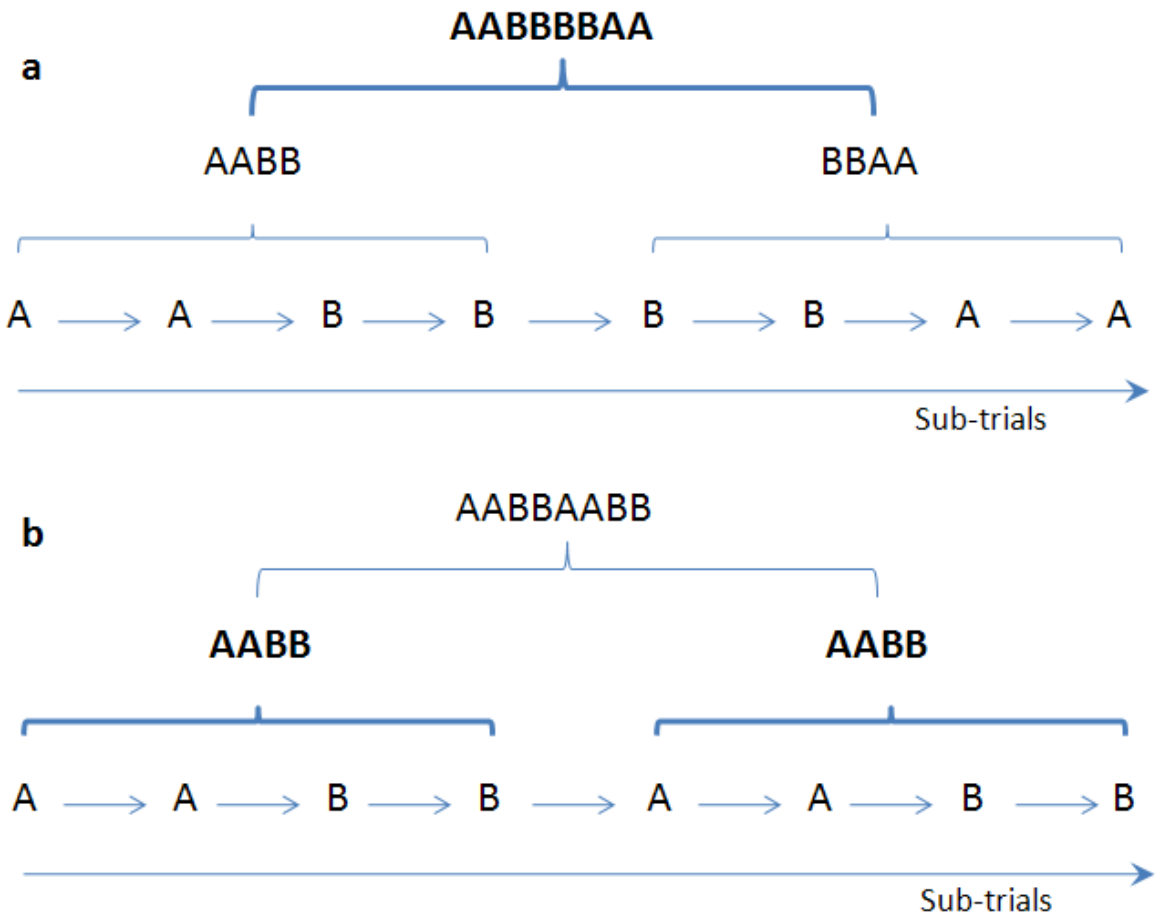

Figure 4.2. Possible schema of trial type 1 (a) and trial type 2 (b). In trial type 1, the entire trial was expected to be organised as a unified task episode with a strong representation of the highest sequence level representation, and with relatively weak evidence of the distinction between the two chunks; trial type 2, on the other hand, was expected to be organised more as two task episodes and the sequence representation was expected to be relatively weaker, while the chunk representations were expected to be relatively stronger than in trial type 1.

The behaviour at the fifth subtrial, which marks transition across the two chunks, has different predictions depending upon the kind of organisation of the two trials. Transition across hierarchically organised episodes of behaviour has been shown to require additional time, compared to transits within the episode (Lien and Ruthruff, 2004; Schneider and Logan, 2006 and 2007). Therefore, the greater the hierarchical difference between the two episodes, the greater would be the cost of transiting across them. So if the subtrials 1 to 4 and 5 to 8 in trial type 2 are

organised more as separate task episodes, the reaction time (RT) on the fifth subtrial would be greater than in the case of trial type 1 where the two chunks are represented more as part of the same task episode.

On the other hand, if the organisation of the two trial types was the same, the behaviour at the fifth step would depend upon how the cued sequence has been represented. If sequence has, nonetheless, been chunked then the fifth RT in trial type 1 would be expected to be higher than the corresponding RT in trial type 2 since it involves recalling from a different chunk and creating a new behavioural episode, whereas the relevant chunk remains the same in trial type 2, and the previous episode needs to be repeated. This assumption is supported by the findings of Schneider and Logan (2006 and 2007). They had subjects memorise two distinct sequences (called α and β) of four task units each. On every trial they cued the sequence and the position (e.g. α3) from which subjects were to access and execute the relevant task. They found a distinct switch cost whenever a different sequence was cued compared to the previous trial. In a different study, they had subjects execute these sequences in the order α α β β. Thus, the sequences were alternately repeated and switched. They found a distinct switch cost whenever the sequences switched. Finally, if the sequence was not chunked in either, the fifth RT would not be different from other RTs in both trial types.

In earlier chapters it was shown that completion of defined task episodes correspond to an increase in activity across many fronto-parietal regions, and task phases closer to such completions have greater fronto-parietal activity compared to earlier phases. It was hence predicted that if trials are organised as two sub-episodes, fronto-parietal activity at the fourth subtrial marking the completion of the first episode would be higher compared to other subtrials. Further, if this organisation as two sub-episodes is greater for trial type 2, then the fourth subtrial activity in it would be higher than in case of trial type 1.

What would be the pattern of phase effect across the two trial types? From the results of the previous chapter, it was not clear if the increase in activity is limited to the final phase or includes any subsequent phase compared to the earlier one. If the increase is pervasive and stepwise across the phases of the plan, trial type 1, being a more unified episode, can be expected to show a pervasive stepwise increase in activity along the eight subtrials, while trial type 2 may show separate increase for the two sub-episodes i.e. parallel increase along subtrials 1 to 4 and 5 to 8 (Fig. 4.3). If the increase was limited to the final phase, trial type 1 would show greater activity only at the eighth subtrial, while both, the fourth and the eighth subtrials can be expected to have relatively higher activities.

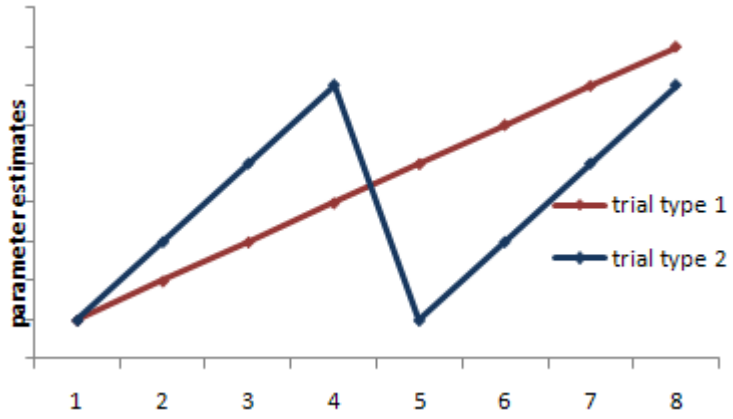

Figure 4.3. Predicted pattern of activity over the eight subtrials of the two trial types if the phase effect was pervasive and not just limited to the final phase. Trial type 1, being a unified task episode, was predicted to show a gradual increase in activity across the length of the trial in fronto-parietal regions sensitive to the phase. Trial type 2, on the other hand, was expected to show separate increase in activity for the first and the second group of the four subtrials.

While the two trial types (*organised* trials) described above were organised through specific learned sequences, a third kind of trials (called *unorganised* trials) also consisting of eight subtrials did not have an organising sequence cued at the beginning of the trial. Instead the relevant rule was cued on each subtrial by the

colour of the margin outlining the stimuli. These trials were identical to the organised trials in all other aspects. Such trials provided an opportunity to test a very speculative thesis. Since subjects had been biased to organise the subtrials in groups of four through the organising sequences, they might organise a series of eight steps in groups of four even in absence of explicit organising sequence. In this case it would be a case of abstract conceptualisation of the episode affecting its organisation. If the activity in some fronto-parietal regions was sensitive to such abstract conceptualisation of the structure of task episode, it can be expected that those regions would show an increase in activity during the fourth subtrial of this trial type. Note that the lack of organisation in case of the *unorganised* trials is only relative to the other two more (explicitly) organised trial types. Otherwise even the *unorganised* trials are organised. Subjects know that these have an explicit start, consist of eight steps and have an explicit end etc.

Experiments in the previous two chapters showed patterns of fronto-parietal activity contrary to what would be expected from the requirements of cognitive load. In chapter 2, the last target detection event did not require the recall of a new target letter and beginning of a new visual search, unlike the previous target detections. Nevertheless, fronto-parietal activity was maximum for the last target detection. In chapter 3, as compared to the initial phases, the final phase had lower working memory load and did not require the maintenance of branched information (Koechlin, 2007) nevertheless activity was higher for the final phase in all areas of the fronto-parietal cortices. The current experimental design allowed for greater exploration of these issues.

Hierarchical task episodes require hierarchical cognitive control (Schneider and Logan, 2006; Lien and Ruthruff, 2004) – higher level control accesses and maintains higher level representations while the lower level control accesses and executes the subtasks. For example, in an experiment by Schneider and Logan (2006)

multi-element chunks composed of two tasks (e.g. AABB, AABBAA) were repeatedly executed (e.g. AABB-AABB-AABB). The RT on the first element of the chunk was always found to be the highest and the degree of increase in RT at the first element was determined by the complexity of the chunk being executed (e.g. it was greater for AABABB compared to AABB), suggesting that the beginning of the chunk required an additional chunk level control, which was greater for more complicated chunks. That this control was hierarchically higher than that required for the execution of the individual elements was evident from the absence of any lower level task switch cost at the first element. For example, when AABB was repeatedly executed (AABB-AABB-AABB) the first element involved a task switch (B to A), but the RT on this was not greater than in the case of the first element when ABBA was repeatedly executed (ABBA-ABBA-ABBA) even though the latter did not involve a task switch. Indeed the RT in the latter case was higher than the former, possibly because the latter sequence was more complicated as it involved a greater number of task switches within the chunk (two in ABBA compared to one in AABB). The higher chunk level control seems to subsume lower level control requirements.

From this perspective, the organised trials required multiple levels of control. Lowest level control to perform the subtrials, chunk level control between the fourth and fifth subtrials to instantiate control structure required for the new chunk, and at the highest level, a control structure associated with the higher level sequence instantiated at the first subtrial. These control levels have a nested hierarchy since the highest sequence level representations (AABBBBAA) are required to access the chunk level representations (AABB or BBAA) from which in turn are accessed the elemental representations (A or B). Regions instantiating these higher order control can be predicted to have phasic activity at the first and fourth (or possibly fifth) subtrials.

Some recent theories on the functional organisation of the prefrontal cortex (PFC) suggest that control at progressively higher levels result in more anterior activity within the PFC (Koechlin and Jubault, 2006; Badre and D'Esposito, 2009). These would predict that the first subtrial would be associated with the most anterior locus of prefrontal activity since the highest sequence level control is instantiated at this point, followed by the fourth (or fifth) subtrial when chunk level control is instantiated, followed by other subtrials.

The scheme of organised trials also made explicit predictions from the perspective of working memory. Working memory load decreased along the length of the trial as the representation relevant to the particular subtrial could be discarded after its completion. Brain regions sensitive to the working memory load can hence be expected to show a stepwise decrease in activity along the length of trial type 1. In trial type 2, such decrease may happen only between subtrials 5 to 8, since only during these can the relevant representations be discarded.

In summary, divergent predictions exist from the perspectives of organisation, cognitive control and working memory load.

The results showed that the two organised trial types were indeed differently organised. Compared to other subtrials, the fifth subtrial had longer RT in trial type 2 than in trial type 1 suggesting that chunking was more intense in the former. Neuroimaging results agreed with this. Relative to other subtrials, the fourth subtrial completing the first chunk elicited greater activity in trial type 2. Most intriguingly, in trial type 3 a number of fronto-parietal regions showed increased activity for the fourth subtrial. These results provide compelling evidence that higher level representations are indeed causal in determining the organisation of the task episode. The eighth subtrial in all three trial types, predictably, elicited the greatest activity across most fronto-parietal regions.

The results also confirmed that the phase effect is not specific to the final phase, rather all subsequent phases elicit greater activity compared to the earlier phases. The exact pattern of this was shown to depend on the organisation of the episode. Trial type 1, organised as a single episode, with poor evidence of chunking, showed most widespread evidence of a monotonic increase in activity across the length of the trial. This was weaker in trial type 2 and weakest in trial type 3. Indeed, the pattern of increase in activity in trial types 2 and 3 were qualitatively different from trial type 1. Instead of a monotonic increase across the length of the trial, these trial types showed parallel increase along subtrials 1 to 4 and 5 to 8, suggesting that the dominant pattern of organisation of these trial types was bi-episodic.

Lastly, no brain region showed greatest activity during the first followed by the fourth (or fifth) subtrials of organised trials, as would have been predicted for regions involved in instantiating hierarchical control. Neither did any region show a decrease in activity across the length of the trial as would have been predicted for regions sensitive to working memory load.

4.2 Methods

4.2.1 Stimulus

The stimulus consisted of a letter and a number placed at the different corners of a square (Fig. 4.4). Responses were made according to one of the two tasks: task A - categorise the letter as a vowel (index finger) or a consonant (middle finger); task B - categorise the number as even (ring finger) or odd (little finger) using the next two fingers. The exact position of the letter and the number within the square changed randomly. The letter and the number were in black Arial font and subtended a visual angle of half a degree. The outlining square was black in colour unless specified otherwise (see below) and had a visual angle of ~ 1.5 deg.

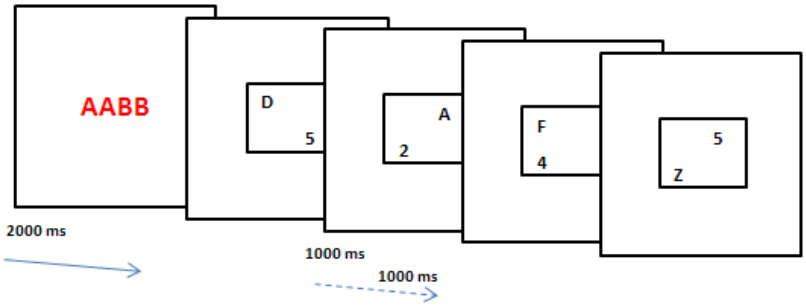

Figure 4.4. Scheme of trials during the first part of pre-scan practise session (see text for details).

4.2.2 Pre-scan

Prior to imaging, subjects did a 20 minute practice session (Fig. 4.4). Trials began with the presentation of a four-letter string (for 2 s) composed of all combination of letters A and B (e.g. ABAB, AABA, BBAA, BABB, AAAA etc.). This letter sequence was to be kept in working memory for the duration of the ensuing four subtrials. On each of the subtrials the stimulus described above appeared for 1s. The identity of the letter in the remembered sequence corresponding to the number of current subtrial determined the task-relevant to that subtrial; e.g. if the sequence was AABB, task A was to be done on the first two subtrials followed by task B on the next two. Inter subtrial interval was 1000 ms giving subjects 2 s to respond. After the fourth subtrial, the next trial began (inter-trial interval (ITI), 1s). Subjects did a total of 80 such trials. This aim of this session, apart from familiarising the subjects with some aspects of the task, was to bias the subjects towards representing four letter strings as a chunk.

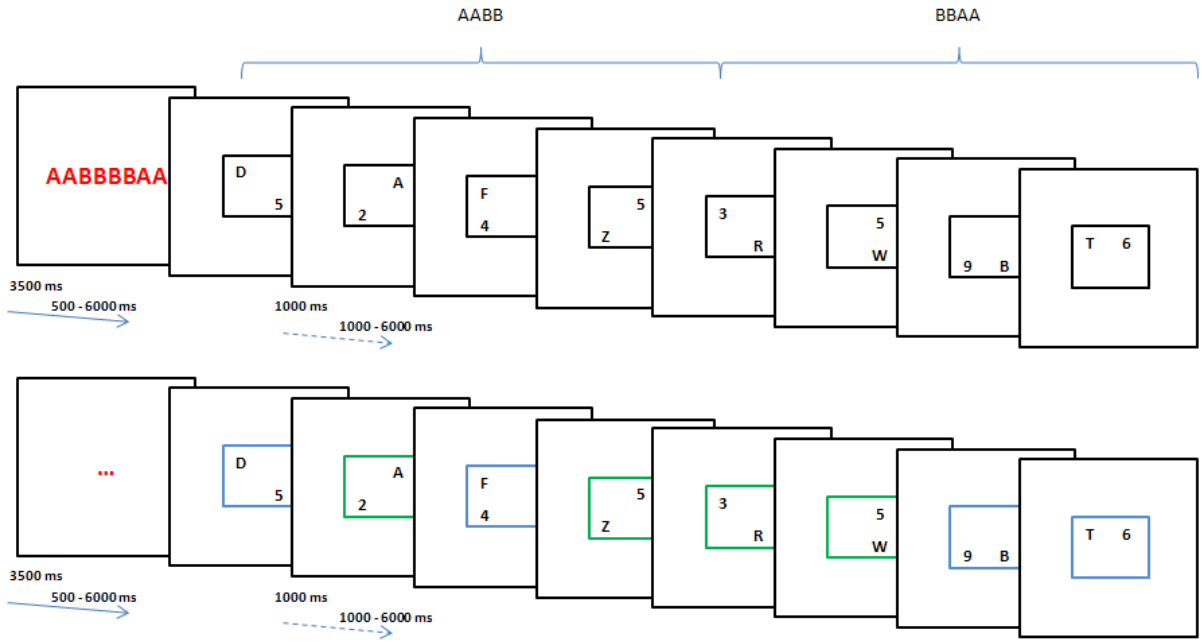

Figure 4.5. Organised (top) trials began with a cue display of the eight letter sequence organising the rest of the trial. On the ensuing eight subtrials, the corresponding letter in the sequence was to be recalled and the subtrial executed accordingly (A: letter task; B: number task). In the unorganised trial (below) no letter sequence was cued, and the relevant task on any subtrial was determined by the colour of the square (blue: letter task; green: number task).

4.2.3 Imaging session

Trials began with the presentation of the relevant eight letter sequence in the organised trials (Fig. 4.5) or with "…" in unorganised trials (presented for 3500 ms). The sequence was made up of random combination of letters 'A' and 'B' (e.g. 'AABABBAB'). This was followed by a series of eight subtrials separated by an average interval of 3.5 s (jitter 0.5 to 6 s). On every subtrial the stimulus described above appeared for 1s. In organised trials (trial types 1 & 2) the colour of the margins of the square was black throughout the eight subtrials, while during the unorganised trials (trial type 3) the colour of the margins signalled the task-relevant on the particular subtrial, blue for task A and green for task B. Since in organised trials subjects had an a priori knowledge of the rule relevant on any subtrial, they could start preparing for the upcoming subtrial immediately after responding to the previous one. In unorganised trials, to allow for the same preparation time, the colour of the outlining

margin changed immediately after the offset of the stimuli of the previous subtrial (and after the offset of the cue for the first subtrial). Valid responses could be made up till 2 s following the stimulus appearance. At the end of the eighth subtrial a blank square with black margins remained on the screen till the beginning of the next trial. Average ITI was 4 s (jittered 2 to 7 s). (See Appendix B.4 for exact instructions).

Trials were of three kinds – organised trials with different four-letter sequences (e.g. AABB and BBAA, presented as 'AABBBBAA'), hence called type 1 trials; organised trials with the same four-letter sequence repeated (e.g. 'ABABABAB'), hence called type 2 trials; and unorganised trials or type 3 trials. These three trial types were delivered in equal proportion. The total duration of the imaging session was 68 minutes, which was divided into four separate scanning runs. Subjects did 30 trials in each of the four sessions.

4.2.4 Participants

15 participants (7 female, mean age = 22.5 ± 3.6 years) were recruited from the MRC-CBU volunteer panel. Participants were right handed and had normal or corrected vision. Informed consent was taken and the participants were reimbursed for their time. The study had the approval of the Hertfordshire Local Research Ethics Committee.

4.2.5 Acquisition

Same as in previous experiments.

4.2.6 Analysis

The fMRI data were analysed using SPM5. Prior to statistical analysis, all EPI volumes were slice-time corrected using the first slice as a reference, and then realigned into a standard orientation using the first volume as a reference. These

realigned images were then normalised into the Montreal Neurological Institute (MNI) space and spatially smoothed using an 8 mm full-width half-maximum (FWHM) Gaussian kernel. During the normalisation stage, voxels were resampled to a size of 3 × 3 × 3 mm. The time course of each voxel was high-pass filtered with a cutoff period of 90 s.

For each trial, the cues and the subtrials [eight x three trial types] were modelled using epoch regressors, of width equal to the duration of the respective events. Each subtrial was modelled from the offset of the previous stimulus (or cue in case of subtrial 1) to the offset of the current stimulus. The stimulus events onsets within each subtrial were modelled with one event regressor of no duration to regress out the common phasic activity related to the presentation of the stimulus. All of these regressors were convolved with a basis function representing the canonical hemodynamic response. Movement parameters and block means were included as covariates of no interest. Parameter estimates for each regressor were calculated from the least squares fit of the model to the data, and estimates for individual participants were entered into a random effects group analysis.

Whole-brain comparisons were performed using paired t-tests on the relevant contrast values from each participant's first-level analysis. Unless otherwise specified, all results are reported at a threshold of $p < 0.01$, corrected for multiple comparisons using the false discovery rate (FDR) method. Coordinates for peak activation are reported using an MNI template.

Apart from the MD ROIs used in the previous experiments, three additional ROIs were created from the three major clusters of prefrontal voxels that showed a significant effect of phase in the previous experiment (chapter 3, Fig. 3.9). These clusters were from the ventrolateral (centered at MNI coordinates 47, 32, 0 and -47,

20, -4) and medial superior regions (0, 37, 42) of the PFC. These ROIs will subsequently be referred to as the phase-sensitive ROIs.

4.3 Results

Behaviour

The hierarchical organisation of a task episode can be inferred from the time taken to access the representation at different points, since it is expected that more time will be needed to access representations from a different level of hierarchy than from the same level (Schneider and Logan, 2006 and 2007; Lien and Ruthruff, 2004). In the organised trials, if the cued sequence was chunked into groups of four, the RT on the fifth subtrial can be predicted to be higher than others (except the first RT). Further, the magnitude of this increase would depend upon the magnitude of chunking.

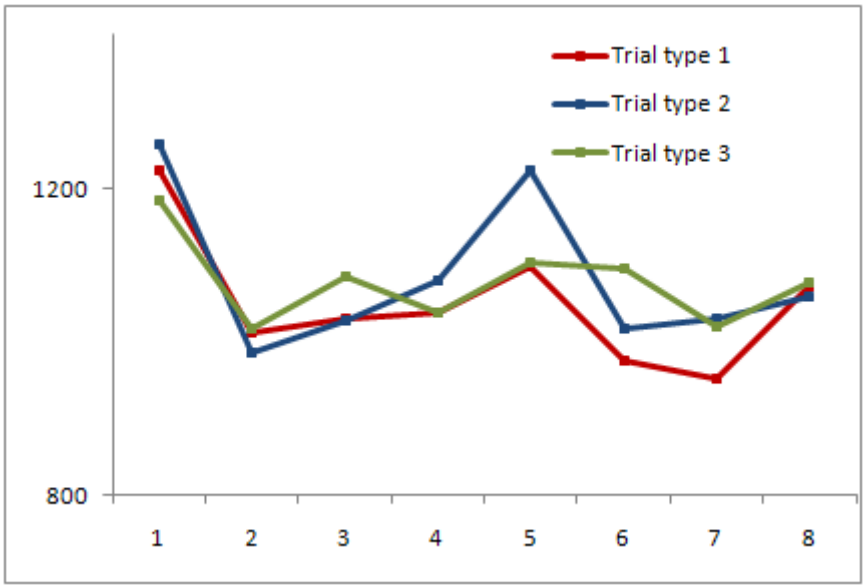

Figure 4.6. Profile of RTs across the eight subtrials in the three trial types. The first RT was highest of the eight subtrials in trial types 1 and 3, but not in trial type 2. Only trial type 2 showed a high RT at the fifth position. As predicted the first RT was greater than fifth only in trial type 1. The fifth RT of trial type 2 was significantly greater than the fifth RT of trial type 1.

Fig. 4.6 shows the profile of reaction times in the three trial types. In trial type 2, the first RT was higher than the other within-chunk reaction times (t_{14} > 4.2, p<0.001), however it was not higher than the fifth RT (t_{14}=0.8, p=0.4). The fifth RT was higher than the other within-chunk reaction times (t_{14} > 4.1, p<0.01), unlike the case in trial type 1. As discussed earlier, this pattern of RTs is strongly suggestive of chunking of first and second groups of four subtrials. The results were different for trial type 1; the first RT was the highest (paired t_{14} > 3.2, p< 0.01, for all comparisons). The fifth RT was not higher than other reaction times, suggesting that chunking was poor in this trial type.

The profiles of the RTs in the two trial types were directly compared to see if the effect of position on the RT was different in the two trial types. This was indeed the case ($F_{(7,91)}$=6.1, p <0.01). Lastly, a direct comparison of the fifth RT on the two subtrials showed that this was significantly greater in the trial type 2 (t_{14} =3.5, p <0.01). The behavioural profile of the two trial types shows that these were organised differently. Indeed the pattern suggests that two chunks of subtrials in trial type 1 were organised as parts of the same episode, while those in trial type 2 were organised relatively independently. As discussed earlier, this profile is consistent with stronger chunking in trial type 2 than in trial type 1.

In trial type 3, there was an effect of position on the RT ($F_{(7, 91)}$ = 3.2, p <0.01), which was driven by the first RT being higher than the rest (t_{14} > 4.0, p<0.001); other RTs did not differ.

Imaging

Depending upon how the representation of the trial episode is organised in the particular brain region, distinct patterns of activity are to be expected. Regions, where the trial representation is organised strongly as two episodes (i.e. has a strong chunk level organisation), can be expected to show greater activity at the completion

of the first chunk i.e. the fourth subtrial (chapter 2). Regions that represent the trial as a single episode would show increased activity only at the eighth subtrial.

Regions that are sensitive to the phase of the task can be expected to show an increase in activity across the sequential phases. However, the exact pattern of activity across the eight subtrials in these regions might depend upon the organisation of the trial, since which subtrials constitute the sequential phases of the *same* episode would get determined by the way the trial is organised. In those organised as a single episode, the eight subtrials constitute their eight sequential phases; whereas in those that are organised as two episodes, subtrials 5 to 8 do not belong to the same episode as subtrials 1 to 4. Regions representing the trial as a single episode may show a monotonic increase in activity across the eight subtrials, whereas those representing it as two episodes will may show parallel increase across the first and the second chunk of subtrials, but without any increase across the two chunks.

The activities between analogous phases of the two chunks were compared. Subtrials 2 and 3 (first chunk) and 6 and 7 were chosen. Subtrials 4 and 8 were excluded to avoid confounds from activity related to the completion of episodes. Regions representing the trial as single episodes may show an increase across them such that activity in subtrial 7 > 6 > 3> 2, i.e. both within-chunk (3>2, 7>6) and across-chunk phase effects (6 & 7 > 2 & 3). Regions representing trials as dual episodes may show 3 > 2 and 7 > 6, but not 6 & 7 > 2 & 3, i.e. only within-chunk but no across-chunk increase in activity. This was tested using an ANOVA which modelled these as sequential phases within the two chunks, and looked at within-chunk effect i.e. 2 and 6 vs. 3 and 7, and across-chunk effect i.e. 2 and 3 vs. 6 and 7.

	4 > int.	8 > int.	8 > 4	(3 & 7) > (2 & 6)	(6 & 7) >(2 & 3)
Single Episode	-	+	+	±	+
Two Episodes	+	+	-	+	-

Table 4.1. Different set of predictions for the two kinds of trial organisation. If the trial is organised as a single episode, subtrial 8, finishing the entire episode, can be expected to elicit maximum activity. Subtrial 4, in such cases, should not differ from other intermediate subtrials. Finally, the activity in some regions can be expected to increase monotonically across the eight subtrials, which can be quantified as greater activity during subtrials 6 and 7 compared to that at 2 and 3. On the contrary, if the trial episode is chunked, both subtrials 4 and 8, completing the two chunks, can be expected to have higher activities than other intermediate subtrials. Further, subtrials within a chunks may show increase in activity (3 > 2 and 7 > 6) but this increase may not be present across the subtrials of the two chunks.

The various predictions are summarised in table 4.1. It is possible that a trial type may be organised in different ways in various regions. One or the other pattern may dominate across these regions depending upon the predominant way the representation of that trial type is organised. Some regions may show a mixture of these two patterns of representation.

4.3.1 Trial Type 1

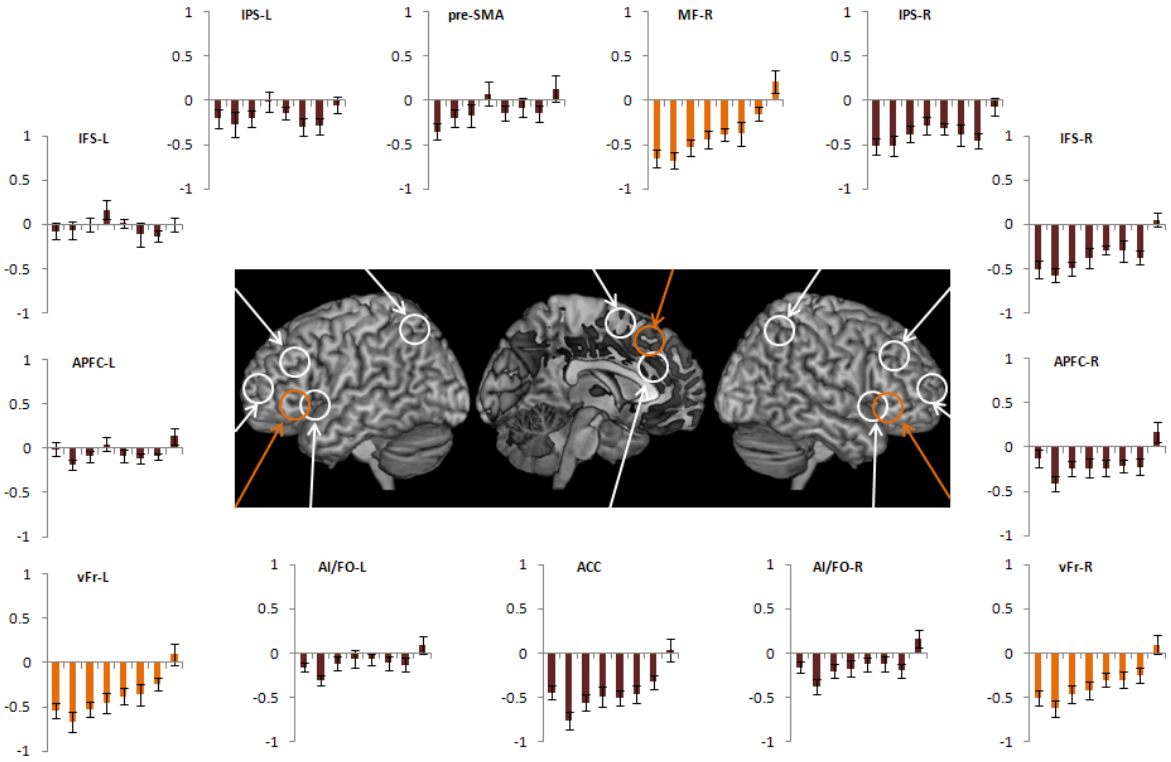

Figure 4.7. Pattern of activity across the eight subtrials of trial type 1 in the various ROIs. The three ROIs in light brown are the phase-sensitive ROIs created from the results of the last experiment (chapter 3).

Figure 4.7 illustrates the profile of activity across the subtrials of trial type 1 in the different ROIs. As is evident, the level of activity varied across the eight subtrials, however, the actual pattern of variation seems to vary across the different ROIs. This was confirmed by a repeated measures ANOVA with subtrial position and ROI as main factors, which showed a significant interaction between the two ($F_{(84,1092)} = 4.6$, $p < 0.001$). Two distinct trends can be discerned. Some ROIs (e.g. medial frontal, ACC, bilateral ventrolateral frontal regions) show a gradual increase in activity over the

eight subtrials, but no additional increase in activity in the fourth subtrial. Activity during the fourth subtrial in these ROIs is not greater than other within-chunk subtrials (e.g. fifth and sixth subtrials). In contrast, some regions (e.g. IPS-L, IFS-L, and APFC-L) do not show an increase in activity across the eight subtrials, but have higher activity for subtrials marking the completion of the two chunks, i.e. the fourth and the eighth.

To find areas affected by the chunk level organisation the activity in the fourth subtrial was compared with that during the within-chunk subtrials. This was significant in left APFC and IFS, bilateral IPS and pre-SMA ($t_{(13)}$>2.7, $p < 0.02$). As can be expected from Fig. 4.7, the activity during the eighth subtrial was greater than that during the intermediate subtrials in all ROIs ($t_{(13)}$>3.1, $p < 0.01$), except left IFS and left IPS. Most of such ROIs also showed greater activity during the eighth compared to the fourth subtrial (table 4.2).

A repeated measures ANOVA compared the activities during subtrials 2, 3, 6 and 7 to look at phase effects within chunks (3 & 7 > 2 & 6) and across chunks (6 & 7 > 2 & 3). Only in ACC was the within-chunk phase effect significant ($F_{(1,13)} = 12.7$, $p <0.01$), others ($F_{(1,13)} < 2.9$, $p > 0.1$). Across-chunk phase effect was significant in right IFS, AI/FO, ACC and the phase-sensitive ROIs ($F_{(1,13)} > 6.8$, $p < 0.02$), and nearly significant in left AI/FO ($F_{(1,13)} = 4.5$, $p = 0.05$).

These results confirm the two distinct pattern of behaviour in the ROIs. Left APFC, IFS, IPS and pre-SMA showed greater activity at the completion of the first chunk but did not show any effect consistent with regions representing the trial as a single episode - activity during subtrial 8 was not greater than during subtrial 4, there was no increase in activity across the two chunks. Other ROIs did not show any evidence of the chunk level organisation, but showed evidence consistent with sequence level organisation – Activity at subtrial 8 was higher than other subtrials,

but subtrial 4 was the same as intermediate subtrials; there was increase in activity across the two chunks. The summary of the above results are in table 4.2.

	4 > int. $t_{(13)}$	8 > int. $t_{(13)}$	8 > 4 $t_{(13)}$	(3 & 7) > (2 & 6) $F_{(1,13)}$	(6 & 7) >(2 & 3) $F_{(1,13)}$
APFC-R	0.5	**5.5****	**3.9****	1.8	2.2
APFC-L	**2.7***	**3.2****	1.0	2.5	0.3
IFS-R	0.6	**6.2****	**3.6****	0.01	**12.6****
IFS-L	**2.5***	0.6	-1.2	0.04	1.2
AI/FO-R	0.7	**5.5****	**4.4****	0.5	**6.8***
AI/FO-L	1.5	**4.4****	-2.0	2.1	4.5
Pre-SMA	**2.7***	**2.2***	0.4	0.1	0.8
ACC	0.7	**5.5****	**4.3****	**12.8****	**10.9****
IPS-R	1.7	**3.0***	1.5	0.1	0.2
IPS-L	**2.4***	1.7	-0.23	0.2	0.3
vFr-R	0.1	**7.2****	**5.7****	1.8	**11.2****
vFr-L	0.07	**7.8****	**5.0****	2.1	**26.0****
MF	0.2	**7.6****	**6.5****	2.9	**18.8****

Table 4.2. (Trial type 1). Summary of ROI analyses. The significant t and F statistics are in bold. (* 0.05 > p >0.01; **: p < 0.01)

None of the ROIs showed significantly greater activity at the first subtrial, neither did any ROI show a decrease in activity across the subtrials as would have been expected if their activity was commensurate with the working memory load.

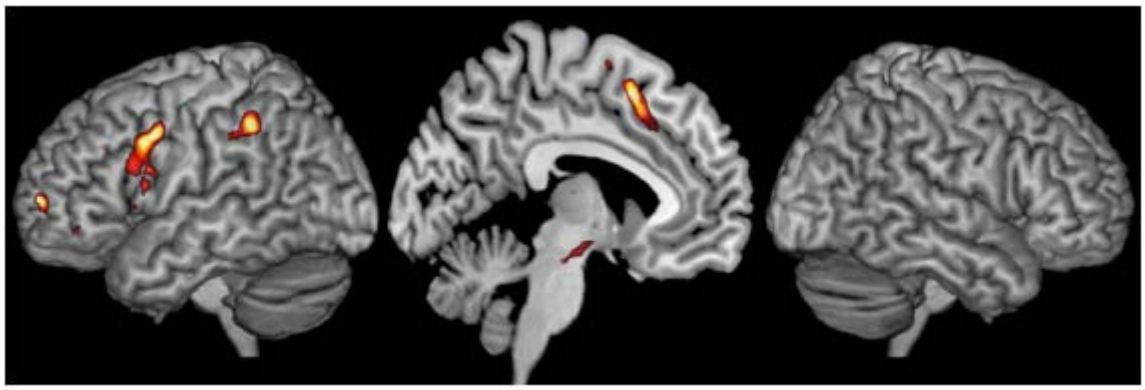

Figure 4.8. (Trial type 1) Whole brain render of a contrast looking at regions where activity during the fourth subtrial was greater than during other intermediate subtrials. Significant activity can discerned in left IFJ, pre-SMA and left APFC. (uncorrected $p<0.001$)

To look at other brain regions showing the above effects, a number of contrasts were done. The first contrast looked at brain regions that were affected by the chunk level organisation. Since this would predict higher activity during the fourth subtrial, the fourth subtrial was contrasted against the intermediate ones (2, 3, 6, 7). No brain region survived the whole brain FDR correction at $p<0.05$. However, at an uncorrected threshold of $p<0.001$, the results paralleled those from the ROI analysis (Fig. 4.8); areas of significance were present in the left inferior frontal junction (slightly posterior to the left IFS ROI), left pre-SMA and left IPL, along with clusters of activity in the left APFC and the midbrain.

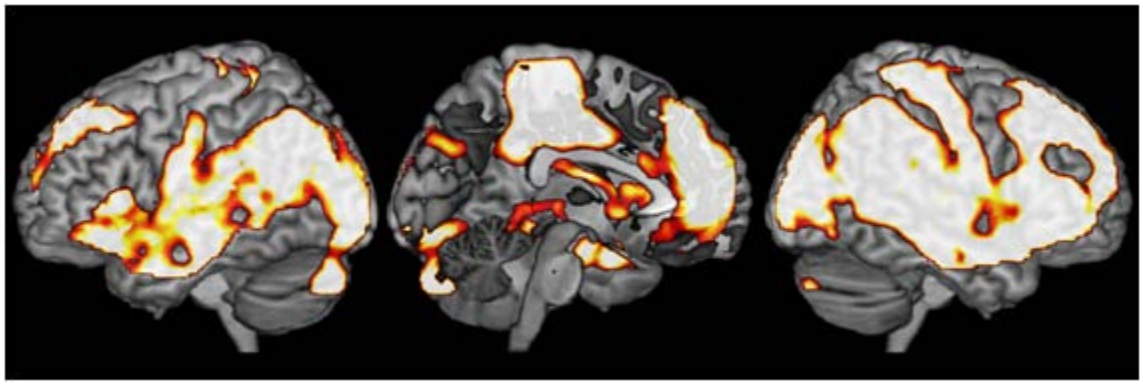

Figure 4.9. (Trial type 1) Areas with greater activity during the eighth subtrial compared to earlier subtrials.

In a second contrast, the final (eighth) subtrial was contrasted against the earlier ones. Significantly greater activity for the eighth subtrial was found in many brain regions (Fig. 4.9, coordinates in table A4.1 of appendix). In the frontal cortex, this was visibly greater on the right, where it involved the entire extent of the lateral PFC except for a small region in the middle of the middle frontal gyrus; medially the entire medial prefrontal cortex extending posteriorly to the anterior parts of ACC and pre-SMA was involved. In the left PFC, the effect was limited to the anterior insula, frontal operculum, parts of ventrolateral PFC and the superior frontal gyrus. In parietal cortices, bilateral inferior parietal lobules, angular gyrii, temporo-parietal junctions, precunei and posterior cingulate were involved along with parts of the right somatosensory regions. Huge swathes of areas in occipital and temporal cortices also showed this effect, along with hippocampus and parahippocampus, lateral cerebellum, the caudate nuclei and the amygdala.

A contrast of average activity during subtrials 6 and 7 against that during 2 and 3 revealed very similar results (Fig. 4.10), showing that the increase was not limited to the eighth subtrial but subsequent phases other than the final phase also showed greater activity than the earlier phases. The results were identical when a linear contrast was done across subtrials 2 to 7 which looked at regions where the activity increased in a stepwise manner.

In summary, the results of this trial type thus show that in widespread regions, this trial type was organised as a single coherent task episode. The evidence of chunk level organisation, though present, was weak.

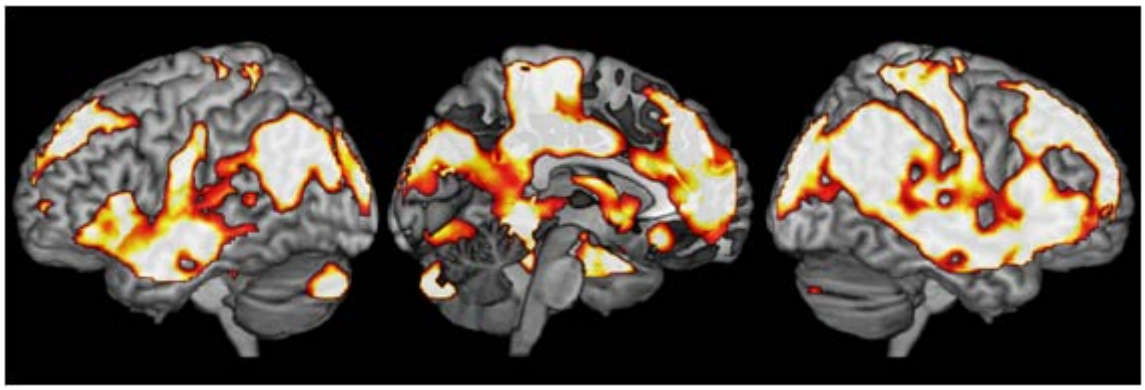

Figure 4.10. (Trial type 1) Regions with greater activity during the subtrials 6 & 7 compared to subtrials 2 & 3. Note that the regions are identical to those showing greater activity during the eighth compared to the earlier subtrials.

An important feature of the above results is its confirmation that the phase effect, described in the last chapter, is not limited to the final phase, but extends across the entire sequence of phases making up the task episode. As in the previous chapter, this effect was widespread. Indeed in the fronto-parietal cortices, areas considered to be core regions of task control (Right APFC, bilateral AI/FO, ACC; Dosenbach et al., 2006) show robust increase in activity along the length of the trial along with others (e.g. posterior cingulate, IPL, TPJ, superior temporal sulcus) that are considered part of the default mode network and tend to deactivate during the trial blocks (Raichle, 2003). These two groups have hitherto been considered to be anti-correlated with each other (Fox et al., 2005), but in the above result behave in a very similar manner.

No region showed a decrease in activity along the length of the trial, even at a very liberal threshold of uncorrected $p<0.05$, suggesting that nowhere was the

activity directly correlated with the working memory load. Lastly, compared to other subtrials, nowhere was the activity greater activity during the first or the fifth subtrials, even though, as manifested in RT profile, these were the points at which greater control was required.

4.3.2 Trial Type 2

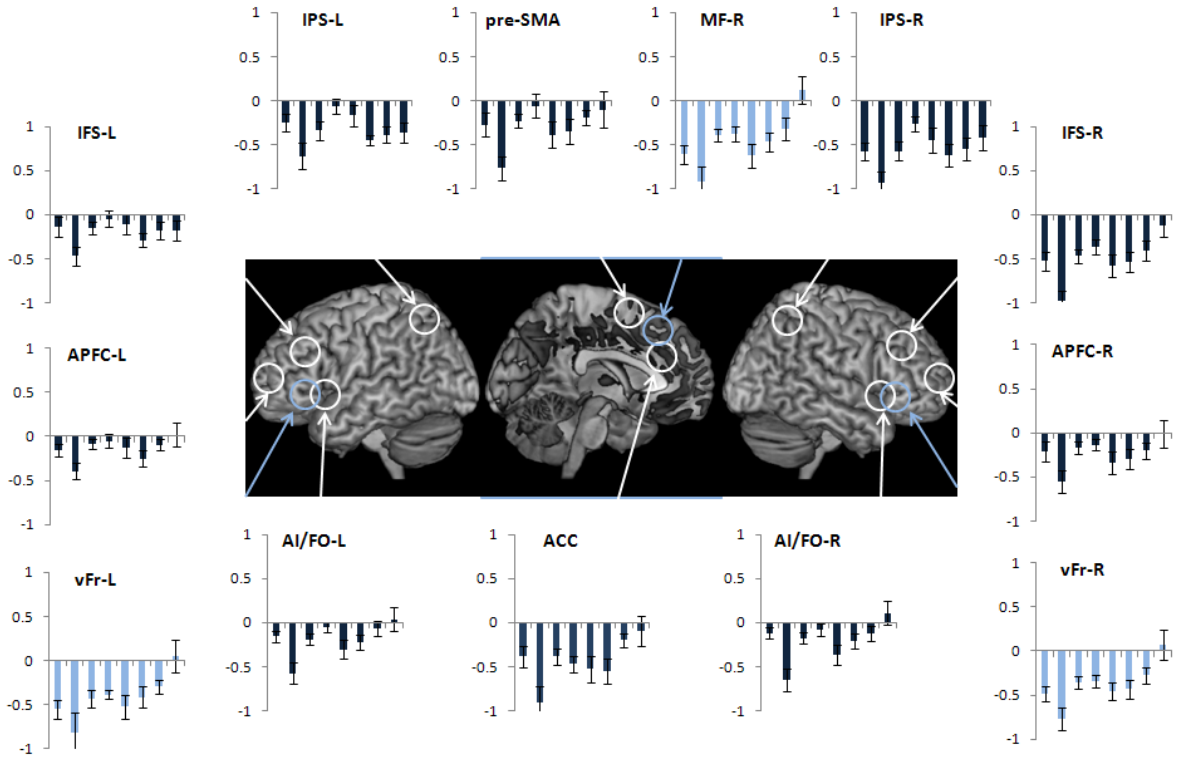

Figure 4.11. Pattern of activity across the eight subtrials of trial type 2 in the various ROIs.

Fig. 4.11 shows the profile of activity in the various ROIs during type 2 trials. As in the earlier trial type, the level of activity differed across the eight subtrials across ROIs ($F_{(7,91)} > 2.3$, $p < 0.03$), but the actual pattern of this was different for the different ROIs ($F_{(84,1092)} = 3.7$, $p < 0.001$). The comparison of activity during the fourth subtrial with the within-chunk subtrials showed greater activity during the fourth subtrial in bilateral IFS, IPS, AI/FO, right APFC, vFr and pre-SMA ($t_{(13)}>2.2$, $p<0.05$),

showing that the effect of chunk completion was more widespread than trial type 1. The activity during the final subtrial was then compared to that during the earlier subtrials. The final subtrial had significantly greater activity in bilateral AI/FO, right IFS, APFC, ACC and the three phase-sensitive ROIs ($t_{(13)}$>2.4, p<0.03). However, the activity during the eighth subtrial succeeded the fourth only in ACC and the phase-sensitive ROIs ($t_{(13)}$>2.5, p<0.02).

As for the previous trial type, activities across subtrials 2, 3, 6 and 7 were compared for within and across-chunk phase effects. Unlike trial type 1, many ROIs showed a significant within-chunk phase effect without a concomitant across-chunk phase effect. This was the case in bilateral APFC, left IFS and IPS, and pre-SMA ($F_{(1,13)}$ > 5.6, p < 0.03). Whereas in right IFS, AI/FO, vFr, ACC and MF both kinds of phase effects were significant (within-chunk: $F_{(1,13)}$ >12.0, p < 0.01; across-chunk : $F_{(1,13)}$ > 5.1, p < 0.05). The results are summarised in table 4.3.

	4 > int. $t_{(13)}$	8 > int. $t_{(13)}$	8 > 4 $t_{(13)}$	(3 & 7) > (2 & 6) $F_{(1,13)}$	(6 & 7) >(2 & 3) $F_{(1,13)}$
APFC-R	**2.9***	**2.4***	0.9	**9.1***	1.5
APFC-L	1.9	2.2	0.5	**13.0****	1.0
IFS-R	**2.7***	**4.4****	2	**17.5****	**7.4***
IFS-L	**2.4***	0.5	-1.2	**8.4***	0.9
AI/FO-R	**4.2****	**4.6****	1.5	**13.0****	**11.1***
AI/FO-L	**4.3****	**3.3****	0.8	2.1	4.5
Pre-SMA	**3.1***	1.7	-0.2	**18.1****	3.6
ACC	1.1	**3.6****	**2.5***	**13.2****	**6.8****
IPS-R	**3.0***	1.5	-0.9	4.4	1.1

IPS-L	**3.3***	0.4	-1.9	**5.6***	0.3
vFr-R	**2.2***	**3.8****	**2.9***	**12.2****	**6.0***
vFr-L	1.3	**4.5****	**3.2***	4.5	3.7
MF	1.7	**5.9****	**3.5****	**12.0****	**5.1***

Table 4.3 (Trial type 2). Summary of ROI analysis. (* 0.05 > p >0.01; **: p < 0.01)

The pattern of results (subtrial 4 > intermediate, within-chunk phase effect, absence of subtrial 8 > 4, and lack of across-chunk phase effects) in bilateral APFC and IPS, left AI/FO and IFS and pre-SMA suggest that these represented trial type 2 as two distinct chunks. ACC, left vFr and MF show patterns consistent with regions representing the trial as a single episode (subtrial 8 > 4; across-chunk phase effects). Right IFS, AI/FO and vFr show greater activity at subtrial 4 and within-chunk phase effect, but also show across-chunk phase effect.

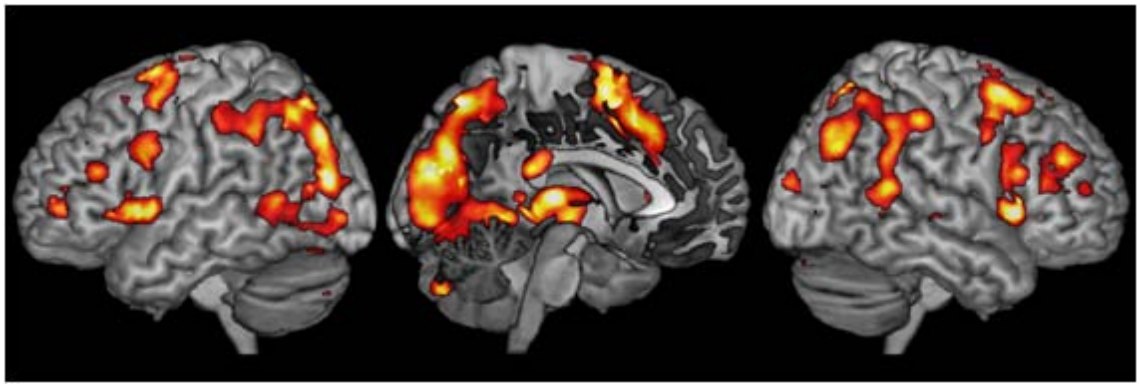

Figure 4.12. (Trial type 2): Areas with greater activity during the fourth subtrial compared to the with-chunk subtrials. Note that the results are widespread and more intensive than those obtained from trial type 1 (Fig. 4.8).

Whole brain contrasts were then done to look into other regions showing these effects. As in the previous trial type, to look at the brain regions affected by the chunk level organisation, activity for the fourth subtrial was contrasted against those during the intermediate subtrials. Fig. 4.12 shows a whole brain render of this contrast (coordinates in table A4.2). Compared to the case in trial type 1, where the

results were weak and could not survive correction for multiple comparison and could only be documented at an uncorrected p <0.001, the results here were more extensive. As in the ROI analysis, many regions in the fronto-parietal cortex showed greater activity during the fourth subtrial. These included AI/FO, areas along the antero-posterior axis of IFS up till APFC, premotor regions – ventral and dorsal extending medially onto SMA, pre-SMA and ACC. Notably, the locus of activity in ACC was posterior and dorsal to the locus of ACC that is one of the ROIs, which would explain the absence of this result in the ACC ROI. Posteriorly, the IPS was involved, along with superior parietal lobule, IPL, supramarginal and angular gyri, and the TPJ along with the posterior regions of superior and middle temporal gyrii. Medially, precuneus, cuneus and parts of the primary visual cortex were also involved.

The brain regions where the activity was greater during the eighth compared to the intermediate subtrials differed from those showing greater activity for the fourth subtrial in interesting ways.

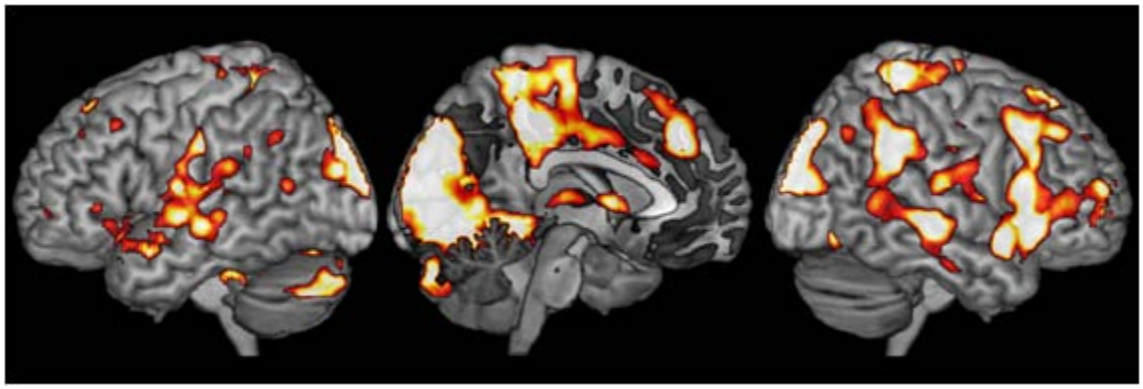

Figure 4.13. (Trial type 2) Areas with greater activity during the eighth subtrial compared to the within-chunk subtrials.

As is evident in Fig. 4.13, activity during the eighth subtrial was greater in the right prefrontal (premotor, AI/FO, VLPFC extending onto the APFC, and right superior medial regions) and the right parietal cortex (IPL, TPJ, angular gyrus, parts of primary sensory and motor cortex, medially mid and parts of posterior cingulate, paracentral

lobule and precuneus). Various regions in temporal, occipital and cerebellar cortex also showed greater activity during the eighth subtrial (Coordinates in table A4.3).

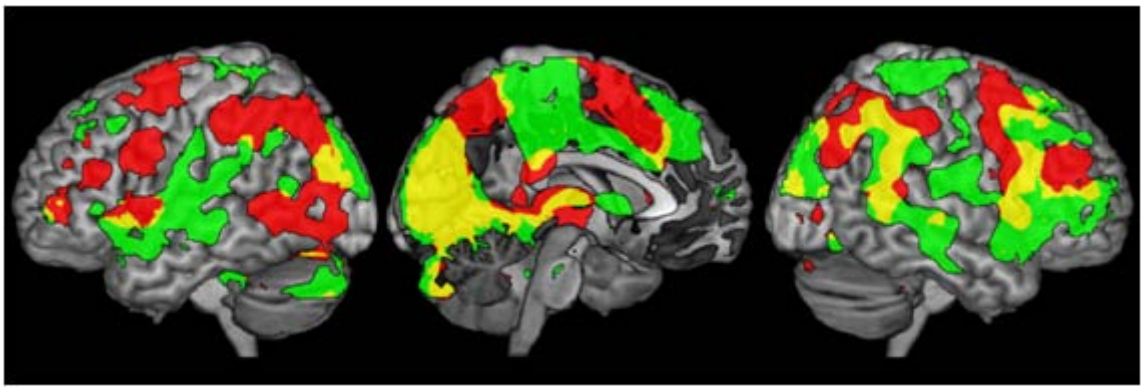

Figure 4.14. (Trial type 2) Contrasts comparing areas with greater activity for the eighth subtrial over the within-chunk subtrials (green) rendered on a common template with areas having greater activity for the fourth over the within-chunk subtrials.

The relation of these areas with those that showed greater activity for the fourth subtrial is interesting. Fig. 4.14 shows the two contrast images rendered on a common template (the images are at a slightly liberal threshold of FDR corrected $p<0.05$). Maximum overlap can be seen in the medial occipital regions (areas 17, 18, cuneus), right IPL and areas around the right TPJ, right IFS and the right AI/FO (MD region). Beyond this overlap, a clear cut separation is evident especially in the prefrontal cortex. Areas in the middle frontal gyrus show greater activity for the fourth subtrial, whereas those in the superior and inferior frontal gyri, especially in the right PFC, have greater activity for the eighth subtrial. Medially there is an antero-posterior distinction, with anterior regions of medial PFC being more sensitive to the eighth subtrial. Similarly anterior regions of the medial parietal regions have greater activity during the eighth subtrial whereas its posterior regions are more active during the fourth. All three phase-sensitive prefrontal regions (venterolateral frontal and superior medial regions), expectedly, prefer the eighth subtrial. Many of

these distinctions turned out to be statistically significant in directed contrasts between the fourth and the eighth subtrial activities (Fig. 4.15).

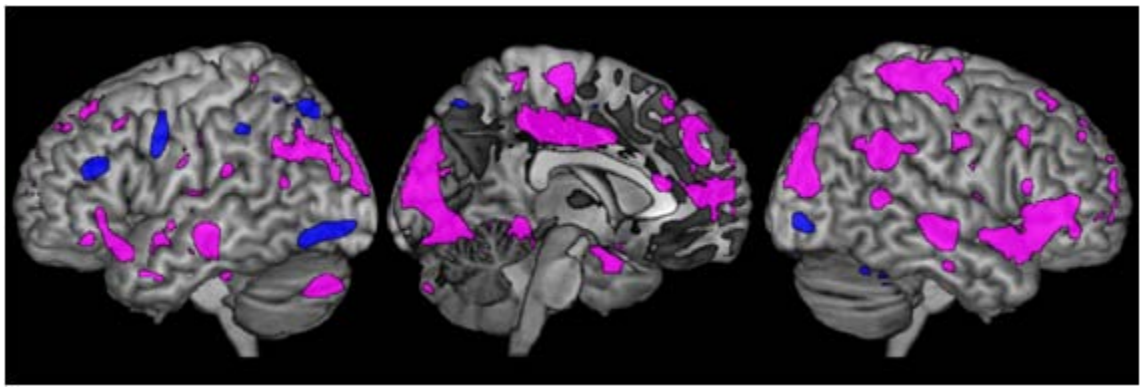

Figure 4.15. (Trial type 2). Purple: areas where the eighth subtrial had greater activity than the fourth, blue: areas where the fourth subtrial elicited greater activity than the eighth.

Lastly, the brain regions showing an increase across the two chunks were examined (Fig. 4.16). A comparison of mean activity between subtrials 2 & 3 and 6 & 7 showed higher activity for the subtrials 6 & 7 in the bilateral APFC, and the right AI/FO, extending into VLPFC, and medial PFC. Other regions showing this effect included parts of the superior and middle temporal gyri including STS, posterior insula, pre and paracentral gyri (especially on the right), head of caudate nucleus and secondary visual areas. As would be expected, a reverse contrast showed that no region had greater activity for the first chunk compared to the second. Note that the locus of APFC in this contrast is anterior to the ROI based on APFC that was defined a priori and used in the ROI analysis.

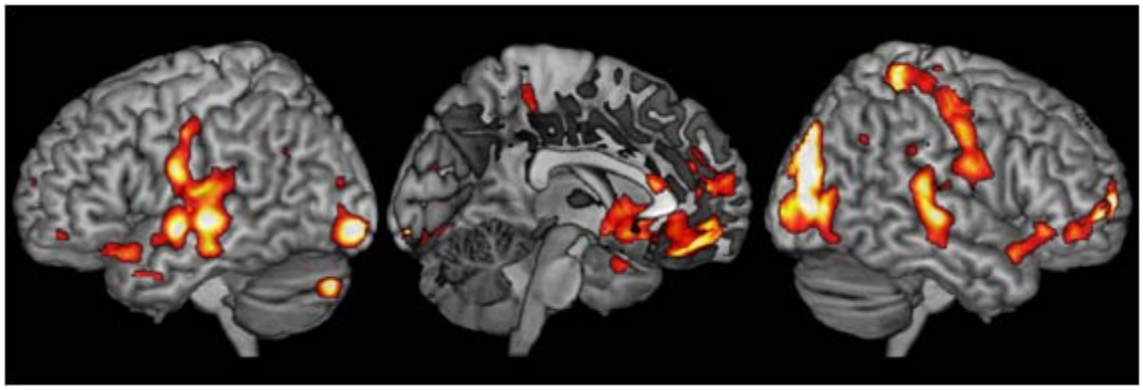

Figure 4.16. (Trial type 2) Regions with greater average activity during the subtrials 6 & 7 than during the subtrials 2 & 3. Note the paucity of the results compared to that in the earlier trial type (Fig. 4.10)

4.3.3 Comparison of the two trial types

The results from the analysis of the two trial types show that in a hierarchical task episode with multiple levels, brain regions are affected by organisation at all levels. Some regions seem to be affected only by the lower level organisation to the exclusion of any effect from the higher level organisation, and vice versa. Left APFC, IFS, IPS and pre-SMA in both trial types were only affected by the lower chunk level organisation, whereas the left vFr, MF and the ACC were only affected by the higher sequence level organisation i.e. these regions showed the effect expected if the entire trial was a single episode in both trial types. The behaviour of other regions, however, was variable across the two trial types. For example, right prefrontal MD ROIs were affected only by the sequence level organisation in trial type 1, but in trial type 2 they were mostly affected by the chunk level organisation.

It was found that trial type 2 showed greater evidence of organisation as two separate chunked episodes than trial type 1 but weaker evidence of organisation as a single sequence of subtrials compared to trial type 1, and vice-versa. Right APFC, bilateral AI/FO and right IFS showed evidence of chunk level organisation in trial type 2, but of sequence level organisation in trial type 1. To make a direct comparison between the two trial types, a repeated measures ANOVA looked at the interaction of the trial type with the effect of the position of the subtrial in the chunk (chunk level effect) and with the effect of the position of the chunk in the sequence (sequence level effect). The effect of position within the chunk differed across the two trial types in right IFS, right AI/FO and pre-SMA ($F_{(1,13)} > 4.9$, $p<0.01$); however in none of the ROIs did the difference of the effect of the position of the chunk in the sequence, across the two trial types, reach statistical significance ($F_{(1,13)} < 3.4$, $p>0.06$).

The whole brain comparisons supported these conclusions. The fourth subtrial showed greater activity in more widespread areas in trial type 2, whereas increase in activity across the two chunks was more evident in trial type 1.

The results so far, show that the two trial types are differently organised, and that subtle difference in organisation of higher level representations do indeed change the organisation of the identical lower level behavioural events. The unorganised trial (trial type 3) allowed for the test of this hypothesis at an even more subtle level. Recall that there was no sequence guiding the performance on and organising the eight subtrials of this trial type. Instead the relevant rule was signalled on each subtrial by the colour of the margins of the stimulus. Thus there was no explicit organising higher level representation. However, since this trial had the same structure as the organised trials, it was predicted that there would be transfer of the implicit higher level representations from the organised trials. For example, as the subjects had become biased to chunk the organised trials into groups of four, it was expected that they would *see* the unorganised trial as having two groups of four subtrials, even though this had no role in the execution of the subtrials of this trial type.

4.3.4 Trial Type 3

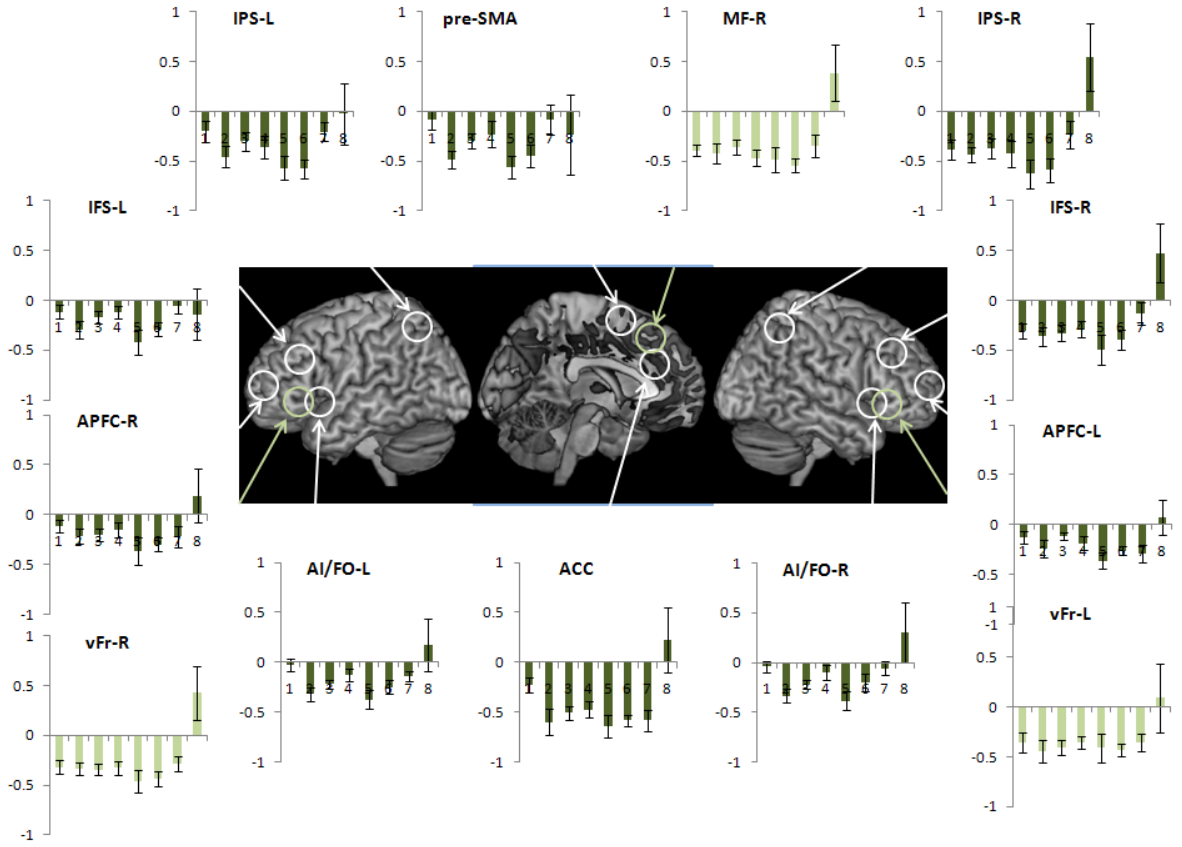

Figure 4.17. Pattern of activity across the eight subtrials of trial type 3 in the various ROIs.

As in the earlier trial types, the activity varied across the eight subtrials (Fig. 4.17) in all ROIs ($F_{(7,91)} > 2.2$, $p < 0.04$), and the ROIs differed in their differential response to the eight subtrials ($F_{(84,1092)} = 3.2$, $p < 0.001$). Since, unlike the previous trial types, this trial episode was not further organised by chunked letter strings, the subtrials 2 to 7 were equivalent. However, as is evident in Fig. 4.17, the activity is, nonetheless, variable across them. The fourth and eighth subtrials were specifically investigated.

The fourth subtrial elicited greater activity compared to the intermediate subtrials in bilateral AI/FO ($t_{(13)}>2.4$, $p<0.02$), and close to significance in left IFS

($t_{(13)}$>2.0, p<0.06). The activity at the eighth subtrial was the highest in most ROIs, it, however, reached significance only in ACC and right IPS, right IFS, right VLPFC, and MF ($t_{(13)}$>2.2, p<0.04). Other aspects of organisation into two chunks were also present. Many regions showed within-chunk phase effect – bilateral IFS, AI/FO, IPS and pre-SMA ($F_{(1,13)}$> 6.3, p< 0.03), however none showed a significant effect across the two chunks. These are summarised in table 4.4.

	4 > int. $t_{(13)}$	8 > int. $t_{(13)}$	8 > 4 $t_{(13)}$	(3 & 7) > (2 & 6) $F_{(1,13)}$	(6 & 7) >(2 & 3) $F_{(1,13)}$
APFC-R	1.4	1.6	1.3	1.3	0.3
APFC-L	0.9	1.7	1.5	0.8	1.2
IFS-R	0.7	**3.2***	**2.5***	**9.3***	0.8
IFS-L	2	0.4	-0.1	**16.1****	0.6
AI/FO-R	**2.4***	1.9	1.5	**9.7***	4.6
AI/FO-L	**2.6***	1.8	1.2	**6.3***	1.1
Pre-SMA	1.8	0.3	0	**18.1****	3.6
ACC	1.9	**2.2***	2	0.3	0.1
IPS-R	0.2	**2.9***	**2.5***	**7.0***	0.0
IPS-L	0.5	1.3	1	**12.1****	0.0
vFr-R	0.9	1.5	1.3	2.0	0.1
vFr-L	0.7	**3.0***	**2.8***	1.6	0.1
MF	-0.6	**2.8***	**2.7***	1.9	0.3

Table 4.4. (trial type 3) Summary of ROI analysis. (* 0.05 > p >0.01; **: p < 0.01)

Unlike the case in the previous trial types, the activity in the first subtrial was greater than that in the intermediate subtrials in a number of ROIs – ACC, pre-SMA and bilateral AI/FO ($t_{(13)}$>2.8, p<0.01).

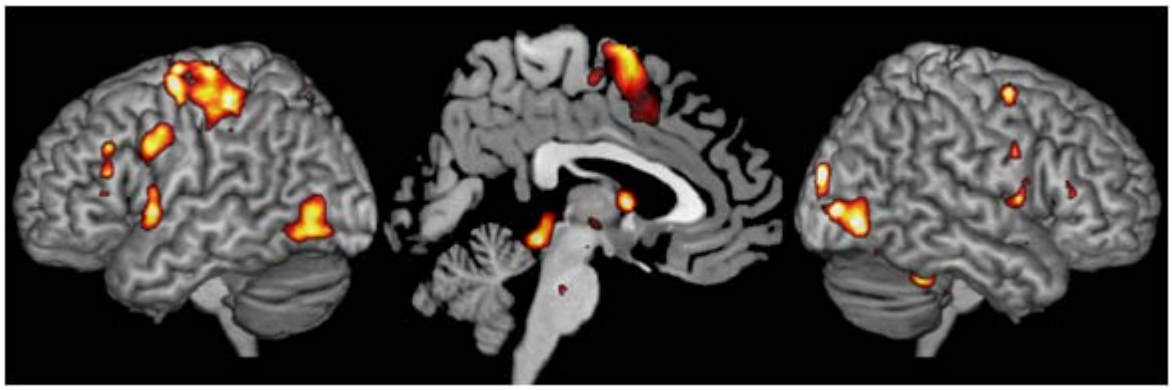

Figure 4.18. (Trial type 3). Areas with greater activity at the fourth subtrial compared to the within-chunk subtrials.

Following the ROI result, a contrast was made between the fourth against and the other intermediate subtrials (Fig. 4.18). Greater activity for the fourth subtrial was found in a number of brain regions - bilateral insula, left IFJ, left IFS (near the MD IFS region), left dorsal premotor extending into motor cortex, SMA and pre-SMA, and bilateral inferior occipital regions (coordinates in table A4.4). Thus, the *abstract view* of this trial as having two groups of four subtrials did indeed affect the activity in certain fronto-parietal and occipital regions. A noteworthy point in this regard is that this abstract conception of task plan did not affect activity in anterior prefrontal regions but rather in the posterior prefrontal regions which are usually thought to be concerned with concrete task demands.

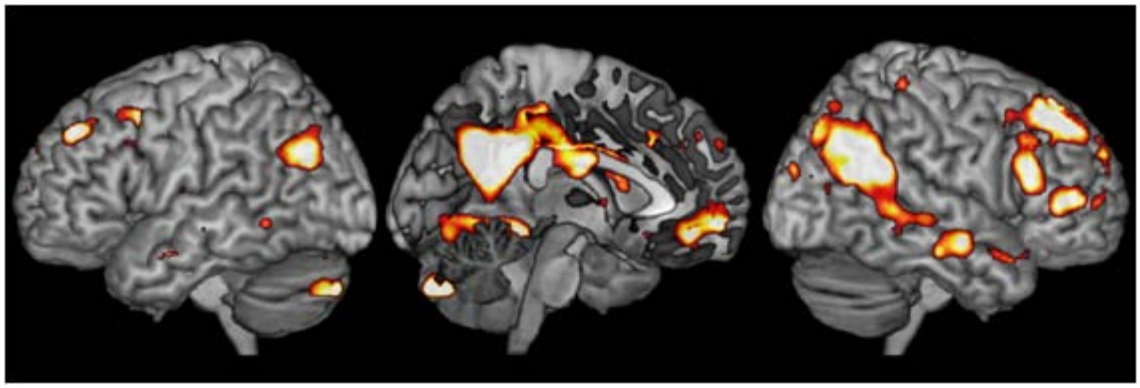

Figure 4.19. (Trial type 3) Areas with greater activity at the eighth subtrial compared to the within-chunk subtrials.

As in previous trial types, widespread areas showed greater activity for the last subtrial (Fig. 4.19), the extent of which, however, was limited compared to the previous trial types. In the prefrontal cortex this involved similar right lateralised regions– right inferior and superior frontal gyri, dorsal premotor and medial frontal regions. Similarly, the posterior clusters like in earlier trial types included regions around the TPJ, posteror to mid regions of the superior and middle temporal gyri, and posterior cingulate (table A4.5).

4.3.5 Comparison of trial types 1 and 3

Both of these trials consisted of eight sequential subtrials or phases, however the subjects' knowledge of the phase they were in was less explicit in trial type 3, as their performance was not guided by an explicit sequence. Accordingly it was expected that the activity related to the phase of the trial, seen as a sequential increase in activity across the task phases would be more prominent in trial type 1. Trial type 3 being structurally similar to trial type 1 also becomes a control condition to demonstrate that the stepwise increase across the phases of a trial is not an artefact of modelling sequential events.

ROIs	Trial Type 1	Trial Type 3
MF-R	1.02	0.22
VLPFC-R	0.92	0.33
VLPFC-L	0.86	0.19
ACC	0.63	-0.05
IFS-R	0.70	0.45
AI/FO-R	0.53	0.26
pre-SMA	0.37	-0.05
IPS-R	0.37	0.33
APFC-R	0.35	0.13
AI/FO-L	0.45	0.10
APFC-L	0.14	-0.05
IFS-L	0.05	0.03
IPS-L	0.05	-0.06

Table 4.5. z-transform of the Pearson's correlation coefficient between the subtrial number and its elicited activity. Note that the correlation coefficient was higher for trial type 1 in all ROIs.

In this regard, the linear trend in the profile of activity across the eight subtrials was quantified. Regions that are most sensitive to the phase of the task element can be expected to show a stepwise increase in activity across the eight subtrials and so have a stronger linear trend than those less sensitive. A Pearson's correlation was done between the subtrial number and its elicited activity to test the strength of the linear relationship. The correlation coefficients obtained were converted into z values using Fisher's z-transform to enable parametric statistical comparisons (Table 4.5). As apparent in Table 4.5, the z-transform of the Pearson's coefficients were greater for trial type 1 than for trial type 3 in most ROIs, being significant in all ($t_{(13)} > 2.3$, $p < 0.03$) except bilateral IFS and IPS. As is apparent, the linear trend was strongest in the phase-sensitive ROIs. Amongst the MD ROIs, right hemisphere regions had stronger linear trend than those on the left ($F_{(1,13)} = 56.6$, $p<0.001$), again confirming that effect of the sequence level organisation was dominant in the right hemisphere MD regions.

A contrast looking at brain regions where the linear increase in activity across the length of the trial was greater for trial type 1 than trial type 3 revealed significant results in widespread regions (Fig. 4.20), which were largely similar to regions that showed greater activity at the completion of the trial episode or that showed the effect of sequence level organisation in the trial types 1 & 2 (c.f. Figs 4.9, 4.10 and 4.13). It involved the phase-sensitive regions of the prefrontal cortex – ventrolateral PFC and superior medial frontal gyrii, along with superior frontal gyrus, parts of middle frontal gyrus, APFC, dorsal premotor, pre-SMA. Posteriorly, regions around the TPJ, posterior cingulate and the precuneus showed significant results.

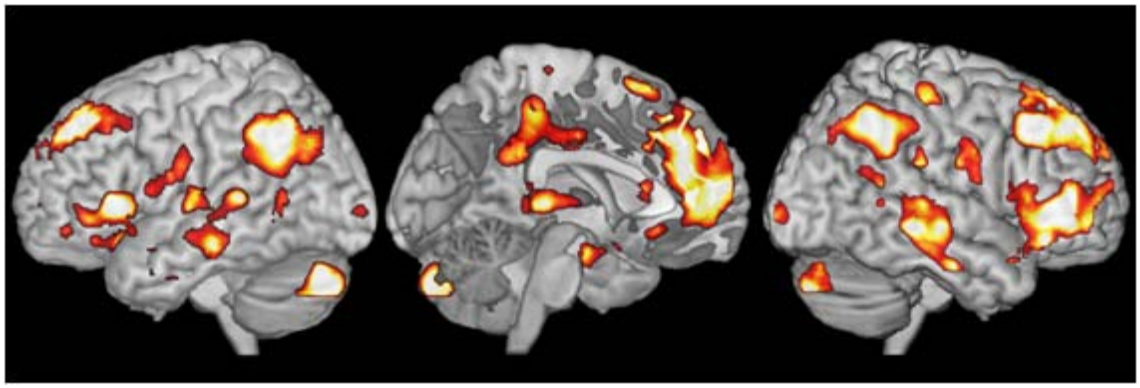

Figure 4.20. Regions where the linear increase in activity across the length of the trial was greater in trial type 1 compared to trial type 3.

While linear increase in activity was more prominent in trial type 1, the disparity between the activities during the final versus intermediate subtrials was largely similar across the two trial types (Figs. 4.7 & 4.17). This was confirmed by repeated measures ANOVA that compared the activity in the eighth subtrial with that collapsed across the intermediate subtrials, across the two trial types, in the various ROIs. In none of the ROIs was there a significant main effect of the trial type ($F_{(1,13)} < 3.1$, $p > 0.1$) or an interaction of the disparity in activities between the final and intermediate subtrials with the trial type ($F_{(1,13)} < 1.7$, $p > 0.2$). Neither did any brain region show a significant effect in the whole brain contrast looking for regions where

the difference between the activities of the final and intermediate subtrials was different across the two trial types. Thus the organisation of the task episode seems to affect only the increase in activity across the various phases of the episode but not the activity at the end of the episode.

Results so far

Results so far show that the pattern of fronto-parietal activity elicited by the various elements of the sequential task episode depends upon the way the episode has been organised. Trial types 1 and 2 were organised by an eight element sequence. Subjects had been primed to chunk this sequence into two four-element chunks. Behavioural evidence suggested that this chunking was stronger in the trial type 2. Contrary to these trial types, trial type 3 did not have any organising sequence, but like other trials consisted of eight sequential subtrials.

The results showed three kinds of brain regions. The first set of regions (hence called Group I regions), which included the three phase-sensitive brain regions (bilateral ventral frontal, superior frontal and superior medial frontal regions) and ACC had the highest activity for the eighth phase, whereas their activity during the fourth phase was not higher than during intermediate phases (2^{nd}, 3^{rd}, 6^{th} and 7^{th}). Further these regions showed a trend towards sequential increase in activity across the eight phases in the first two trial types. Since this was the only pattern of results shown by these regions, it can be said that this set of regions was sensitive only to the sequence level organisation, but not to the chunk level organisation. On the other extreme were such brain regions (Group II regions), which only showed pattern of results consistent with chunked organisation of the trial, i.e. greater activity at subtrial 4 and increase in activity within a chunk, but not across the two chunks. These were the left hemisphere lateral prefrontal regions (IFS, APFC), IPS and the pre-SMA. Additionally, in both organised trial types, eighth subtrial activity in these

regions was not greater than the fourth. Thus, these regions were sensitive only to the chunk level organisation but never to the sequence level organisation.

The behaviour of these two set of brain regions was more or less consistent across the two organised trial types, however, the behaviour of right hemisphere MD regions varied (Group III regions) across the two. These regions showed both – a weak chunk and but robust sequence level effect in trial type 1, but showed the chunk level and some sequence level effect in trial type 2. The variability in these regions was probably related to the way the representations of these trials were organised in the subjects' minds. Subjects' behaviour in these two trials showed that the effect of chunking was stronger in trial type 2 than in trial type 1. It is plausible that these regions are most sensitive to the subtleties of the plan of organisation.

The results in trial type 3, showed three distinct points. The activity was the highest during the final phase in the first set of regions mentioned above, and in some of the third set of regions (right IFS and IPS). The sequential increase in activity across the phases was much less compared to the organised trials, suggesting that the increase in activity across the phases of a task episode is dependent upon the organisation of the episode. Finally and most interestingly, the activity across the intermediate (2^{nd} to 7^{th}) phases varied, and in certain regions (left IFS, left premotor, bilateral AI/FO and pre-SMA) was higher for the fourth phase than the average of other intermediate phases, which probably meant that there was some transference of the neurocognitive structures of the organised trials to the unorganised trial. The last result also makes it unlikely that the increase in activity in IFS, IPS and pre-SMA at the completion of the first chunk during the organised trials is solely the result of the control requirements for accessing and instantiating the representations related to the second chunk, as was speculated earlier.

4.3.6 Cognitive Load and Phase Effect

An interesting aspect of the above results from whole brain analysis is the co-activation of the MD regions with another set of fronto-parietal regions that have recently been referred to as the default mode network (DMN, reviewed in Buckner, 2008) or task negative regions (Fox et al., 2005). Like MD regions, these regions also show modulation of activity across a wide variety of task situations (Dosenbach et al., 2007). However, in most studies their modulation tends to be anti-correlated with the MD regions, hence the name 'task negative' as opposed to the 'task positive' MD regions (Fox et al., 2005). Supporting this view, mind wandering (Mason et al., 2007) and attentional distractions (Weissman et al., 2006) have been shown to increase activity in default mode regions while causing a decrease in MD region activity, whereas the execution of task blocks is associated with an increase in MD activity but with a decrease in the DMN activity (Dosenbach et al., 2007). Further, this anti-correlation seems to increase with increased cognitive demands (Singh and Fawcett, 2008). Thus, an increase in the task demand further increases the MD activity while further decreasing the activity in the DMN regions. In the light of these, the current findings of correlated pattern of activity in the MD and the DMN regions needs further investigation.

This becomes even more important when it is noted that some aspects of the above results seem to suggest that the increase in MD activity seen across the steps of organised trials could be the result of the accompanying decrease in the working memory load. For example, both trial type 1 and 3 consist of eight steps, but the increase in activity across these eight steps was greater in trial type 1 compared to trial type 3. Since these two trial types differ crucially in the decrease in WM load that occurs across the steps of trial type 1 but not of trial type 3, it is plausible that this increase in activity was due to the decrease in the WM load.

To investigate these issues in greater detail, additional ROIs were created at three regions considered to be the core of DMN (Greicius, 2011) - posterior cingulate (placed midline to include both hemispheres) and bilateral TPJ. As apparent in Fig. 4.21, the pattern of activity in these regions on the three trial types was qualitatively similar to the fronto-parietal regions investigated earlier. In trial type 1, all regions showed a gradual increase along the length of the trial. In none of the regions was the activity during the fourth subtrial higher than that during the within-chunk subtrials, whereas all regions had significantly greater activity during the eighth compared to the earlier subtrials ($t_{(13)}$ >4.1, p<0.01). The average activity across the two chunks was also different ($F_{(1,13)}$ > 11.2, p<0.001), showing that the increase in activity across the trial was significant.

In trial type 2, all ROIs showed a trend towards increased activity during the fourth subtrial, however, the direct comparison showed significantly higher activity only in bilateral TPJ ($t_{(13)}$ >2.1, p<0.05). In all ROIs, there was a significant increase in activity between the second and the third phases of the chunk ($t_{(13)}$ >4.2, p<0.01). At the sequence level, however, there was no difference in average activity between the two chunks ($F_{(1,13)}$ > 3.3, p>0.06).

Together the above results suggest that these ROIs were affected by the sequence level organisation in trial type 1, but by the chunk level organisation in trial type 2. The behaviour across the two trial types was directly compared in a repeated measures ANOVA looking at the main effects of the position of the subtrial in the chunk, and of the position of the chunk in the trial, across the two trial types. The effect of position in the chunk of the subtrial was significantly different across the two trial types in posterior cingulate ($F_{(3,39)}$ = 2.8, p < 0.05), while the disparity between the average activities of the two chunks was different across the two trial types in bilateral TPJ ($F_{(1,13)}$ > 8.3, p<0.01); showing that the effects of sequence level organisation were significantly different for these two trial types in these regions.

The behaviour of these regions, thus, was very similar to that of right prefrontal MD ROIs and left AI/FO (group III regions). The above results also show that in the regions known to be sensitive to change in working memory load, load in and of itself is insufficient as an explanation of their behaviour in the two trial types. For example, the working memory load remains constant in the first four subtrials of the trial type 2, however, the activity in both TPJ and posterior cingulate vary across these subtrials.

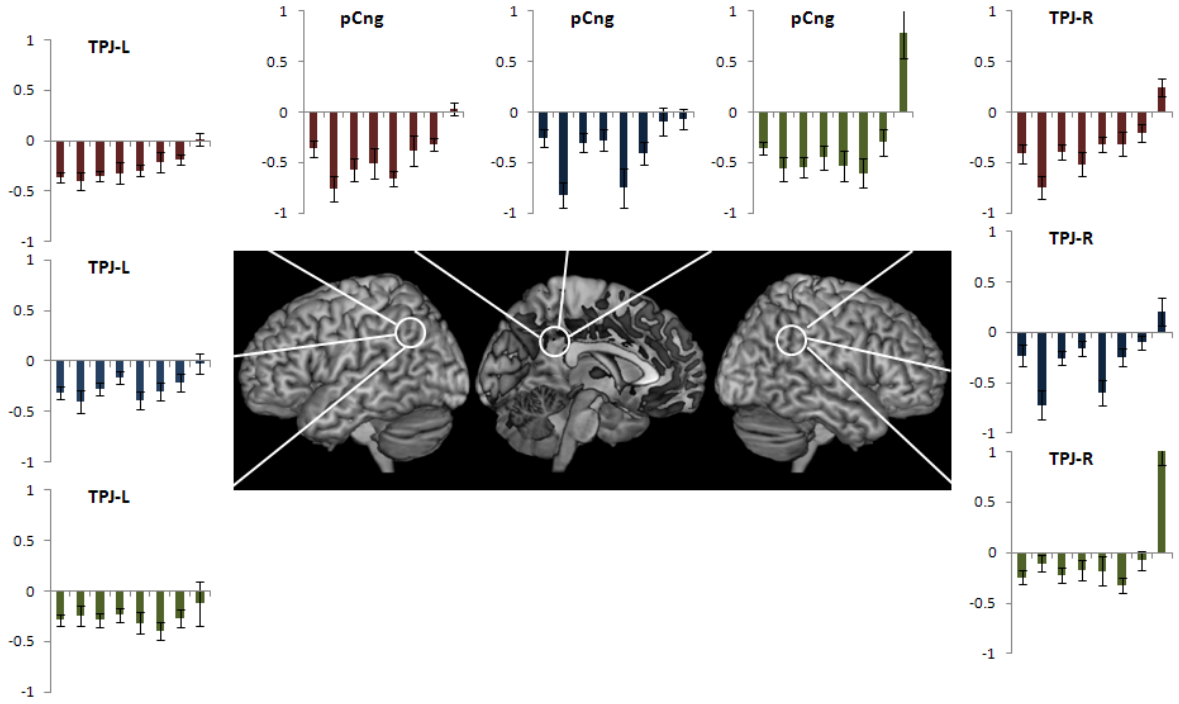

Figure 4.21. Pattern of activity in the three trial types (1, red; 2, blue; 3, green) in bilateral TPJ and posterior cingulate. Note the distinct patterns across the three trial types, especially 1 and 2.

To further test for the possibility of change in the working memory load as an explanation of the various kinds of phase effects, following comparisons were made:

- Average activity in trial type 1 versus average activity in trial type 3: The former had high working memory load compared to the latter, and the thesis about cognitive load will predict that the activity during trial type 1 in these regions would be *lower* than in trial type 3.
- Average activity during the first chunk of trial type 1 with the first chunk of trial type 2: subjects maintained two chunks during the former, but one during the latter; so the activity during the former can be predicted to be *lower* than the latter

Both of these analyses yielded null results across all ROIs. Average activity in trial type 1 did not differ from that during trial type 2 ($F_{(1,13)}$ <0.8, p>0.4). Average activity during trial type 1 was lower than that during trial type 3 only in the right TPJ ($F_{(1,13)}$ =11, p=0.005), suggesting that the working memory load did affect the activity of this region. This was not the case in any other region, indeed contrary to the above prediction, group II ROIs (IPS, IFS, APFC and pre-SMA) showed greater average activity for trial type 1 than trial type 3 ($F_{(1,13)}$ >4.7, p<0.05). Activity did not differ in other ROIs ($F_{(1,13)}$ <2.4, p>0.12,).

These findings suggest that it is unlikely that the increase in activity across the phases is related to the stepwise decrease in cognitive load per se. It is however possible that the stepwise decrease in cognitive load emphasises the knowledge of transition of phases to the subjects and that may contribute to the increase in activity. While further studies would be required to resolve this issue, the current conclusion is supported by the visibly present trend of increase in activity across the last four phases in trial type 3 (Fig. 17), where there was no change in cognitive load.

4.3.7 Cues

The cues preceding the organised trials provided subjects with working memory load, whereas the cue preceding the unorganised trial (type 3) did not have

any working memory load. This provides another way to check if the increase in activity across the phases of the organised trials were caused by the stepwise decrease in cognitive load, since areas which deactivate with cognitive load can be expected to have lower activity during the cues preceding the trial types 1 and 2 (referred to as Cues 1 and 2) than the cue preceding trial type (Cue 3). Note that since the cues were presented for 3.5 s and cue related activity would include both the encoding and maintenance related activity. The current design did not allow for the separation of these two.

As is apparent in Fig. 4.21 the activity during Cues 1 and 2 were higher than Cue 3 in all MD and phase-sensitive ROIs. However, in DMN ROIs, the pattern was reversed and the activity during Cue 3 was the highest. This result shows that activity in the fronto-parietal regions considered in the above analyses increased during the encoding and maintenance of working memory. Further, the results also show that MD and phase-sensitive regions show qualitatively different response compared to DMN regions during these conditions.

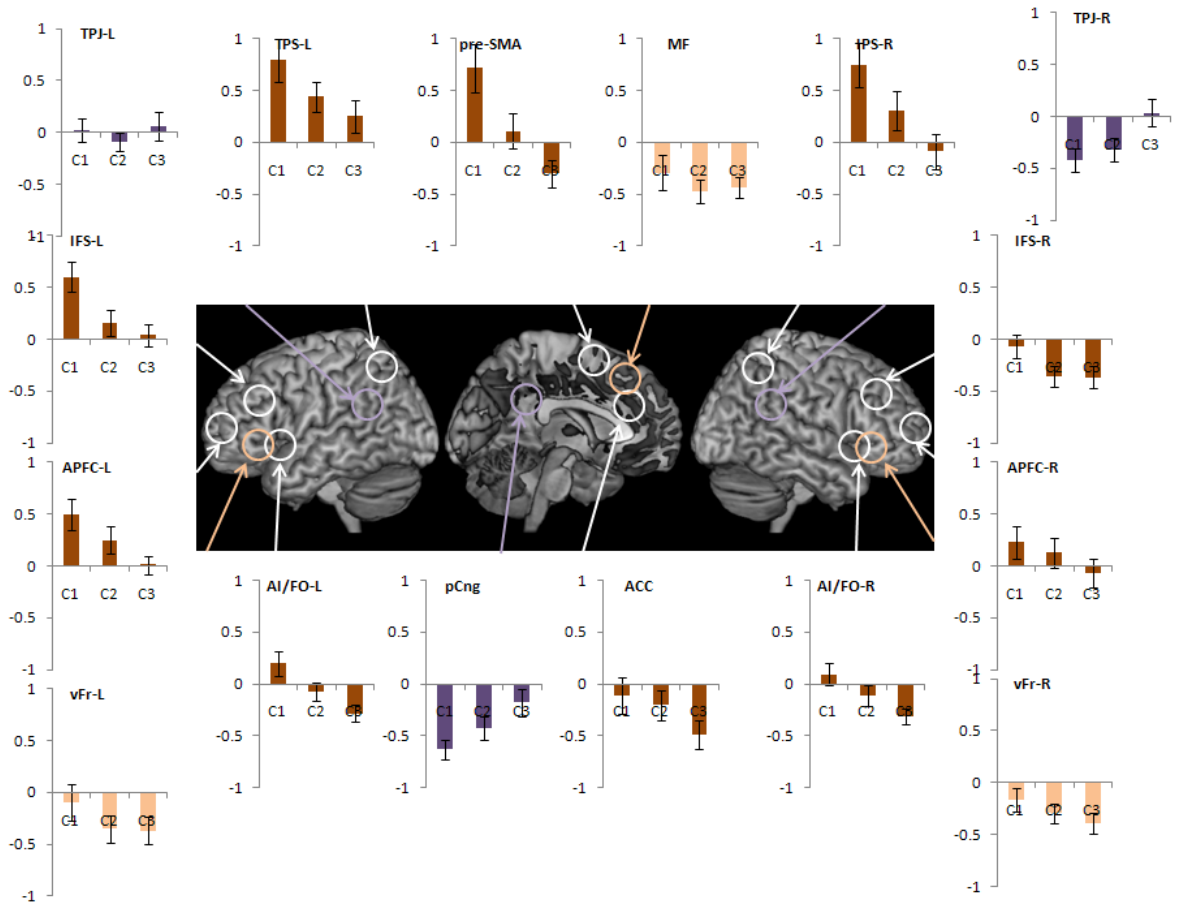

Figure 4.22. Activity in response to the cues preceding the three trial types. The working memory load in the order Cue 1 > Cue 2 > Cue 3. MD regions (brown), phase-sensitive regions (light brown) and default mode regions (purple) gave very different response patterns.

4.4 Discussion

The results of this chapter add to those of earlier chapters and show that neural activity elicited by a task element is determined by its position in the overall organisation of the task episode. End of episodes were associated with maximal activity across many regions, steps closer to the end were associated with greater activity in many regions compared to those earlier. These effects were not limited to the fronto-parietal regions but included almost all major parts of the brain. The

results also support the thesis introduced in the introduction that the nature of higher level representations determine the nature of the task episode.

4.4.1 Effect of Organisation

The findings of earlier experiments are confirmed in a new experimental framework. Current findings confirm that the effect of phase is not limited to the final, but all subsequent phases elicit greater activity compared to the preceding ones. The set of fronto-parietal regions that showed this effect in the current experiment is very similar to those in the previous chapter. Both were lateralised to the right prefrontal cortex and involved the entire extent of inferior (including insula) and superior frontal gyri, premotor regions along with most parts of middle frontal gyrus and medial prefrontal cortex. The ROIs created from the core regions showing the effect of phase in the last chapter, showed very consistent phase effects in the current results.

The pattern of additional activity elicited by the final subtrial of the trials was different from the pattern found for task episode completion in chapter 2. But it is obvious that the two cases are different. In chapter 2 the focus of investigation was additional phasic activity elicited by the event marking the transition across task episodes, while the current experimental design only allowed for an assessment of total activity at the final phase of the episode. Plausibly, the activity related to the last phase does include the activity related to the transition across episodes, but in addition also includes other activities, most prominent of which is that related to phases closer to the completion of the task episode. That might explain the very widespread nature of the current results.

The experimental design does not allow for a definite conclusion regarding the reason behind the stepwise increase in activity. Nonetheless, it is worthwhile to consider some possibilities. A critical aspect of the current design was the sequence

that was to be kept in memory (in organised trials) and then executed out. This meant that, in such trials, at every step some representation would be discarded from the working memory. It is possible that the increase in activity across the trial was related to the parallel decrease in cognitive load. Indeed the whole brain results did include the various default mode regions that are known to deactivate with cognitive load (Fox et al., 2005).

Many aspects of the results are difficult to reconcile with this view. Firstly, during cue presentation, most fronto-parietal regions showed *increased* activity with increasing load unlike the default mode regions that showed opposite pattern. Thus, the former do not normally deactivate with working memory load, consonant with the vast literature that shows increased activity in these regions during periods of maintaining greater working memory load (reviewed in Cabeza and Nyberg, 2000; Duncan, 2006). Secondly, many of these regions (including some of the DMN regions) showed increase in activity across the first four subtrials of trial type 2. However, since subjects, arguably, maintained only four letters and repeated them, none of these could be discarded across the first four subtrials. Hence, based on the working memory predictions, this increase should not occur. Thirdly, the magnitude of average activity during the various trial types was not inversely correlated with the accompanying working memory load. Trial type 1 required maintenance of eight letters, while trial type 3 required none. But the average activity across these trial types did not differ in any of these regions (including some DMN regions). Finally, many of these regions also showed increased activity across the first four and second four subtrials of trial type 3 (table 4.4), which had no such load.

An important finding is the widespread nature of the regions that show the effect of phase. These included not just the multiple demand fronto-parietal region, but also the default mode regions, sensory and motor cortices, cerebellum and subcortical regions. The correlated activity in the MD and DMN regions is

noteworthy, since these regions are supposed to be anti-correlated (Fox et al., 2005). Some studies have noted correlated activity across parts of these networks during internally focussed but goal-directed cognition (Spreng et al., 2010; Christoff et al., 2009). However to my knowledge their correlated activity in case of an externally directed task and involving such vast extent of the two networks has not been demonstrated. The current results cannot say, if the underlying reason for the modulation of activity in the various regions showing this effect is the same. It is quite possible that, for example, the DMN and MD regions activate across the trials for very different reasons.

Current findings show some interhemispheric functional differences across the prefrontal regions. As noted earlier in the chapter, left prefrontal MD regions consistently showed a pattern of activity consonant with chunk level organisation of the trial across all three trial types. However, this was not the case for the right prefrontal MD regions, which were plastic in their behaviour. In trial type 1 they did not show any effect of chunking, but in trial type 2 and 3 they did. But in all trial types, in comparison to left prefrontal ROIs, these showed more effects consistent with the organisation of the trial as a single episode.

This could point towards a speculative possibility. As discussed earlier (Fig. 4.2), representation of a trial as chunked or as a single episode could involve representing different aspects of hierarchy. Temporally abstract episodes by definition constitute a higher rung in the behavioural hierarchy, for e.g. episode of preparing breakfast is more abstract than preparing coffee. Right prefrontal regions might be biased towards representing temporally more abstract episodes, and hence, in the current experiment, preferentially represented the trial as a single episode. This may also explain why the increase in activity at the end of the task episodes has been greater in the right prefrontal regions in all experiments described in this thesis.

Similar hemispheric differences have been reported earlier but with regard to the temporal abstraction of control. Braver et al., (2003) imaged subjects while they performed blocks of single task and blocks in which they switched between two tasks. They found left fronto-parietal activity on task switch trials, which was modulated by changes in response speed, suggesting that this activity was sensitive to the transient demands on switch trials. However, when pure trial blocks were compared with task switch blocks, sustained activity was found in right anterior prefrontal region, which co-varied with the mixing cost, suggesting that this sustained activity was related to more sustained control requirements like maintaining two task sets. Thus, while left fronto-parietal regions instantiated transient control across trials, sustained control extending over the trial block was instantiated by right anterior prefrontal region.

Velanova et al., (2003) found very similar results on an episodic memory retrieval task. Transient left prefrontal activity was present on trials when subjects had to recall less recallable events compared to trials on which recall was easy. Right anterior prefrontal regions, on the other hand, showed sustained activity when blocks of less recallable events were compared with blocks of easily recallable events. A number of other studies have reported similar distinctions (Koechlin et al., 1999, 2003; Badre & D'Esposito, 2006; Pisapia et al., 2007; Reynolds et al., 2009).

4.4.2 Higher Level Representations in Organisation

The current results show that organisation of higher level representations does determine the organisation of the ensuing task episode. To an extent it could be regarded as a truism. After all, the task instructions and higher level representations ensuing therefrom, do determine how the task is to be executed and hence how the behaviour is to be controlled (Logan and Gordon, 2001). But how will the ensuing behavioural episode be parsed into episodes of specific cognitive focus, is usually

ignored. The various mental/behavioural events which will constitute a task is always assumed, while its determinants have not been looked into. This issue was investigated in chapter 2, wherein it was shown that task instructions could determine the parsing of a set of target detections into distinct task episodes. The current experiment showed that biasing the way higher-level representation like the task list is organised changed the way the ensuing behaviour was organised. This raises the possibility that the behaviour could be organised by sequencing and arranging abstract task related concepts, and it is the nature of these concepts that determines the details of the ensuing task episode.

Behaviour

Although both of the organised trial types had a list of eight letter strings to be executed, the strings used in trial type 2 involved a repetition of a four-letter sequence. In the resulting behaviour chunking was concluded on the assumption that it would take longer to execute a step if it involved moving across hierarchically organised episodes (Schneider and Logan, 2006, 2007; Lien and Ruthruff, 2004). Fifth RT was more prolonged in trial type 2, suggesting that the chunking was greater in this trial type. It is likely that this was due to the fact that in case of trial type 2, representing just four letters was more parsimonious, while trial type 1 required the representation of all eight letters. It is interesting to note that a strategy to reduce working memory load, on its flip side, resulted in greater behavioural effort manifesting in greater RT at the fifth subtrial.

The current results bear an interesting comparison to those of Schneider and Logan (2006). They had subjects memorise and execute two similar four letter chunks (‘AABB’ and ‘ABBA’) in the order –

chunk1 chunk1 chunk2 chunk2 chunk1 chunk1 chunk2 chunk2…

The chunks were thus alternately repeated and switched. They found that RT was greater on the first element of the chunk when the current chunk was a switch from the earlier one than when it was a repeat. The authors interpreted this as a chunk switch cost, akin to task switch cost seen when a new task has be to executed (Rogers and Monsell, 1995). The current experiment tried to replicate this finding using just two chunks. However, just the opposite was found, chunk repeats resulted in greater cost of transition than chunk switches.

It is possible that the findings of Schneider and Logan (2006) do not stem from chunk switching per se, but from the peculiarities of the organisation of the resulting sequence. Their design required subjects to execute the chunks in the above-mentioned sequence ad infinitum. In such a case subjects could have hyper-chunked the chunk repeats, and the series –

chunk1 chunk1 chunk 2 chunk2 chunk1 chunk1 chunk2 chunk2

could have been organised as –

chunk1 chunk1 – chunk2 chunk2 – …

It is possible that the same contributes to the task switch cost seen in many studies that use a similar design of alternate task switches and repeats (Rogers and Monsell, 1995). Indeed, the finding that organising such tasks into sequences alters task switch costs (De Jong, 1995; Lien and Ruthruff, 2004) supports this speculation.

Imaging

The elicited pattern of activity in the two trial types were different in a way that would be expected from their different organisation, with trial type 2 being organised more as two episodes, and trial type 1 as a single episode. Greater and more widespread activity was seen at the fourth subtrial in trial type 2 than in trial

type 1. More widespread result was seen in trial type 1 for regions showing maximum activity at subtrial 8. Secondly, more areas in trial type 1 showed increase in activity across the eight phases, whereas trial type 2 had more frequent cases of areas showing increase in activity across the four phases of the chunk, but no increase in activity across the two chunks.

Why does the organisational structure of the episode affect the fronto-parietal activity? The immediate answer would be that the organisation determines the nature of control structure required to execute the episode. The behavioural profile suggests that the control requirements did indeed vary with the trial organisation. However, the pattern of activity across the trial does not seem to agree with what would be expected if fronto-parietal regions were sensitive to the control demands resulting from the organisation. For example, the peak of activity occurred at the eighth or the fourth subtrial and not at the first or the fifth. While it could be argued that the activity at the fourth subtrial represents the control demands related to switching across the chunked episodes (which also manifests in increased RT at subtrial 5), it does not explain why the activity in many of such regions is equally high (if not higher) at subtrial 8. Even more difficult to reconcile is the finding that many of these regions also show an increase in activity between the second and third phases of these chunks, and in some cases, across the two chunks.

The findings most inexplicable from this perspective are those of trial type 3. Recall that trial type 3 was different from others since it had no higher-level representations or list of task elements to organise. Subjects just faced a series of eight subtrials and were told on each of them what to do. However, given that it had the same eight subtrials as components, the bias of representing them in groups of four acquired in other trial types, plausibly, led subjects to *perceive* these trials as having two episodes of four subtrials each. They showed two effects that were also seen in trial type 2 – increased activity at the fourth subtrial and an increase in

activity between phases 2 and 3 of the two chunks. Even more interesting is the fact that these effects were most manifest in IFS, AI/FO, IPS, pre-SMA - regions that are mostly associated with representing more concrete task related information (Woolgar et al, 2011; Sigala, 2008; Badre, 2009), unlike more anterior regions that have usually been implicated for abstract representations (Christoff et al., 2009).

CHAPTER 5

Representation of Task Information across the Sequential Phases of the Task Episode

5.1 Introduction

Previous chapters showed that the neural activity elicited by a task element depends upon the position of that element in the overall task organization. In this chapter, the influence of task organization on the pattern of elicited activity across voxels of the region will be investigated, which, arguably, can be considered as a better reflection of actual representation of the task element.

Conventional fMRI methodologies look at the measures of the magnitude of activity at individual voxels; these are averaged across all the voxels of the region to conclude if it is sensitive to the particular mental/behavioural event. While this allows one to conclude if the region is involved in the processing of the mental event, it gives relatively little clue about the actual representational content of the region during the processing of that mental event. This limits the epistemic value of such methodologies since any mental event has many cognitive concomitants, and conventional analyses cannot tell which one of these was actually represented by the neurons of the region in question. For example, a particular task event may be lead to a switch in rule, a switch in response, completion of a task episode, completion of a goal and so forth. Univariate analysis can tell if the task event, as a whole, elicited activity in the region of interest, but not which of the associated cognitive event was actually represented in that region.

Methodologies based on the analysis of voxel patterns offer relative advantages in this regard. Neural representations in a variety of studied domains are

based on the activity of a population of neurons (Georgopoulos et al., 1986; Young et al., 1992; Averbeck, et al., 2006). Given the relative size of brain areas with respect to the voxel size used in fMRI studies, in most cases the relevant population of neurons would be distributed over a number of voxels. Hence qualitatively different neural population codes could elicit qualitatively different patterns of activity across the voxels of the region. In other words, qualitatively distinct representations in a region will elicit qualitatively distinct patterns across voxels.

This allows for greater inferential judgements about the kind of information represented in the region of interest. Consider the previous example of the task event with many cognitive accompaniments; if the region represents relevant task rules, distinct rules will have different population codes and hence will elicit qualitatively distinct patterns of voxel activity. On the other hand, if the region does not have information about the response made, different responses will not elicit different patterns of activity. Thus, the degree of distinction between the patterns elicited by the exemplars of a cognitive category can be a measure of the representation of that category. This distinction can be measured in terms of the accuracy with which the patterns in question can be differentiated by learning algorithms that have been trained to distinguish such patterns on a separate but representative part of the data (Norman et al., 2006), or in the measures of dissimilarity between the different patterns like 1-correlation, Euclidean distance and Mahalanobis distance (Kriegeskorte et al., 2008).

Earlier studies using pattern analysis methods have shown that the fronto-parietal regions represent working memory contents (Harrison et al., 2010), decision results (Haynes et al., 2006), planned action (Haynes et al., 2008), task rules (Woolgar et al., 2011), contents of imagery (Stokes et al., 2009) and so forth. Indeed these regions seem to represent any information that is relevant to the task at hand (Duncan, 2010). More importantly, the magnitude of such representations increases

with increase in cognitive control and attentional resources invested (Woolgar et al., in press). It was seen earlier (chapter 3) that the effect of phase on the overall activity of these regions was similar to the effect of cognitive control. This raises the question whether the effect of phase on the representation of task elements will also be similar.

In this chapter, different aspects of the effect of task organization on the representation of the individual task elements will be discussed. The first issue pertains to the effect of phase on the representation of the task elements. Recall that the task elements representing phases closer to completion elicited greater activity in many regions. However, the conclusion that the increase in activity was due to the phase per se was confounded. It is possible that the increase in activity was the result of some non-specific modulatory effect resulting from increasing expectation of the imminent goal completion, or the increase was unrelated to the task event and was the result of disinhibition due to the accompanying decrease in cognitive load, akin to default mode activity. However, in both of these possibilities, the actual representation of the task elements would not be affected, since these predict a mass effect in terms of an increase in activity of the whole population of neurons, which would not affect the informational content of the task elements' representations. Thus these possibilities would get ruled out if regions showing increased activity were also shown to have greater information about the later compared to the earlier phases.

In organized tasks, the information associated with the constituent task elements can be considered to belong to two distinct categories. The first category pertains to the information that is associated with a task element by virtue of its position in the task plan, e.g. its being the penultimate step of the task plan or it completing the task. For example, in the series of actions that constitutes the steps of a hypothetical breakfast: boil water, add coffee, add cream, toast bread, and

spread butter; ‘spread butter’ represents the last step of the task of preparing breakfast and ‘toast bread’ represents the penultimate step. Information of this kind associated with the event constitutes the first category of information. Earlier chapters have shown that task events that differ in this category also differ in elicited neural activity.

Task elements also have another kind of information associated with them which pertains to the information about how they are to be executed, e.g. in case of ‘spread butter’ how the butter is to be spread, the rules relevant to it and so forth. This constitutes the second category of information. It is an open question if the representation of both of the above kinds of information is affected by the position of the task element in the plan of the task episode.

Another issue pertains to the change in the representation of a cognitive event by the very fact of its being a part of an overarching task compared to the case when the cognitive event is assessed as a task in itself. For example, consider the series of actions: boil water, add coffee, add cream, toast bread, and spread butter. Each of the individual steps can be conceived as a task in itself (e.g. ‘boil water’) or as one of the steps in the overall task of preparing breakfast. Will the representation of these events be different in these two cases?

Varying predictions can be made in this regard. The representation of a task event as part of an overall task means that its representation would be ensconced under higher level representations organizing more extended task episodes. If the same limited capacity neurocognitive reserves are required for representing both the higher and the lower level representations, it can be predicted that less resources would be available for representing lower level task elements, and hence, the representation of these would decrease when they are part of a strongly organized hierarchical task episode, compared to the situation when they are weakly organized.

From a different perspective, the individual task elements of an organized task episode may require more distinct representation since the individual steps now have to be more thoroughly distinguished from other steps. Indeed electrophysiological studies have shown that successive phases of behaviour are represented by very distinct neural states (Abeles et al., 1995; Sigala et al., 2008). Recording from a population of prefrontal neurons, Sigala et al. found that the neural activity across the sequential phases of a trial (cue – delay – target) was orthogonal. From this perspective it can be predicted that the representation of the task elements representing successive phases of a higher order task would be more distinct than successive task elements which do not represent successive phases of a higher order task.

The experimental design discussed in the last chapter allowed for the investigation of these issues. Recall that each trial consisted of eight subtrials which were identical in all aspects except for their sequential position. Accordingly, the distinctness of the pattern of activity elicited by subtrials at different position can be the measure of the representation of the phase information associated with the subtrial (category 1 information). Further, the subtrials at each position could be of one of the two types (referred to as A and B) depending upon the relevant rule that decided the correct response. The distinctness of the patterns associated with subtrials A and B at any position was therefore the measure of the representation of category 2 information, which was not related to the phase per se but was related to the execution of the individual subtrial. These measures of information could be compared across the eight positions to get the effect of position on the representation of information related to the task elements. If the increase in activity across the phases of the task plan, seen in earlier chapters, were accompanied by an increase in the representation of the corresponding task elements, the informational content in the relevant brain region can be expected to increase across the eight positions.

The different kinds of trial types that differed in organization, allowed for an investigation into the effect of organization on the representation of the individual subtrials. The organized trials were characterized by organizing (higher level) representations that sequenced the eight subtrials, while in the unorganized trials such higher level representations were absent and these subtrials were not explicitly organized. Note that the latter trials were not unorganized in the true sense of the term. The subjects did have an explicit knowledge that these trials consist of the eight subtrials. However, since there was no organizing representation and subjects' knowledge of their position in the trial was less explicit, these trials can be said to have been less organized than the organized trials. Thus, the term 'unorganized' is used in a relative sense. Given the difference between these two kinds of trials, comparison of the representation of different aspects of the lower level task elements across these trial types would show the effect of organization on them. Note that even though the rule relevant on any subtrial was cued differently across the organized and the unorganized trials (identity of the letter in the memorized string sequence in the organized trials, colour of the stimulus margin in the unorganized trials), this was irrelevant, since the critical point is the distinctness of a subtrial compared to other subtrials of the same trial type - whatever the mode of conveying the relevant rule, it was the same for all the subtrials of a given trial type.

The results, in summary, showed that the representation of information related to the subtrials increased across the eight positions in many brain regions. This increase was seen in both categories of information. Organization of the trial seemed to decrease the representation of individual subtrials, since the distinctness of pattern elicited by them was greater in trial type 3 than trial types 2 and 1.

5.2 Methods

The details of experimental design and fMRI acquisition have been described in the previous chapter. The multivariate analysis described here is based on representational similarity analysis (described in Kriegeskorte et al., 2008). Two distinct general linear models were used to get the univariate estimate of activity in response to each event at each voxel.

The first general linear model (GLM) was the same as the one used in the analysis described in the previous chapter. The individual subtrials were separately modelled according to the trial type they belonged to and the position they occupied. They were modelled with epoch regressors having a width equal to the duration of the subtrial (beginning from the moment the stimulus of the previous subtrial disappeared and ending at the moment when the stimulus of the current subtrial disappeared) convolved with the canonical hemodynamic basis function. For all subtrials, the common activity related to the appearance of a stimulus was modelled out using a common event regressor of no duration.

In the second GLM, each subtrial was additionally characterised by the rule relevant to it beyond being characterised by their trial type and the sequential position they occupied. Other aspects of this GLM were the same as the previous one. In both cases, the first level analysis was done using unnormalised and unsmoothed pre-processed images.

5.2.1 ROI analysis

These were done at the same ROIs that were used in the previous analysis – multiple demand regions, phase sensitive regions and two nodes from the default mode network. For all ROIs, a corresponding ROI in subjects' native space was created. The parameter estimates (β values) of activity for the different subtrials were extracted for each voxel of the ROIs and a vector array of length equal to the

number of voxels was created. The individual cells of the array had a value equal to the β value of the corresponding voxel. These arrays then quantified the pattern of activity elicited by the various subtrials in that ROI. To measure the distinctness of activity pattern elicited by the subtrials, correlation distance (1-linear correlation between the corresponding arrays) was calculated across every pair of subtrials. Correlation distance was used as the measure of dissimilarity since it normalized for the mean level of activity, which, as seen in the last chapter, was variable across the subtrials. The dissimilarity values for all pairs of subtrials were assembled in a matrix (called representational dissimilarity matrix, RDM), the rows and columns of which represent the subtrials, and the value in the cell represented the dissimilarity between the corresponding subtrials (Fig. 5.1). The matrix was symmetric about the diagonal that represented the dissimilarity of a subtrial with itself and so had a value equal to zero. For visualisation, the dissimilarity values within each matrix were rank transformed and scaled between values of 0 and 1, and represented by a colour code. This was constructed for each ROI and for each subject and then subjected to further tests.

The example in Fig. 5.1a depicts an ROI RDM for the eight subtrials of trial type 1. The diagonal represents the dissimilarity between each subtrial and itself which obviously would be zero. The first column represents the dissimilarities between the first and the other subtrials, the second column likewise has the dissimilarities between the second and the other subtrials. Since the matrix is symmetric across the diagonal, the eight rows represent the same dissimilarities as the eight columns. Such ROI RDMs allow for the visualisation of the pattern of dissimilarities. For example, it is apparent in this sample that the dissimilarities between subtrial 8 and the rest are higher than dissimilarities between earlier subtrials. This example also shows that the dissimilarity between a subtrial and the following one (i.e. between subtrials 2 and 3, 3 and 4, 4 and 5 and so forth) increases with the position of the subtrial

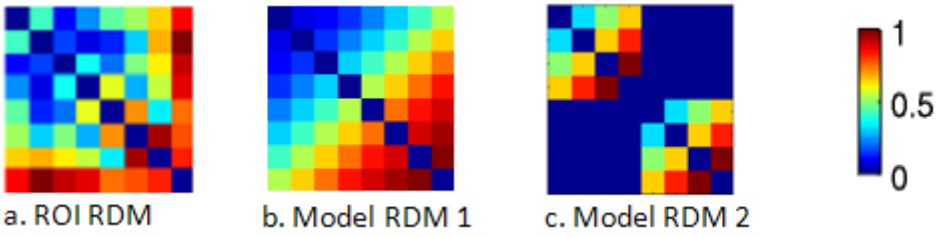

Figure 5.1. (a) A sample representation dissimilarity matrix (RDM) of left IPS from a subject representing the pattern of dissimilarities between the eight subtrials of a trial type. The rows and columns represent the eight subtrials and each cell represents the dissimilarity (1-correlation) between the corresponding two subtrials. For each matrix, the values have been rank transformed and scaled between 0 and 1. The diagonal represents a vector of zero (represented by deep blue) corresponding to the dissimilarity between each subtrial and itself. The ROI RDM represents the actual pattern of dissimilarities extracted from an ROI. (b) This model RDM represents the predicted pattern of dissimilarity based on the hypothesis that the later subtrials were more dissimilar compared to the earlier ones i.e. the dissimilarities increase with the position of subtrials in the trial e.g. dissimilarity between 1 &2 < 1 & 3 < 2& 3. (c) Model RDM (2) with dissimilarities increasing between subtrials 1 to 4 and 5 to 8 in parallel. Note that the dissimilarities between the subtrials of 1st and 2nd chunk have been masked and excluded from the analysis.

Two methodologies were used for statistical analysis of the pattern within the ROI RDMs. For analysis involving populations of values that were not completely independent of each other or for testing a pattern that involved a broad aspect of the ROI RDM, model RDMs were constructed. They had the same structure as the ROI RDM that they were being used to test i.e. they had the same rows and columns. The values in the cells of the model RDMs were based on the idealised case of the hypothesized pattern. For example, consider the model RDM1 in Fig. 5.1b, that models the idealised case when the dissimilarity between the subtrials increases with the position of the subtrial, so the values in the rows and columns of this RDM increase from row 1 to row 8 and from column 1 to column 8. Since the rows and columns stand for the eight subtrials, this pattern means that in the idealised case of this hypothesis, the dissimilarity between subtrial 1 and others will increase with the

position of the subtrials i.e. greater dissimilarity between subtrials 1 and 8 compared to that between 1 and 7, which is in turn greater than 1 and 6, and so on; similarly the set of dissimilarity between subtrials 2 and others will be higher than that between subtrials 1 and others. Like the ROI RDM, the model RDM is also rank transformed and scaled between 0 and 1.

The model RDM need not make a hypothesis about the entire extent of the ROI RDM, instead it can test for a pattern over a part of the ROI RDM. In such cases the model RDM is congruent only with the part of the ROI RDM that it seeks to test, while other parts are masked. For example, the model RDM in Fig 5.1c seeks to test the pattern in the top left and bottom right quadrant of the ROI RDM in Fig 5.1a.

To measure the likeness between the patterns in the ROI and the model RDMs, Spearman's rank correlation coefficient was obtained between the two. Note that Spearman's rank correlation was used because a linear relationship cannot be assumed between the model and the ROI RDMs. To assess the statistical strength of this coefficient, the rows and columns of the ROI RDM were randomly reordered and then correlated with the model RDM. This step was repeated 105 times, to get a distribution of correlation coefficients expected if the two RDMs were unrelated. If the correlation between the actual ROI and model RDMs fell within the top 5% of the simulated null distributions, the null hypothesis that the matrices were unrelated was rejected.

The second methodology was used for dissimilarity values that were independent of each other. In such cases the relevant correlation coefficients were extracted from the ROI RDMs, z-transformed (using Fisher's z transform) and then subjected to further statistical parametric tests.

5.2.2 Whole Brain Analysis

Whole brain searchlight analyses were done to look for other brain areas showing predicted patterns of representation. A spherical searchlight (9 mm radius) was sequentially moved across the entire set of brain voxels (brain images used were unsmoothed and unnormalised). At each step, the pattern of the activity across the voxels within the searchlight was computed for each event. The relation between the patterns elicited by the various subtrials was computed as described above for ROI analysis, and a searchlight RDM assembled. This RDM was then correlated with the model RDM. The resulting correlation coefficient was z-transformed. This value was assigned to the voxel that served as the centre of the searchlight. A brain map was thus obtained for each subject with voxel values equal to the z- transformed correlation coefficient between the model RDM and the searchlight RDM based on that voxel. These were then normalised into MNI space, smoothed and fed into a random effects analysis to look for the behaviour of brain regions across the subject population. The results were thresholded to correspond to an FDR correction of $p < 0.05$.

5.3 Results

5.3.1 Effect on Category 1 information

ROI analyses

The representations of the eight subtrials making up a trial were compared. As described earlier, the dissimilarities between each subtrial with the other seven subtrials were assembled in RDMs. Fig. 5.2 shows such RDMs of the various ROIs for the three trial types, averaged across the subjects. It is apparent that across all three trial types, many ROIs show greater dissimilarities for the later subtrials with the eighth subtrial being the most dissimilar. This result thus parallels the earlier results of an increase in activity for the later subtrials, with the eighth eliciting the greatest

activity, and suggests that the increase in activity seen earlier was accompanied by an increase in representation of the information associated with the subtrials.

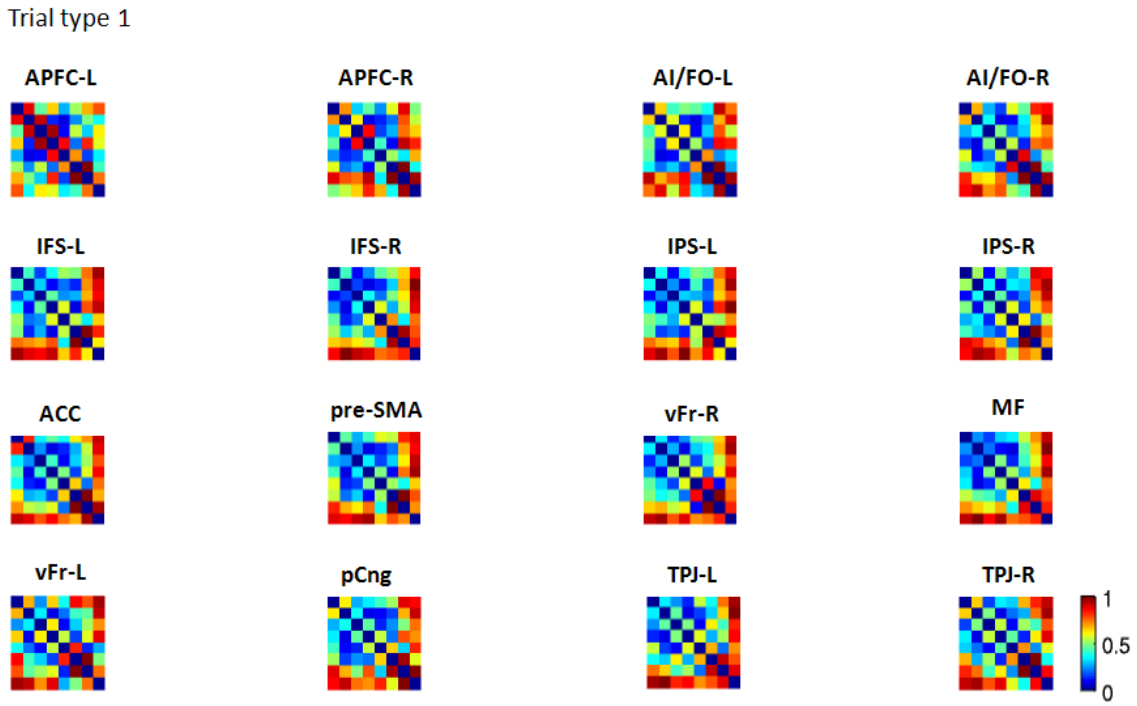

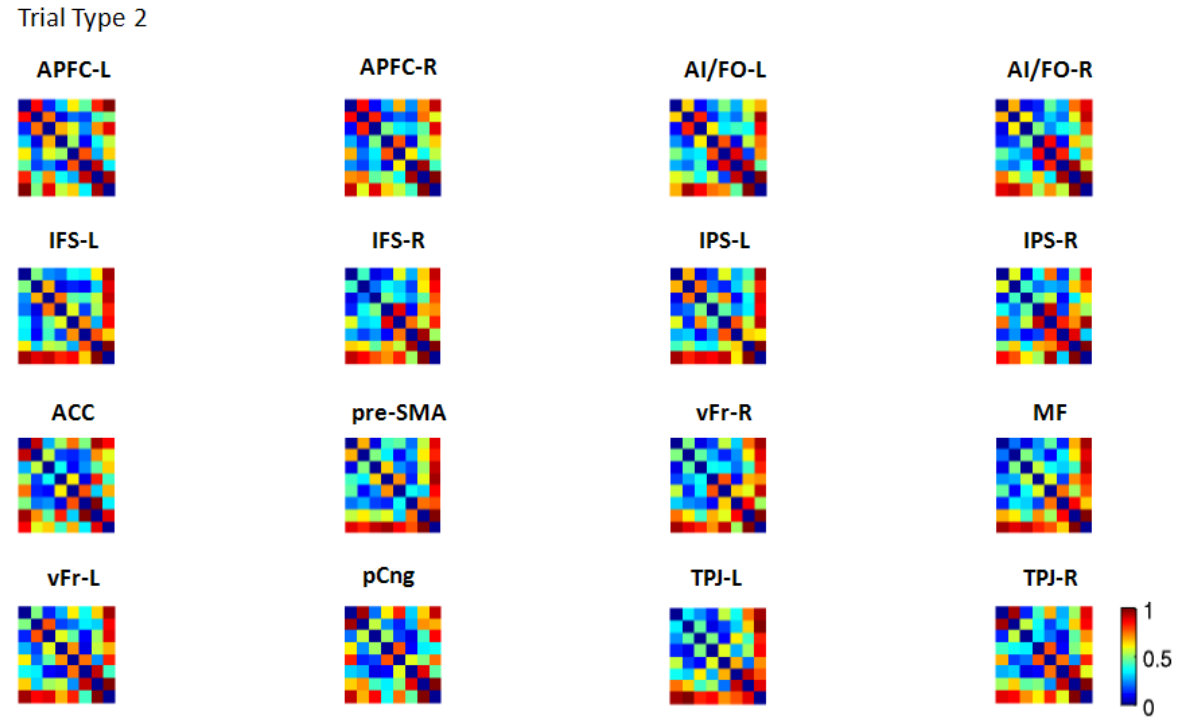

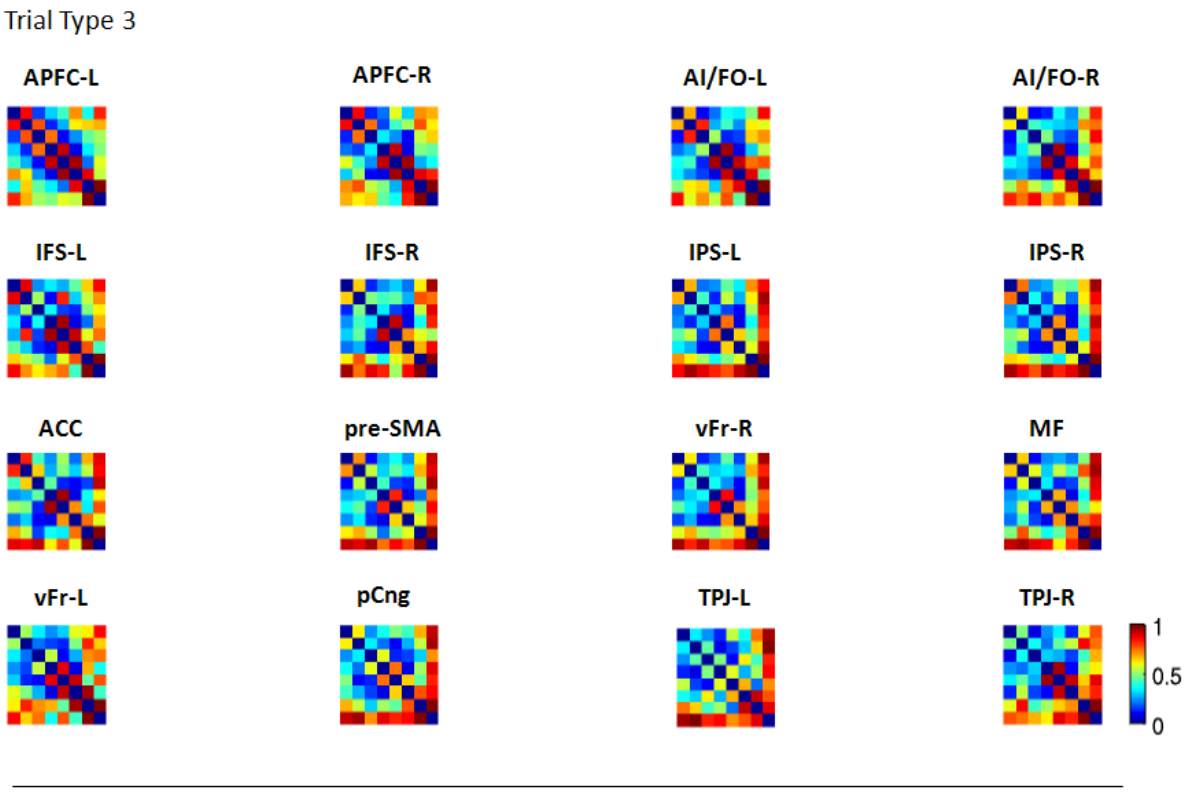

Figure 5.2. RDMs for the three trial types across various ROIs. In most ROIs, across the three trial types, the dissimilarities between the subtrials increased with the position of the subtrial in the trial, showing that the distinctness of the representation of latter subtrials was greater than earlier ones across different ROIs and trial types. APFC – Anterior prefrontal cortex, AI/FO – Anterior insula extending into frontal operculum, IFS – Inferior frontal sulcus, IPS – Intra parietal sulcus, ACC- Anterior cingulate cortex, pre-SMA – Pre-Supplementary motor area, vFr- Ventrolateral frontal cortex, MF – Superior Medial prefrontal, pCng – Posterior cingulated, TPJ – Temporoparietal junction.

To confirm this trend, these RDMs were correlated with a model RDM (model 1, Fig. 5.1b) that predicted a linear increase in dissimilarities between the subtrials across the length of the trial. In this model RDM as described earlier (see methods) the dissimilarity between any two subtrials increased linearly with their position in the trial. The pattern of values in this RDM was one that would be expected if the dissimilarity between subtrials was solely determined by their position in the trial. The model considered the entire trial to be cohered as one plan, and predicted that the chunk level organisation of the subtrials had no effect on their representation.

ROI	Model 1			Model 2		
	Trial type 1	Trial type 2	Trial type 3	Trial type1	Trial type 2	Trial type 3
APFC-L	0.05	0.36	0.20	0.37*	0.20	0.02
APFC-R	0.46*	0.4	0.37	0.5*	0.27	0.02
AI/FO-L	0.48*	0.40*	0.42	0.47*	0.27	-0.03
AI/FO-R	0.60**	0.52**	0.61***	0.37	0.21	0.18
IFS-L	0.74***	0.55**	0.42	0.46	0.47*	-0.14
IFS-R	0.83***	0.65***	0.55**	0.4	0.27	0.25
IPS-L	0.79***	0.56**	0.62**	0.39	0.23	0.25
IPS-R	0.68***	0.55**	0.64**	0.4	0.01	0.31
ACC	0.79***	0.53*	0.42	0.44	0.23	0.17
pre-SMA	0.81***	0.61**	0.53**	0.46*	0.62*	0.24
MF	0.84***	0.71***	0.55**	0.52*	0.43	0.26
vFr-L	0.47*	0.43*	0.52**	0.43	0.28	0.10
vFr-R	0.70***	0.67***	0.66***	0.34	0.32	0.33
pCng	0.63**	0.56**	0.70***	0.45	0.17	0.43
TPJ-L	0.59*	0.54*	0.60**	0.39	0.25	-0.05
TPJ-R	0.73***	0.62**	0.67***	0.40	0.30	0.29

Table 5.1. Spearman's correlation coefficients between the ROI and model RDMs 1(*model 1*) and 2 (*model 2*). The coefficients for trial type 1 are the highest for both models. For all trial types model 1 is the better predictor.

Table 5.1 (model 1) shows the Spearman's rank coefficients obtained at each ROI, across the three trial types. This model was well correlated with the pattern of dissimilarities from all trial types in all ROIs except APFC and left AI/FO. This shows that there indeed was a pattern of increasing dissimilarities across the length of the trial in all trial types. In other words, there was a significant trend of increase in the informational content/representation of the subtrials with increase in their position.

Notably, this trend was also present in trial type 3, even though it was not explicitly organised. Recall that this trial type had shown minimal evidence of increase in activity across its length, from which it was concluded that this trial type did not show the evidence of organisation that other trial types had shown. Current results show that the representation of subtrials in this trial type changed in manner similar to the trials that were more explicitly organised. Trial type 3 was organised in that it was known to the subjects that it consisted of eight subtrials, albeit this organisation was less explicit than other trial types. While such organisation did not affect the mean level of activity between subtrials 1 to 7, it did affect the information content in the pattern of voxels. (For direct comparison across trial types, see below).

In earlier results the extent of increase in activity across the eight subtrials was affected by the organisation of the task episode; specifically, the evidence of such an increase was greater in trial type 1 than in 2. Evidence suggested the former to be organised more as a single task episode, while the latter seemed to be organised more as two task episodes. In the current results, although all trial types show evidence of increase in representation across the eight subtrials, this was strongest in trial type 1, the coefficients for which exceeded those of trial type 2 and 3 in most ROIs.

An additional result is the effect of laterality, with right hemisphere coefficients being generally higher in all three trial types.

While the above findings show parallels with increase in activity across the eight subtrials seen earlier, differences are to be noted. In the last chapter, ROIs like left IFS and IPS did not show any trend of increase in activity across the trial; however both regions were highly correlated with a trend towards increase in representation along the eight subtrials. In the same vein, as mentioned earlier, the third trial type had shown minimal trend of increase in activity across the subtrials, but show good evidence of increase in representation across the subtrials. These results suggest that, while increase in representation has a similar trend to increase in activity, the two are not identical.

Next, the ROI RDMs were correlated with a second model RDM that predicted a different trend in the pattern of dissimilarities across the subtrials. This second model was based on another feature of the pattern of their elicited activities seen in the previous chapter. Recall that activity in some regions (especially. in trial type 2) increased in parallel across the subtrials one to four and five to eight. This was most prominent in trial type 2 in bilateral insula, IPS, left IFS and APFC, and pre-SMA. Since the pattern of monotonic increase of activity seen best across the eight subtrials of trial type 1 was strongly paralleled by a similar monotonic increase in representation, it is plausible to think that the biphasic increase in activity seen across ROIs in trial type 2 might also lead to a similar pattern of change in representation of the subtrials.

To test this, a second model RDM (model RDM 2, Fig. 1c) was created. It predicted a stepwise increase in the dissimilarities of the subtrials one to four and five to eight, but the average representation of the subtrials composing the two chunks was identical, so the dissimilarities of subtrial 2 with subtrials 1, 3 and 4 was same as that of subtrial 6 with subtrials 5, 7 and 8, and similarly subtrial 3 vs 1, 2 and 4 was the same as of 7 vs 5, 6 and 8; in other words, the pattern of dissimilarities across the two chunks were parallel. It was expected that this model RDM would be

more strongly correlated with trial type 2 than with trial type 1. Note that this model RDM did not make a prediction about the entire extent of the ROI RDM, since the hypothesis cannot predict what the dissimilarities between the subtrials of the two chunks would be, relative to the dissimilarities of the subtrials within a chunk. So the corresponding parts of ROI RDMs were masked prior to their correlation with this model RDM.

The Spearman's coefficients are listed in Table 1 (model 2). Note that these are lower compared to that for model 1 in most ROIs across the three trial types. In trial type 2, only pre-SMA and left IFS were significantly correlated with this model; interestingly, for this trial type, in none of the ROIs was this model a better predictor than model 1. However, in trial type 1 the bilateral APFC was better correlated with model 2 than model 1.

Contrary to the expectation from the pattern of the level of activity across subtrials seen in the previous chapter, across most ROIs model 2 was a better predictor of pattern of dissimilarity in trial type 1 than in trial type 2 (pre-SMA being a notable exception). It is likely that this results from the fact that the two model RDMs were correlated with each other, and trial type 1 being highly correlated with model 1 was also better correlated with model 2. This result raises a possibility that the pattern of change in representation of the subtrials was minimally affected by their organisation into distinct chunks.

Whole Brain Analyses

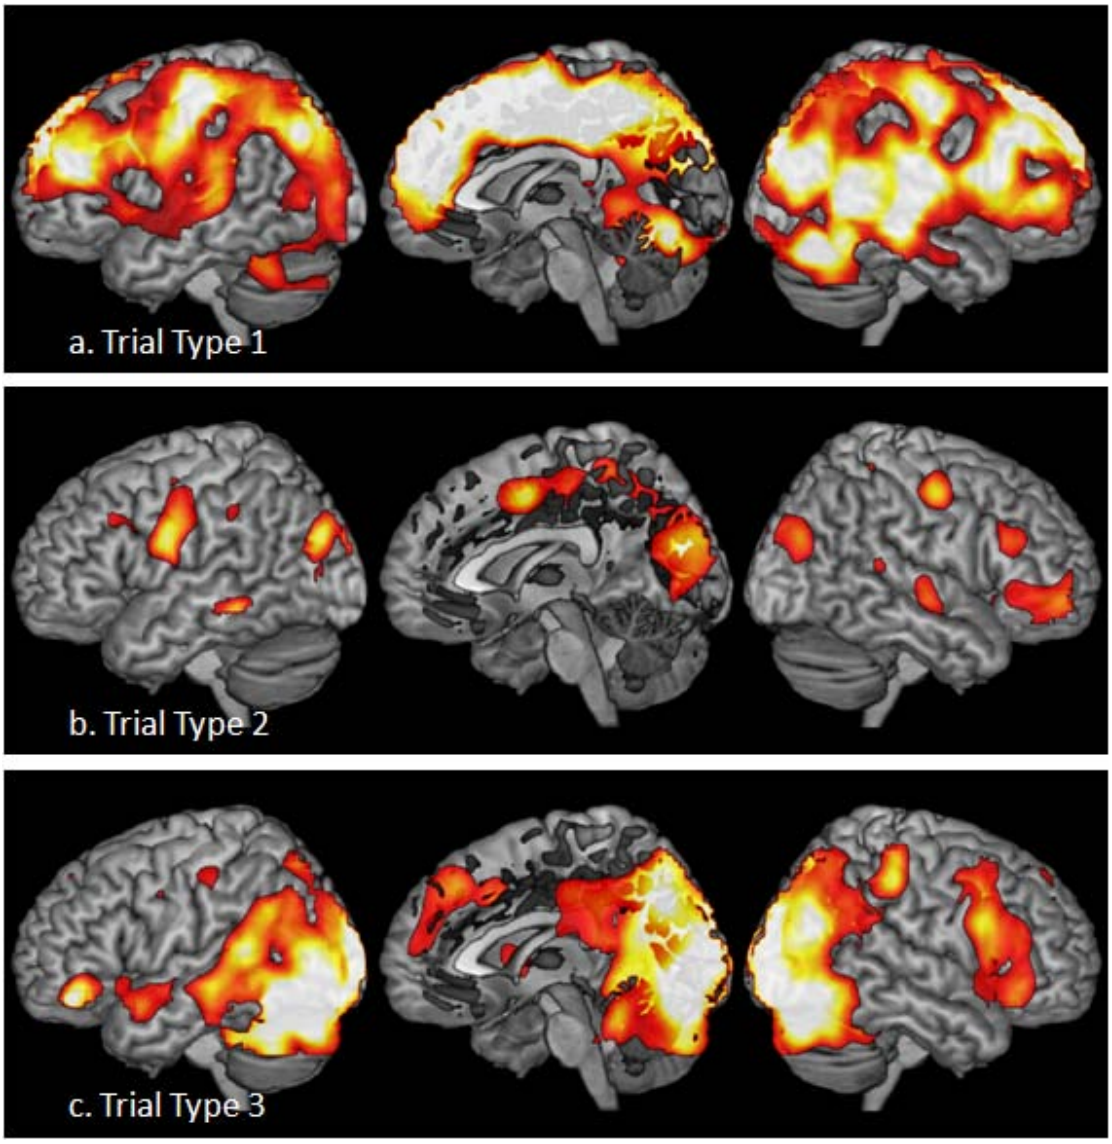

Figure 5.3. Regions where the actual pattern of dissimilarities across the eight subtrials were highly correlated with those in model 1.

To investigate other regions that showed similar effects of organisation on the representation of the subtrials, a whole brain searchlight looked at regions where the activity pattern across voxels was significantly correlated with the two models (see methods). Fig. 5.3 shows the whole brain render of the regions that were highly correlated with model 1, i.e. where the representation of subtrials increased across

the eight subtrials. Note that the results are most extensive for trial type 1 (Fig. 5.3a), and involve very widespread brain regions – bilateral lateral prefrontal, premotor, sensorimotor, and superior parietal along with medial frontal, and parietal regions, right middle and superior occipital, right middle and superior temporal gyri and right lateral cerebellar hemispheres. While these regions are largely coextensive with the regions that had shown increase in activity across the eight subtrials of trial type 1, the difference is nonetheless noteworthy (c.f. Fig 4.10). Left lateral prefrontal, parts of pre-SMA and cingulate and left sensorimotor regions had not shown an increase in activity across the subtrials of this trial type; however, all these regions show an increase in representation of the subtrials across the length of trial type 1. Anterior temporal regions and left TPJ on the other hand had shown increase in activity but do not show any evidence of increase in representation.

The regions showing increase in the representation of the subtrials across trial type 2 (Fig. 5.3b) were far less extensive and were limited to the bilateral premotor cortex, right IFS, ventrolateral prefrontal cortex and APFC along with bilateral superior occipital cortex, middle cingulate, precuneus and cuneus.

The results for trial type 3 were qualitatively most distinct (Fig. 5.3c). The region with the strongest correlation was the occipital lobe and adjoining parts of temporal and parietal cortices and curiously cerebellum. Parts of parietal (IPS, superior parietal lobule, precuneus), frontal (right premotor, IFJ, AI/FO, medial frontal) and right sensorimotor regions also showed significant results.

The second model, on the other hand, did not yield results that were statistically significant at the whole brain level.

Comparison of Trial types 1 and 3

Both trial types 1 and 3 consist of a single trial episode having a sequence of eight subtrials. However, trial type 1 had explicit higher level representations organising the lower level elements (subtrials). A comparison of the representations of the subtrials in these two trial types would therefore show the effect of the presence of higher level organising representations on the representation of lower level elements. In this regard, earlier results (Table 5.1) had shown that the evidence of increase in the representation of the subtrials with the progression of trial was stronger in trial type 1. However that analysis does not tell anything about the actual level of representation of the subtrials in these two trial types. In the analysis described in this section the two trial types were compared directly in a single model.

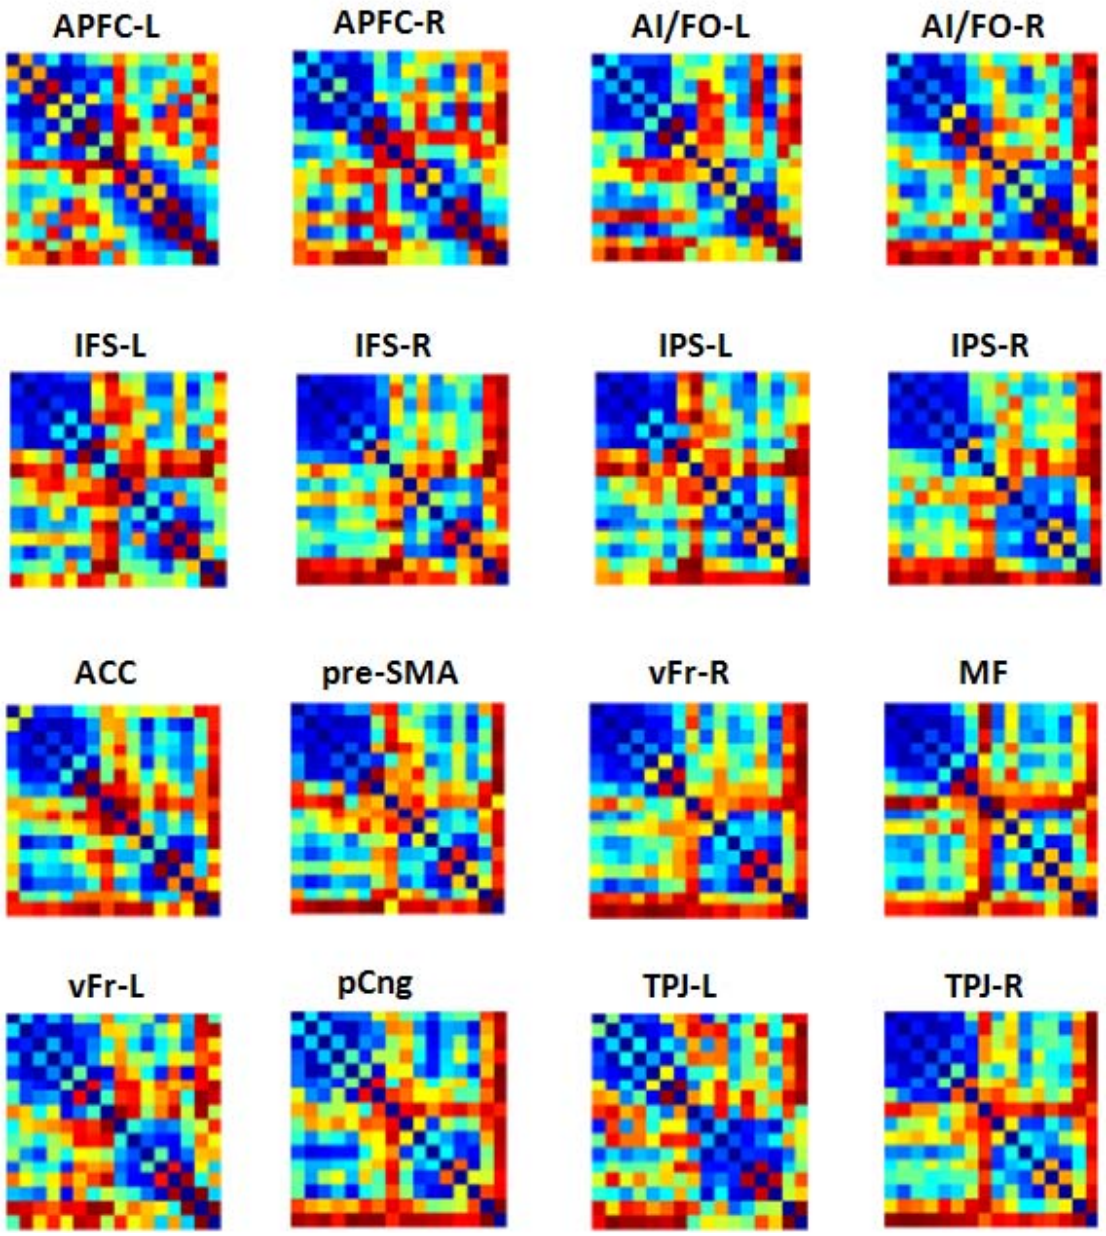

Figure 5.4. ROI RDMs with the dissimilarities between subtrials of trial type 1 and 3. The first 8 rows and columns represent the eight subtrials of trial type 1, the next eight represent those of trial type 3. Note that the dissimilarities between the subtrials of trial type 1 (top left quadrant) are generally lower than the dissimilarities between the subtrials of trial type 3 (bottom right quadrant)

For each ROI the various dissimilarities were assembled into a 16 x 16 matrix, the first eight rows and columns of which represented the eight subtrials of trial type 1, and the next eight rows and columns represented those of trial type 3 (Fig. 5.4).

The top left quadrant of this matrix thus contained the dissimilarities between the subtrials of trial type 1, while the bottom right quadrant contained the dissimilarities between the subtrials of trial type 3. It is visibly evident that the top left quadrant of many of these ROI RDMs is 'bluer' than the bottom right quadrant, suggesting that the dissimilarities in between subtrials of trial type 1 were in general lower than those between subtrials of trial type 3.

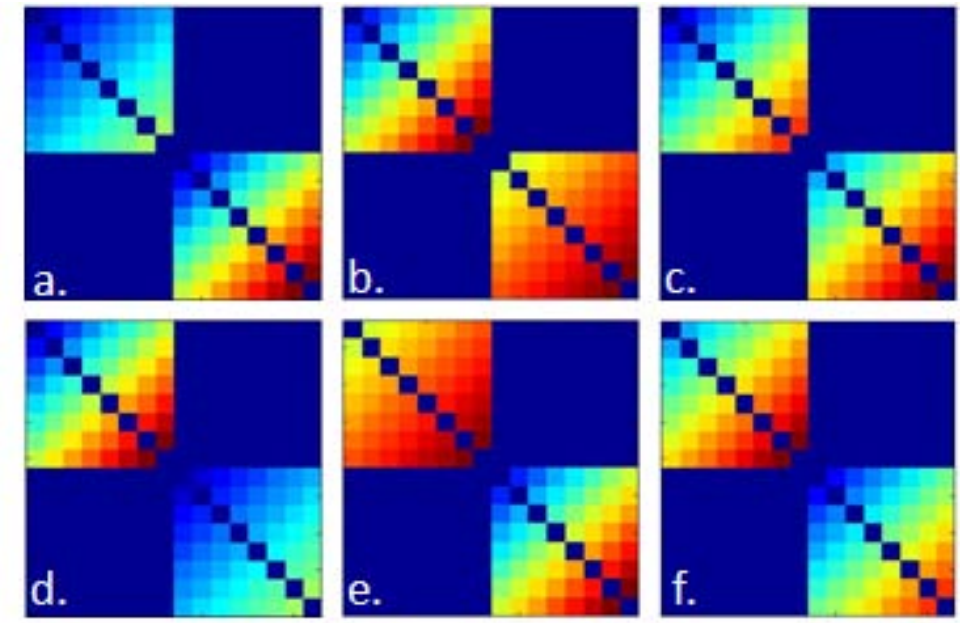

Figure 5.5. Model RDMs for testing the various kinds of relation between the dissimilarities of trial type 1 and trial type 3. (a) to (c) Three models predicting dissimilarity in trial type 3 to be greater than those in trial type 1. (d) to (e) predict the reverse.

Seven possibilities were tested using respective model RDMs (Fig. 5.5). Note that all of these models made predictions only about the dissimilarities in between the subtrials of trial type 1 (top left quadrant of ROI RDMs) and those in between the subtrials of trial type 3 (bottom right quadrant of ROI RDMs). The other quadrants were masked from analysis.

Three of these models predicted the dissimilarities in between subtrials of trial type 3 to be greater than those between subtrials in trial type 1 (Fig. 5.5a to 5.5c):

Model 1: In this model RDM (Fig. 5.5a), the lowest dissimilarities i.e. between the earliest subtrials were same across the two trial types, but the dissimilarities involving the later subtrials were greater for trial type 3. This model therefore also predicted that the rate of increase in dissimilarities with their position is greater in trial type 3.

Model 2: In this model RDM (Fig. 5.5b), the dissimilarities between the last subtrials were similar for the two trial types, but the dissimilarities between earlier subtrials were greater for trial type 3. Hence, as per this model, although the absolute level of dissimilarities was greater in trial type 3, the rate of increase was greater in trial type 1.

Model 3: All dissimilarities in trial type 3 were evenly greater than those in trial type 1 (Fig. 5.5c).

The next three models (4, 5 and 6) were analogous to the above but predicted the level of dissimilarities to be greater in case of trial type 1 (Fig. 5.5d to 5.5f). Model 7 (not shown) predicted the dissimilarities to be equal in the two trial types.

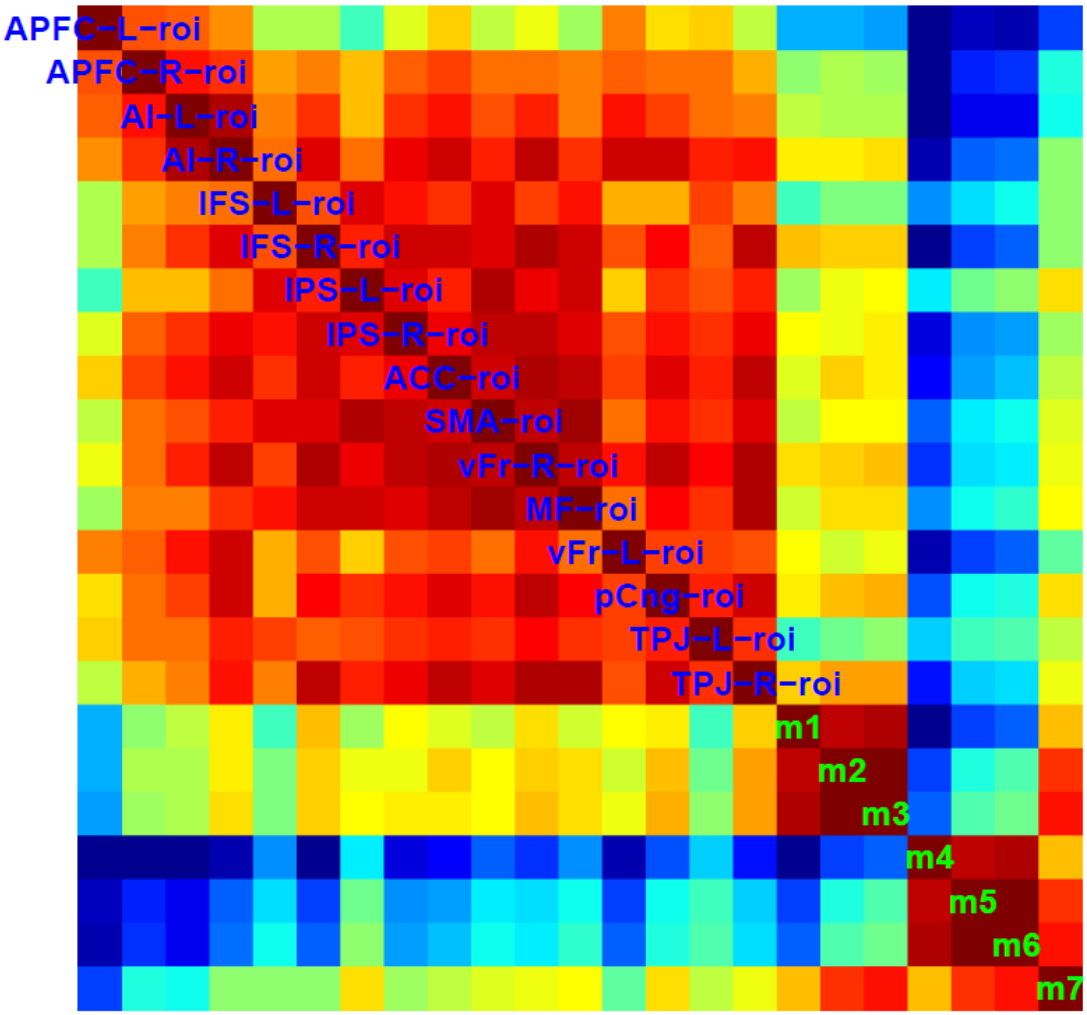

Figure 5.6. Cross correlation matrix depicting correlation coefficients (not dissimilarities) between the various ROI RDMs shown in fig. 5.4 and the model RDMs in fig. 5.5. Most ROI RDMs are intensely correlated with each other. The model RDMs predicting greater dissimilarities in trial type 3 (m1 to m3) are better correlated with ROI RDMs than those predicting the reverse to be the case (m4 to m6). For most ROIs, m2 and m3 are better correlated with ROI RDMs than m7 which predicted the dissimilarities to be identical in the two trial types.

Fig. 5.6 shows a cross correlation matrix showing correlation coefficients between the various ROI RDMs and the model RDMs (marked as m1 to m7). The first sixteen rows and columns represent the sixteen ROIs, the following seven represent the seven model RDMs. As is apparent the three model RDMs predicting greater dissimilarities between subtrials of trial type 3 (models 1 to 3) were far better

correlated with the ROI RDMs than those predicting the reverse (models 4 to 6). In most ROIs the models 1 to 3, especially models 2 and 3, were better correlated with ROI RDMs than model 7 that predicted the dissimilarities to be the same in the two trial types. These results suggest that the dissimilarities were higher in trial type 3 than in trial type 1, and show that the presence of higher level representations may reduce the distinctness of the representation of the lower level elements. Also note that the ROI RDMs are strongly correlated amongst themselves suggesting that the pattern of dissimilarities was very similar across them.

Results till now

Patterns elicited by task events representing the later phases are more distinct than those in earlier phases, and task plans organised as a unified task episode by explicit higher level representations show greater and anatomically more widespread increase in the distinctness of the representation of task elements. While such task plans show greater increase, the actual magnitude of the distinctness of task elements in them is lower compared to those plans that do not have higher level organising representations, suggesting that the organising higher level representations decrease the distinctness of individual lower level task elements. Two possibilities exist – either the organising representation *binds* the lower level representations into a common representational code which results in them being less distinct, or the presence of higher level representations leaves little representational space in the limited capacity neurocognitive reserves for the lower level task elements. This latter possibility could potentially also explain the increase seen in the distinctness of the later task elements, since as greater portions of higher level representations are discarded, more space becomes available for the representation of lower task elements. However, this thesis does not explain the increase in the distinctness of lower level task elements seen across the phases of trial type 3.

Analyses so far have looked into the changes in representation of task elements that represent the phase of the task (category 1 information). It is therefore not clear if the increase in representation also affects information contained in the subtrials that is unrelated to the phase of the task (category 2 information). To look into this issue, the next set of analyses investigates change in the representation of the relevant rule across phases of the task.

5.3.2 Effect on Category 2 information

At each subtrial one of the two rules was relevant (called 'A' and 'B', see methods); on this basis the subtrials at each of the eight positions could further be classified as subtrial A or B. The dissimilarity between the pattern of activity elicited across the voxels of an ROI by subtrials A and B at each of the eight positions was calculated for all three trial types, for all ROIs. These values were then z-transformed. Graphs in Fig. 5.7 shows the measures of dissimilarity between A and B at each of the eight positions of trial type 1. While there is a definite trend towards increase in dissimilarities between A and B across the eight position, the increase, in general, is not as smooth and monotonic as was the case with the level of activity elicited by the subtrial (Fig 4.7, chapter 4) and the dissimilarities in between the subtrials (Fig. 5.2).

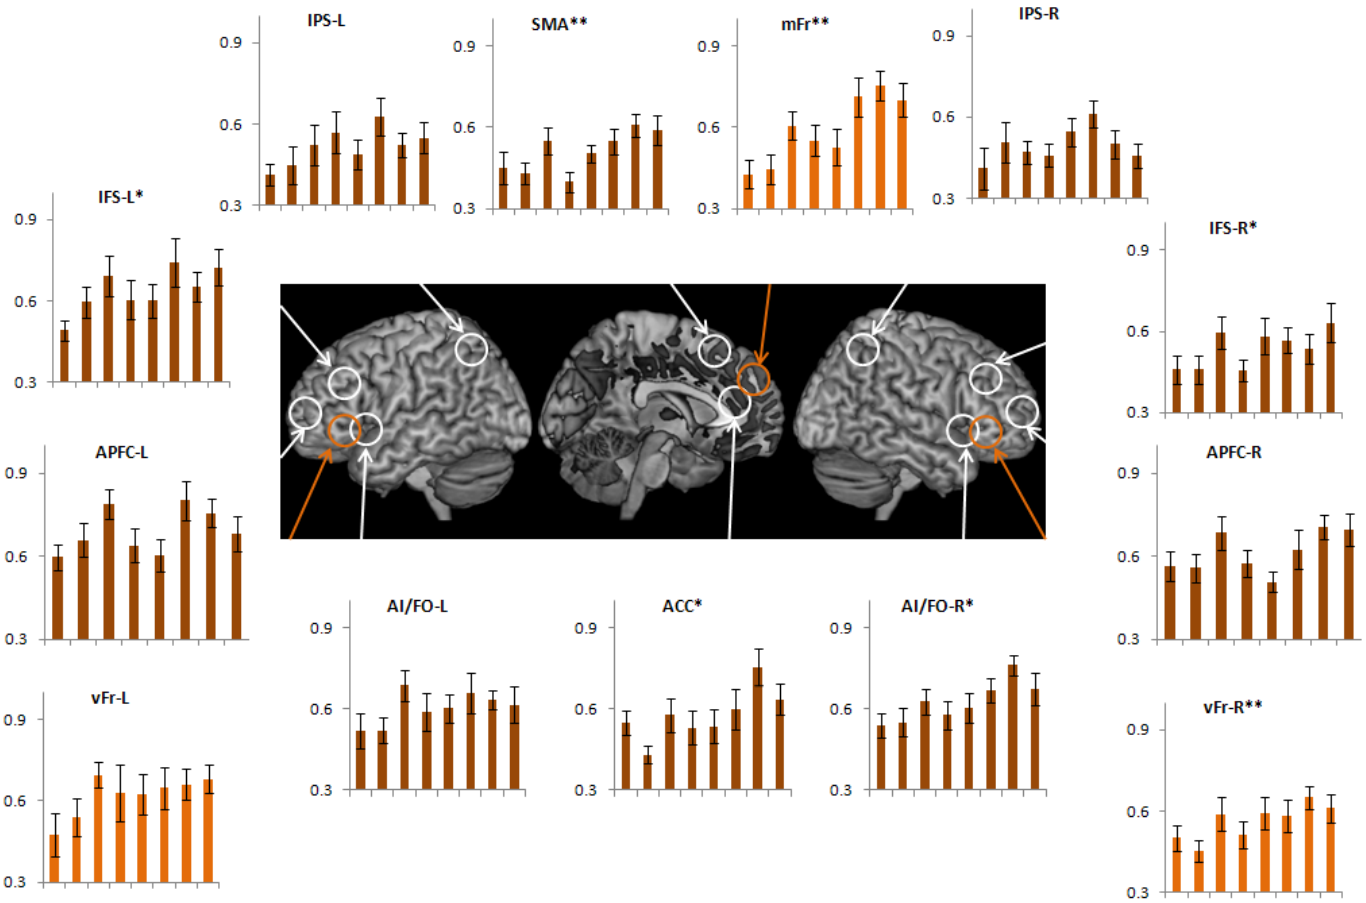

Figure 5.7. Dissimilarities between subtrials A and B at the eight positions in trial type 1. In many ROIs, these were higher for the later positions. Asterisks marks the ROIs where the dissimilarities in the first four positions were significantly lower than those in last four (* : $p < 0.01$; **: $p < 0.001$).

To look for the effect of position on the dissimilarity between the two rules, the dissimilarities in the first four positions were compared with the second four in a repeated measures ANOVA. This was significant in medial ROIs (MF, pre-SMA, ACC), right vFr, and bilateral IFS ($F_{(1,12)}>5.1$, $p<0.05$). Recall that these ROIs were the ones that had shown the strongest evidence of increase in category 1 information (Table 5.1), current result shows that these regions also show the greatest tendency for increase in category 2 information.

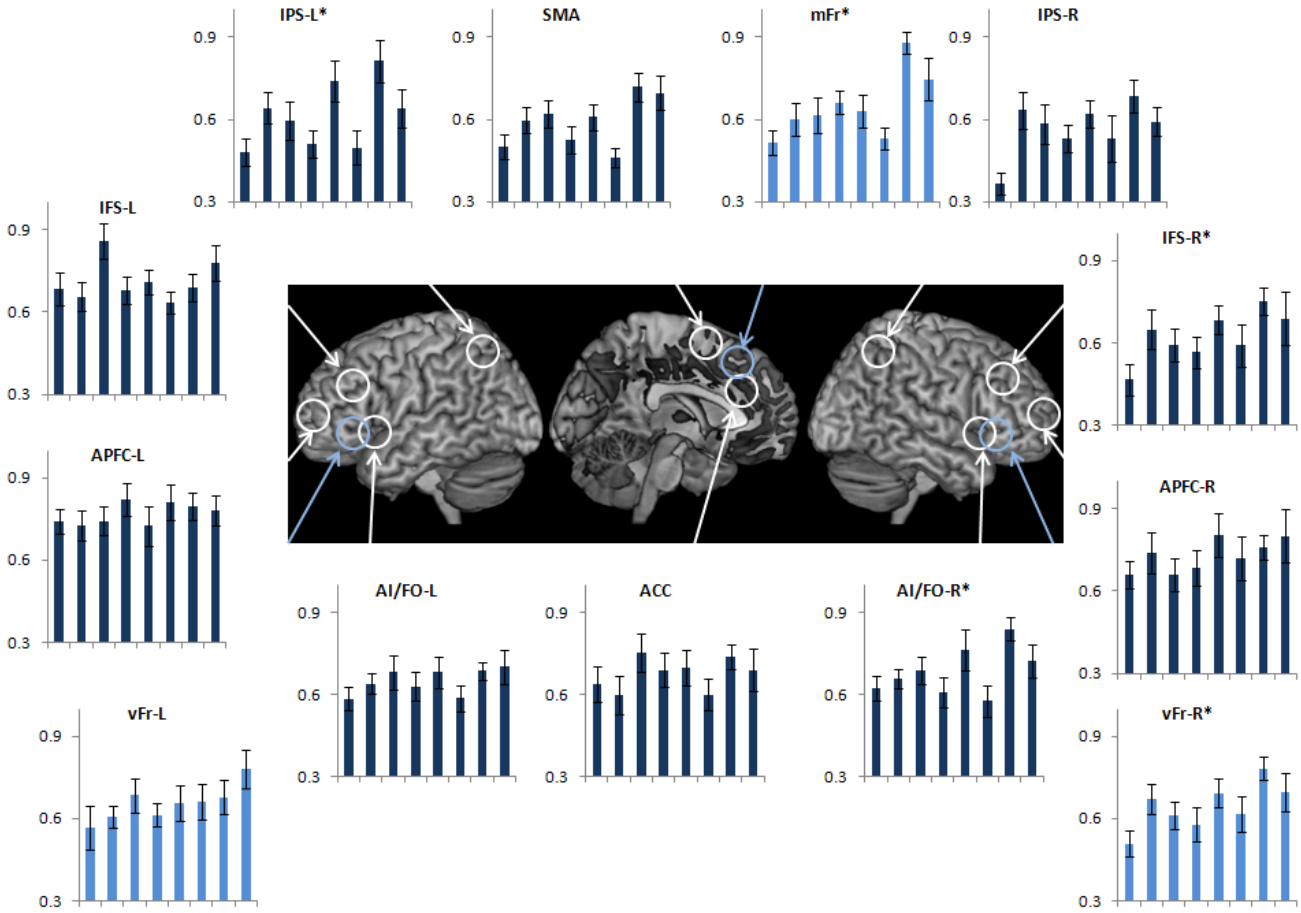

Figure 5.8. Dissimilarities between A and B across the eight positions in trial type 2.

Similar analysis in trial type 2 (Fig 5.8) showed significant results in right IFS, AI/FO, vFr, left IPS and MF ($F_{(1,12)}$>4.7, p<0.05). In contrast to these two trial types, in trial type 3 (Fig. 5.9) none of the ROIs showed significantly greater dissimilarities between A and B during the second half of the trial.

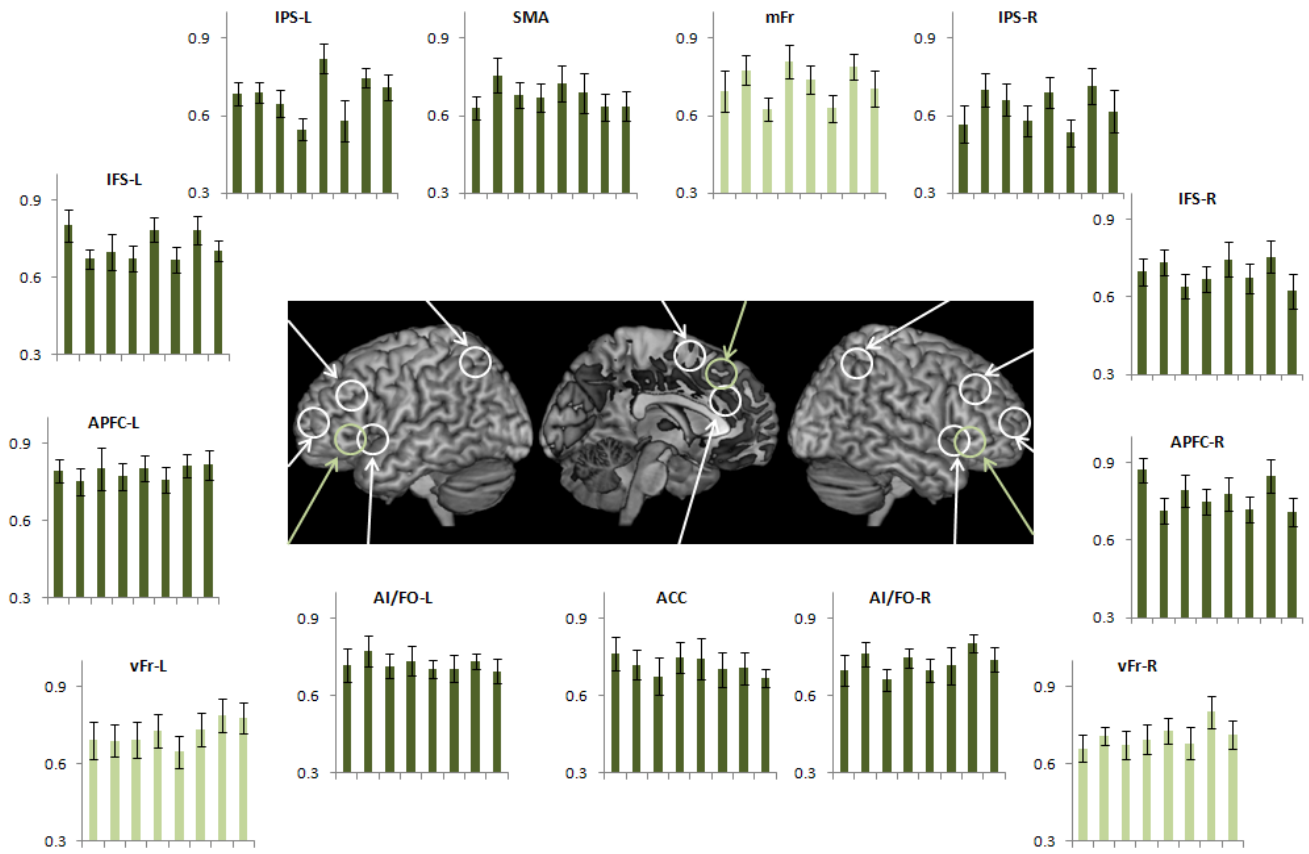

Figure 5.9. Dissimilarities between A and B across the eight positions in trial type 3.

Whole brain analyses

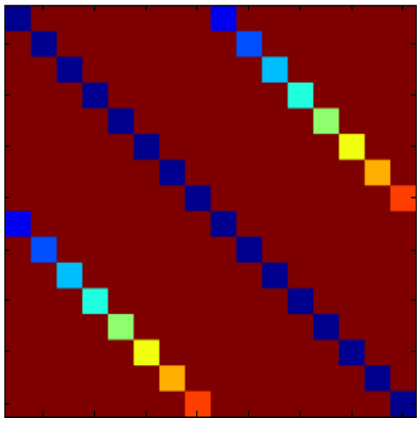

Figure 5.10. Model RDM used for whole brain searchlight for regions where the dissimilarity between subtrials A and B increased across the eight positions. The RDM assembled to correlate with such a model had the first eight rows and columns representing subtrials A across the eight positions, and the next eight rows and columns representing the subtrials B across the eight positions.

For each trial type, a whole brain searchlight was carried out to look for regions where the dissimilarities between rules A and B increased across the eight positions. At each step of the searchlight, an RDM (16 x 16) was created based on the pattern across voxels in the searchlight. The first eight rows and columns of this RDM belonged to the subtrials A across the eight positions, and the next eight rows and columns represented the subtrials B across the eight positions. The model RDM (Fig. 5.10) used was congruent to this RDM. However, all positions in the model RDM were masked except the eighth diagonals representing the dissimilarities between subtrials A and B at the eight positions. The value in these cells increased linearly. This model RDM thus predicted that the dissimilarities between A and B increase across the eight positions. The model RDM was correlated with the searchlight RDM created at every step of the searchlight and the dissimilarity between the two was z-transformed and assigned to the voxel forming the centre of the searchlight.

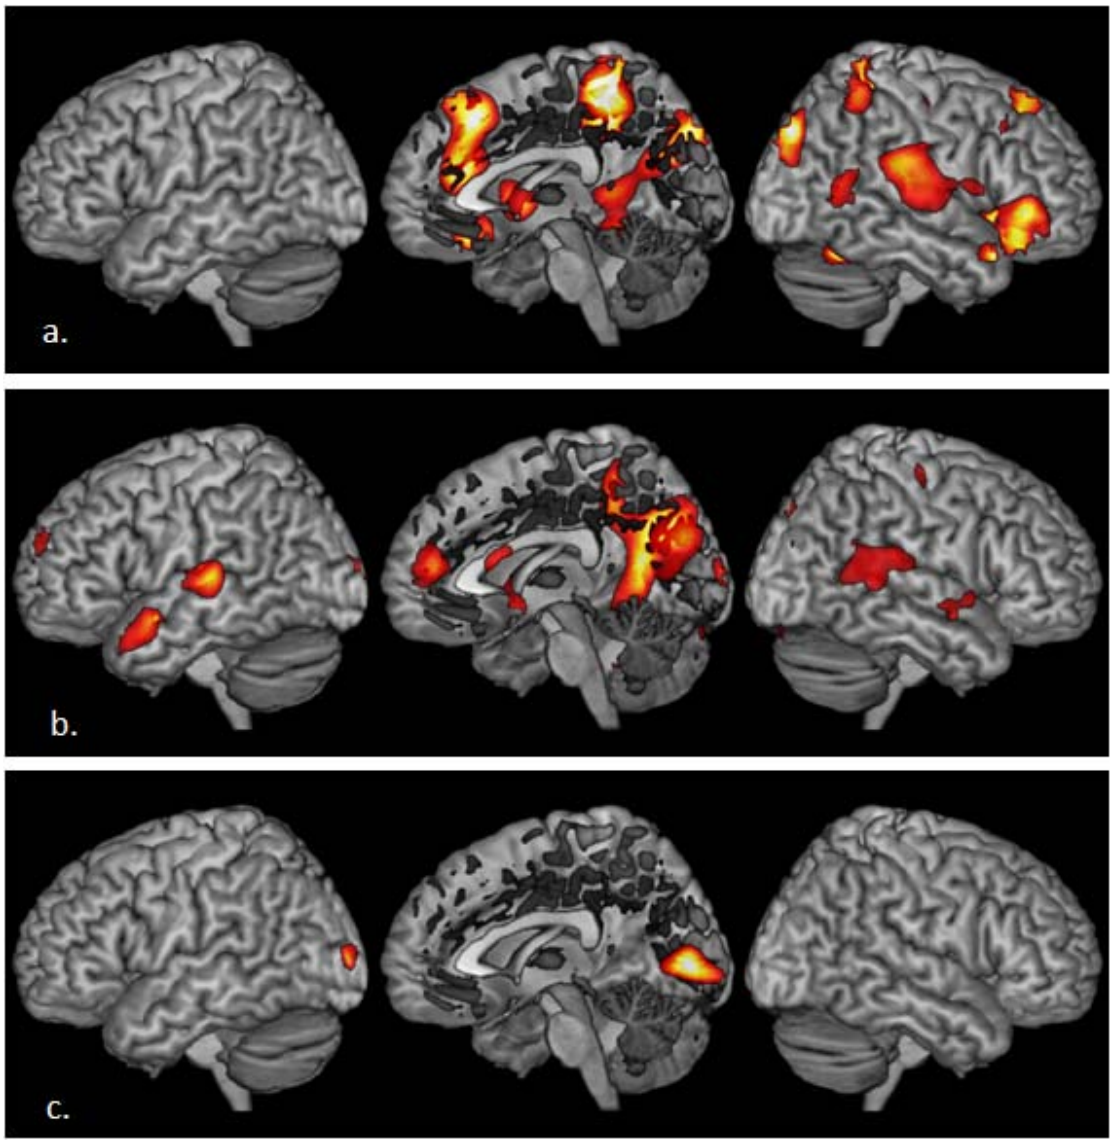

Figure 5.11. Regions showing significant increase in the dissimilarity between A and B across the eight positions in trial types 1 (a), 2 (b) and 3 (c).

Figure 5.11 shows the regions where this model RDM was highly predictive of the pattern in the brain across the subject population. As was found in ROI analysis, the increase was most widespread in trial type 1 (Fig. 5.11a), where it involved the right AI/FO, vFr, pre-SMA, ACC and the precuneus, along with the right angular, posterior superior temporal gyrus, parts of right sensorimotor, superior occipital and

parieto-occipital cortices, and parts of the cuneus (coordinates in table A5.1). These results were right lateralised.

Trial type 2 (Fig. 5.11b), on the other hand, showed significant results most strongly in medial parietal (precuneus and cuneus) and ACC. Loci of significant trend were also found bilaterally along the superior temporal gyrus and sulcus (coordinates in table A5.2). In the case of trial type 3 (Fig. 5.11c), the only region to show significant trend of increasing dissimilarity was the primary visual cortex along the calcarine sulcus.

Comparison of trial types

Lastly, the level of dissimilarity between subtrials A and B was compared across the three trial types to see if the higher level organising representations affected the representation of the relevant rule. For the current analysis, the dissimilarities between subtrials A and B for the first four (1 to 4) and the second four (5 to 8) positions were averaged. These two values are referred to as the dissimilarity between the two relevant rules in the first and the second chunk respectively. Fig. 5.12 shows these two values for the three trial types across the various ROIs, presented in lighter (first chunk) and darker shades (second chunk).

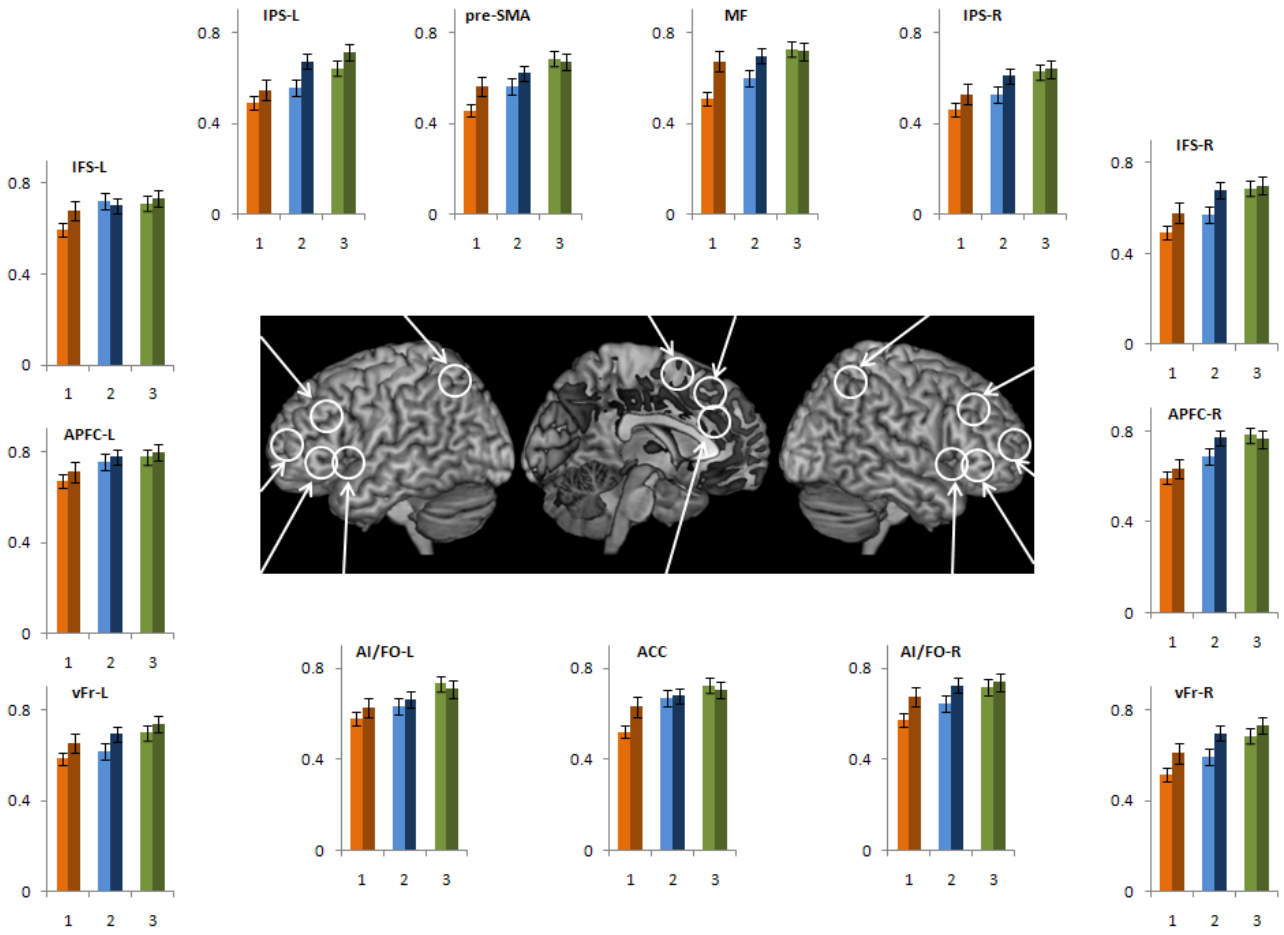

Figure 5.12. Comparison of dissimilarities between A and B averaged across positions 1 to 4 (lighter shades) and 5 to 8 (darker shades) from the three trial types.

Note that in almost all ROIs the level of dissimilarities averaged across the two chunks is highest in trial type 3 followed by 2 and 1. The level of dissimilarity in the first and second chunk differs minimally in trial type 3, while in the other trial types the dissimilarity is higher in the second chunk in many ROIs. A repeated measures ANOVA looking for the effect of trial type was significant in all ROIs ($F_{(2,24)}>3.6$, $p<0.05$) except for the left IFS. This finding parallels the earlier result (Fig. 5.7) of greater dissimilarities in trial type 3 than trial type 1, and adds that the organisation of the episode and the presence of higher level representation also affect the representation of rule information associated with the lower level task elements.

Further, as would be expected from earlier analysis, right AI/FO, IFS, vFr, left IPS, and MF showed a significant effect of chunk ($F_{(1,12)}$>5.8, p<0.03). The interaction between trial type and chunk position reached significance only in MF and ACC ($F_{(2,24)}$>3.6, p<0.04).

As is apparent in Fig. 5.12, across all ROIs the average dissimilarities in trial type 2 were also greater than those in trial type 1. This was significant in right IFS and vFr, left IPS, pre-SMA and ACC ($p < 0.02$). Since both trial types 1 and 2 are organised (albeit differently), and differ in the cognitive load of the organising representations (8 letters vs. 4 letters), the above finding raises the possibility that the decrease in the distinctness of representations is a consequence of higher level representations occupying the limited capacity cognitive resources and leaving less neurocognitive representational space available for lower level task elements. This hypothesis can be tested. The first chunk had greater representational load in trial type 1, whereas the second chunk of the two trial types had similar loads. So if the decrease in distinctness is the consequence of load of higher level representations, the first chunks of the two trial types should differ, while the second chunks should not.

A repeated measures ANOVA between these trial types (1 and 2) with trial type and chunk as factors was carried out. The above hypothesis predicts a significant interaction between the effects of chunk and trial type. While the main effect of trial type was significant in right APFC, IFS, vFr, left IPS, ACC and pre-SMA ($F_{(1,12)} > 5.3$, $p < 0.02$), and that of chunk in right IFS, vFr, left IPS, pre-SMA and MF ($F_{(1,12)} = 6.6$, $p < 0.02$), none of the regions showed a significant interaction between the two ($F_{(1,12)} < 2.9$, $p > 0.1$). A second ANOVA considered all ROIs in the same model along with trial type and chunk as factors. Again, the main effects of trial type ($F_{(1,12)} = 9.8$, $p < 0.01$) and chunk ($F_{(1,12)} = 11.3$, $p < 0.01$) were present. However, neither an interaction between trial type and chunk ($F_{(1,12)} = 0.2$, $p = 0.6$), nor between trial type, ROIs and chunk ($F_{(12,144)} = 2.1$, $p = 0.3$) reached significance.

5.4 Discussion

The results described above show that the representation of task elements is affected by their organisation into coherent task episodes. Two aspects of this were seen. Firstly, the discriminability of information related to the component task elements increased as they got closer to the completion of their respective episodes. While this increase was best seen in trial type 1 which was organised strongly as a unified episode by higher level representations, other trial types that were organised differently – either as two distinct sub-episodes (trial type 2) or weakly organised (trial type 3) – also showed such increase. This finding thus links up with those of the previous chapters wherein many regions had shown an increase in activity for the later subtrials, and proves that this increase in fMRI activity is accompanied by an increase in the task element related information that can be derived from the region. It further makes it unlikely that the increase in activity seen earlier was the result of a non-specific disinhibition consequent to the accompanying decrease in cognitive load (Singh et al., 2008), or due to a non-specific increase in blood flow documented before expected events (Sirotin et al., 2009).

The second aspect of the current results shows that the average level of distinctness of representation of the task elements is lower in organised trials with greater load of higher level representations compared to those less organised and having less organising representations. The distinctness of both categories of information was lowest in trial type 1 and highest in trial type 3.

Many aspects of the current results and those of the previous chapter suggest that the effect of phase could be caused by the dynamics of the representational load of organised episodes. Since these representations form the control structures that control the execution of individual steps, as each step is executed, the representations and control machinery pertaining to that step can be discarded, freeing up the limited capacity neurocognitive resources that can then respond to

and represent information related to the subsequent new step. Likewise, such episodes with greater higher level representational load will have less representational space for lower level elements, resulting in lesser distinction between them.

The amount of representations discarded is clearly maximal in trial type 1, hence the greater and more consistent increase seen in such trials. Similarly, the load of higher level representations was highest in this trial type, hence the lesser distinction between the task events in this trial. In trial type 3, the absence of organising representation meant that fewer representations were discarded at every step; hence less neural space got freed at every subsequent subtrial, therefore the smaller increase in activity across the trial.

Not all aspects of results, however, are easily amenable to the above thesis. It is less clear, why there should be any increase in activity/distinctness across the subtrials of trial type 3. Another less clear finding is that in many ROIs the dissimilarities during the *second* chunk was greater in trial type 2 than in trial type 1 (Fig. 5.12), in spite of the load of higher representations being identical in the two trial types at this stage. Although, in defence of this thesis, it can be argued that the representational load during the second chunk of trial type 2 is still less than in trial type 1. Even though there are equal number of letters maintained, these letters are a *repeat* in trial type 2, and hence their representation might be easier compared to the case in trial type 1 wherein the four letters are entirely new compared to the previous chunk.

Widespread regions showing linear increase in the discriminability of subtrials (especially in trial type 1) is an intriguing result (Fig. 5.3a). If the above mentioned thesis has some truth, it will imply that very widespread regions are involved in representing task related representations. Previous analyses have attributed task related representation to relatively limited regions (Fox et al., 2005; Dosenbach et

al., 2006; Duncan, 2010). The current findings show that beyond such task positive regions, task related information is also present in the default mode network (medial prefrontal, posterior cingulate, precuneus, IPL, TPJ, STS), task-relevant sensory and motor regions (left somatosensory and motor, bilateral occipital) and task irrelevant sensory and motor regions (right somatosensory and motor, bilateral primary auditory) regions. Further, the pattern of change in the discriminability of such information is identical across these functionally varied regions.

The evidence of linear increase was much limited in trial type 2 (Fig. 5.3b). Results of the previous chapter suggested that in many regions these trials were represented as two separate chunks. It is possible that the poorer linear increase in discriminability in such trials is related to their different pattern of organisation. In fact many of the regions that show increase in discriminability of the subtrials of trial type 2 in the current analysis had shown increase in activity across the two chunks in the previous chapter (Fig 4.16).

Results of trial type 3 are least clear. The pattern of these results was very different from those of trial type 1. Intriguingly, these were centred on the occipital regions, although regions with significant effect could be found in frontal and temporal regions. Indeed peak of the effect was in the primary visual cortex, which, in case of trial type 1 had shown no effect in the whole brain analysis. This suggests a possibility that the brain regions involved in representing the organised and unorganised trial types were very different.

Comparison of the two analyses

The results described in the previous and the current chapter largely parallel each other, although interesting differences can be noted. The increase in activity across the length of the trial seen in the previous chapter is paralleled by the increase in representation seen in the current analyses. In both, the increase was

best seen in trial type 1 and relatively less in trial type 2. However, while none of the ROIs had shown an increase in activity across trial type 3, most of them showed an increase in the representation of category 1 information. Dissociations between the two can also be seen in trial type 1. A number of regions had not shown an increase in activity in the last chapter but do show an increase in information. For example, compare parts of left lateral PFC, superior parietal lobule, preSMA and ACC in Fig. 5.3a and Fig. 4.10. It is possible that some of these dissociations could be related to the greater sensitivity of multivariate methods (c.f. Harrison and Tong, 2009).

The results of these analyses have suggested new functional roles of many regions. Default mode network (medial prefrontal, posterior cingulated, precuneus, IPL and other areas around TPJ) is usually considered as a task negative region anti-correlated with the task positive fronto-parietal regions (Fox et al, 2005, Uddin et al., 2009). The current results differ from this view and show that in organised tasks the behaviour of the two can be similar. Indeed, as can be seen in Fig. 5.11, the similarity can go to the extent of both regions showing a parallel increase in the discriminability of the relevant rule.

A related result comes from Rogers et al., (2010). They had subjects do a transitive inference task between two pairs of visual stimuli that were derived from a remembered implicit sequence or were independent. They found greater functional connectivity between regions of multiple demand and default mode networks when subjects had to access the underlying sequence to do the task, as if greater functional interaction is needed between these regions when accessing hierarchically organised representations!

Some earlier studies have found correlated activity in the DMN and task positive regions. Christoff et al., (2009) found the two to be correlated during the episodes of mind wandering; Spreng et al., (2011) found this to be the case during autobiographical planning. In all such cases the increase in DMN activity was related

to the internal direction of mental focus, while the activity in the task positive regions was related to the control exerted on to this internally directed thought. The current results are remarkable in that the correlated activity between the two occur in the context of an external task. Additionally, the increase in the representation of task related information proves that this increase is not due to enhanced mind wandering or an internal shift in the focus of attention.

The presence of task organisation related effect on neural activity and information in the primary sensory regions is another interesting finding. The old view that such regions process only the sensory information coming from external senses has in recent times seen revisions from findings that have shown the complexity of the functional properties of these regions. A number of studies (reviewed in Driver and Noesselt, 2007) have shown that primary sensory regions can respond to a stimulus from a different modality. For example, the primary auditory cortex shows response to visual stimulation alone (Calvert and Campbell, 2003; Kayser el al., 2007), visual enhancement of auditory stimulation (Kayser et al., 2007; Martuzzi et al., 2007) and tactile enhancement of auditory stimulation (Kayser et al., 2005). Primary sensory regions also show sensitivity to task related information, like the attentional context (Watanabe et al., 1998; Li et al., 2004), and the presence of reward (Shuler et al., 2006). In chapter 2, sensitivity of early visual cortex to the completion of task episode was shown. The current finding of the presence of information related to the task *organisation* in primary sensory regions, even those whose modality was completely irrelevant to the task (e.g. primary auditory cortex), is very unique. Further studies would perhaps clarify these issues.

CHAPTER 6

Discussion

The experiments described in this thesis suggest that the fronto-parietal representation of task events is dependent upon their position in the structure of the task episode. Experiments in chapter 2 showed that events completing defined task episodes elicited activity that depended upon the hierarchical level of the episode; for example, completion of tasks elicited greater activity than completion of subtasks. Later chapters showed that in addition to such end of episode activity, subsequent task phases elicit greater activity compared to the earlier phases. Further, such increase in activity was associated with increase in the content of discriminable information. Even more interesting has been the finding that such effects of episode structure were widespread across very diverse brain regions, which suggests that the information about the task structure is present in many regions. The current experiments, however, are insufficient to give a conclusive picture about the neural, cognitive and computational reasons behind these findings. This chapter will discuss the various issues and possibilities raised in the earlier chapters.

6.1 What are Task Episodes?

In much of the earlier discussions, the concept of what constitutes a task was not explicitly defined. It is important to mention that not all usages of 'task episode' would be synonymous with the sense in which this phrase has been used here. For example, the whole experimental session, the whole day (and by extension, all of life) could be regarded as one task episode. Clearly, the effects talked about in the current experiments do not apply to these; else the activity across sequential trials (and across the length of day or even across life) will continue to increase!

Task episode in the usage of this thesis refers to a *continuous* period of *organised* and *sequential* mental activity, elements of which have a means to end relationship and which proceeds towards a defined end. Another characteristic of such episodes is a continuous maintenance of higher level organising representations. The episodes end when such maintenance was no longer needed. Further, the episodes got parcelled into sub-episodes when only a part of these organising representations changed to another, while the rest remained constant across the sub-episodes.

Recall that this was, indeed, the defining feature of all *task episodes* in the previous experiments. The trial in all of these was designed and instructed as a period of continuous and organised mental activity. Chunking of cued representations was used to parse this episode into sub-episodes in chapter 2 and 4, whereas in chapter 3, the trial was inherently designed as two blocks of tasks.

This brings to light another way of conceptualising the current findings. While numerous studies have shown fronto-parietal representation of task-relevant stimuli and responses or information linking them like rules and decisions (reviewed in Duncan, 2010), the current experiments suggest that representations underlying the episode, in the context of which other task-relevant events take place, are also present in these regions.

6.2 End of Task Episode

All experiments showed maximum and most widespread activity at the end of the task episode. This involved widespread cortical and subcortical regions and, interestingly, was right lateralised in the prefrontal cortex. In this regard, the results of the first set of experiments (described in chapter 2) were somewhat different from those of the last experiment (chapter 4 and 5). In the former, the phasic increase in activity at the completion of the task episode was largely centred on the

multiple demand regions, with additional involvement of other regions like precuneus, cuneus and early visual areas seen and when the entire visual search episode ended. The last experiment, looking at the final phase of the episode, found much more widespread activity which in addition to multiple demand regions also involved default mode and sensory and motor regions.

A possibility is that this difference stems from their different designs, as a result of which these experiments looked at slightly different qualitative issues. The first set of experiments looked at the transient increase in activity at the completion of the episode, while in the latter experiments the estimates of activity came from the entire duration of the last step. While, it seems plausible that the regions showing maximum activity at the last step will include those that show an increase at the completion of the episode, they will also include other regions such as those that show sequential increase across the various phases of the task episode.

Another possible source of difference could be that these experiments had different kinds of ends. Recall that the final target detections in the first set of experiments ended the visual search, but not the entire trial, which happened only after the probe was responded to. On the other hand, the final step of the trial in chapter 4 marked the completion of the entire trial, after which subjects could rest. It is possible that with respect to the structure of their respective trial episodes, the end investigated in chapter 2 completed a hierarchically lower episode than was the case with the end investigated in chapter 4.

Findings do suggest that different kinds of ends could be represented differently. Consider, in the case of trial type 2 of chapter 4, the disparity between the activities at subtrial 4, the final step of the lower level episode, and subtrial 8, the final step of the entire episode (Figs. 4.12 & 4.13). Interestingly, regions showing greater activity for the fourth subtrial were largely centred on the multiple demand regions, and involved relatively fewer areas beyond them (Fig. 4.6). In fact the

pattern was close to that seen at the completion of task episode in chapter 2 (Fig. 2.6). On the contrary regions where subtrial 8 elicited greater activity were very different (Fig. 4.13) and closer in pattern to those seen at the final phase of other task episodes like trial type 1 (Fig. 4.9), and the second task block in chapter 3 (Fig. 3.9).

What causes the end of episode activity? Perhaps, the explanation might be different for the two kinds activities found at the end of task episodes. Some regions in certain task episodes showed increased activity only at the end; while in others this increase occurred in the context of a sequential increase in activity across the entire length of the episode. The former is dealt with here, while the latter would be discussed in detail in a subsequent section. While the current findings cannot specify amongst the various plausible reasons discussed below, it is also possible that different combinations of the following reasons are causative in different regions.

Change in the Organising Representations

It was discussed in chapter 2, that the greater increase in activity at the third target detection of experiment 1 was best explained by the fact that the first three target detection events were implicitly thought of by the subjects as belonging to the same task episode, while the fourth was considered a separate task; this meant that the representations organising the first episode terminated at the third target detection, which elicited the increase in activity. This explanation could apply to the ending of all kinds of task episodes. Further, as the case of chapter 2 exemplifies, the representations in question need not even be explicit.

Similar conclusions can be made from the analysis of trial type 3 in chapter 4. Recall that in spite of this trial type not having any explicit letter string organising its execution, the fourth subtrial in such trials elicited greater activity in a number of regions. Here again the cause was likely to be the change in subtle representations

that implicitly organised the first and the second group of four subtrials as separate episodes.

However, also note that the actual brain regions involved in these two cases were different. In the case of experiment 1 of chapter 2, the dominant region was right anterior prefrontal cortex, whereas in trial type 3 of chapter 4, left posterior prefrontal regions (around inferior frontal junction, Brass et al., 2005) showed the maximal effect. This suggests that the pattern of activity related to the change in organising representations is likely to depend upon the details of the task episode since the actual brain representation of these representations can be different across different kinds of episodes.

Salience

Intuitively, the ends of episodes are more salient mental events compared to events within the episode. It is possible that some regions respond to this salience. Indeed many brain regions are known to respond to various kinds of salient mental events (Clark et al., 2000; Downar et al., 2000; Steven et al., 2000; Schall et al., 2003; Seelay et al., 2007). Across various studies this has included multiple demand regions, ventrolateral prefrontal cortex, medial parietal regions, TPJ, and non cortical regions like striatum and cerebellum.

Other studies have looked at issues that can be conceptualised as salient boundaries of mental episodes. Zacks et al. (2001) investigated the phenomenon of event segmentation; they had subjects passively view movies of some episode of purposive behaviour like making a bed; after passive viewing, each subject segmented these videos into episodes of coherent events by tapping a button to mark what they considered were the points where coherent segments began and where they ended. The authors then applied these segment boundaries to the corresponding moments in the passive viewing fMRI data, to identify transient

activity time locked to these boundaries in the passive viewing condition. They found that these boundaries elicited activity in a distributed set of regions – posterior prefrontal, IPS, precuneus, STS and lateral occipital cortices; recall that this pattern of activity is nearly identical to that found in the first set of experiments during the early target detections. Clearly, such event segmentation would be hierarchical; however, the authors did not report the variation of boundary activity with the hierarchy of the episode.

In another study, Sridharan et al. (2008) had subjects listen to musical pieces that had transitions between movements of symphony works. They found activity in multiple demand regions along with deactivation in the DMN regions at such transitions; a pattern very reminiscent of the results obtained in chapter 2. Additionally, Granger causality analysis showed that some of regions from the multiple demand group (ACC and AI/FO) played a critical and causal role in switching between the multiple demand and the default mode pattern of activity.

Default Mode Regions

A set of brain regions are known to decrease their activity during task blocks – medial prefrontal and medial parietal regions, TPJ, STS (Buckner et al., 2008); the end of such blocks would accordingly lead to an increase in activity in these regions. Some of these regions were indeed present in the current results. However, this explanation has to contend with the fact that this increased activity in default mode regions occurs along with increased activity in multiple demand regions. In the current understanding, what activates one deactivates the other. For example, if change in the organising representations and the accompanying salience of the mental event are the elicitor of multiple demand activity, these are also supposed to be the deactivators of default mode activity. Indeed, the default mode explanation for the increased activity at the end of the task episode presumes that these regions were deactivated earlier during the task blocks for precisely these reasons. A

potential, but currently unverifiable explanation could be that the default mode regions do transiently deactivate at the end of task episodes in response to events that cause activity in the multiple demand regions. However this transient deactivation in DMN is so transient that the fMRI estimate of activity is dominated by their succeeding sustained reactivation.

6.3 Phase Effect

It is less clear why the activity in a number of regions should increase across the steps of the task. As discussed in the earlier chapters, none of the currently plausible explanations can account for the results completely. However, across the three sets of analyses on this issue (chapters 3, 4 and 5), the hypothesis that seems to account for most findings is that the decrease in some form of representational load with the progression of the task is responsible for the increase in activity. This hypothesis will now be explained in greater detail.

One of the foundational claims of the cognitive paradigm in psychology is that sequential tasks require representations at multiple levels (Lashley, 1951; Miller et al., 1960) – representation of individual behavioural/mental events making up the sequence at the lower level, which are organised by subsuming representations at higher levels. While the representations at the various levels can be conceptualised as separate, evidence suggests that the higher-level representations could be made up from the elements of lower level representations. For example, executing events marking the beginning of task sequence always takes much longer than executing identical events within the sequence (Allport and Wylie, 2000; Schneider and Logan, 2006). Conceivably, this is due to the extra requirement of constructing the higher-level representation at the beginning. However, across multiple domains it has been found that this increase at the beginning is dependent on the nature of the sequence, i.e. it takes longer to start executing more complicated sequences.

Schneider and Logan (2006) found that during the execution of a list of tasks, the reaction time at the first step is always the highest and is directly related to the complexity of the list. Across many studies looking at simple movements (like finger tapping), the reaction time is always the highest at the beginning of the chunk (Povel and Collard, 1982; Restle and Burnside, 1972; Rosenbaum et al., 1983). In the same vein, simple reaction time initiating a motor action tends to depend upon the complexity of the motor action (Henry and Rogers, 1960) and it takes longer to start speaking a word that has many syllables than a monosyllabic word, the effect being present even when pictures are being named (Klapp et al., 1973). Similar findings come from memory retrieval studies. Anderson and Matessa (1997; see also Anderson et al., 1998) found that when subjects memorise digit sequences that were grouped into chunks during presentation, the pattern of RT during recall showed prolongation at the start of the chunk, the magnitude of which depended upon the length of the sequence.

Common across the findings reviewed above is the fact that complexity of the future sequence determines their initiation times, which suggests that planning the sequence may require access to, and representation of, some information related to each element of the sequence. Plausibly, these form the scaffold for the higher-level organising representation that will organise and control the execution of the episode. In other words the organising representations controlling the execution of the sequence could be built from the information related to individual steps of the sequence.

The analogy of this view with the 'memory drum' theory of action of Henry and Rogers (1960), whose finding was mentioned above, is noteworthy for the current discussion. As per the their view, the information for performing the *entire* motor action is constructed and stored in the form of *motor memory*, and the time to initiate the action is the time taken to transmit this motor memory to the

effectors. More complicated actions require more complicated memory structure, which takes longer to be created and transmitted to the effectors.

Given that across a variety of sequential behaviours – whether it is executing a sequence of motor acts or a sequence of abstract tasks or recalling a list of memorised items – the profile of higher effort and reaction time at the beginning of the sequence is common, it is suggestive that all sequential behaviours require construction of control structures that are based on representations pertaining to the entire sequence. The tasks used in the current experiments required sequencing, not of simple motor acts or stimuli, but of more abstract cognitive sets. It is likely that the beginning of such episodes were associated with the construction of a structure (or a plan) built on the representations from the entire sequence of cognitive sets. While the neural regions representing this structure are not known, it is likely that limited capacity neurocognitive resources that are thought to be engaged by other task control structures (Pashler, 1999), will be involved.

What would be the dynamics of this structure across the task episode? It can again be speculated that as individual steps are executed, the representations pertaining to them can be removed from the main control structure. Consequently, the representational load of this structure would decrease as more and more steps of the sequence are executed, making this structure *lighter* with each step, and hence freeing up more and more of the limited capacity neurocognitive reserves. Hence, more of such resources would be available to respond to and represent the events of each step.

Most of the current findings can be accounted for by this explanation. Not only did the sequential tasks show increase in activity across the sequential steps, the evidence for increase was greater in tasks where greater representational load would be discarded at each step. Hence in chapter 4, the increase was greater in trials with explicit organising representations, trial types 1 and 2, than in trial type 3

that had no such explicit representations. Task steps separated by a greater number of steps showed greater disparity in activity; in chapter 3, the two task blocks were consecutive, hence the disparity in activity between them (Fig. 3.9) is far less than between subtrials 2 and 8 of trial type 1 in chapter 4 (Fig. 4.7). Finally, the average estimates of representation of the individual task elements were lower on trials with greater load of organising representations. Thus in chapter 4, the average discriminability of information pertaining to a subtrial varied across the three trial types in inverse relation to the associated representational load:

trial type 3 > trial type 2 > trial type 1

Note that the term 'representational load' associated with the structure of a task episode in the current usage is not synonymous with the conception of working memory load (Baddley, 2007). As discussed in the earlier chapters, the dynamics of the phase effect cannot entirely be explained by the dynamics of the associated working memory load. A consideration of trial type 2 (chapter 4) makes this point clear. Recall that it had a single four letter string repeated twice, and the accompanying profile of behaviour showed increased RT at the fifth step consonant with chunking of the trial episode into groups of four subtrials. The activity in a number of multiple demand regions increased across the first four subtrials, but after the fourth subtrial showed a large decrease at the fifth subtrial and then increased again sequentially till the last subtrial, i.e. the pattern of activity in a number of regions was of a parallel increase between subtrials 1 to 4 and 5 to 8. Note that subjects had only four letters to remember and none of them could be forgotten across the first four subtrials since they were to be recalled again over the next four subtrials. Thus, clearly the working memory load stayed constant throughout the first four subtrials. Nonetheless, the activity in many multiple demand regions was seen to decrease across these four subtrials. The case of trial type 3 in chapter 4 further supports this notion. Recall that there was no working memory load across this trial

type, nonetheless some fronto-parietal regions showed parallel increase in activity across the subtrials 1 to 4 and 5 to 8, as was the case in trial type 2.

While the dynamics of working memory load cannot explain these, the dynamics of the representational load can. In both of these trial types the subtrials were chunked in groups of four. Hence, the first four and the second four subtrials were likely to be represented as two distinct task episodes. It is plausible that the cognitive structures necessary for the execution of the first four subtrials were constructed at the beginning of the trial, in which case the representational load associated with it can be expected to be greatest at the first subtrial and would decrease stepwise after the execution of each subtrial as the structure relevant for each subtrial is disassembled. This would free up the fronto-parietal resources involved, might manifest as an increased disinhibition of these regions. After the fourth subtrial, a new cognitive structure for the execution of the second task episode would be constructed, the increased load of which manifested in the decrease in fronto-parietal activity seen after the fourth but before the fifth subtrial. Subsequently, each subtrial executed again freed up the corresponding fronto-parietal resources, seen as increased disinhibition of these regions till the last subtrial.

Another aspect of the results that shows a difference between the representational load of episode structure and the working memory load is that increasing working memory load is known to *increase* the activity in multiple demand regions (Duncan, 2006), an effect that was replicated in our results as well. For example, in chapter 3, increasing the working memory load by making subjects maintain additional information pertaining to the second block, *increased* the activity during the first block. Similarly in chapter 4, cues requiring greater information to be maintained increased the activity in multiple demand regions. On the contrary, increased representational load *decreases* the activity in these regions.

Apart from the direction of neural modulation, the two kinds of loads also differ in the extent of brain regions affected. Working memory load in the current and earlier studies (Cabeza and Nyberg, 2000; Duncan, 2006) elicits activity that is mostly limited to the multiple demand regions of fronto-parietal cortices. However, the regions showing change in activity across the phases of sequential tasks and hence across the levels of representational load are much more extensive and included almost all major brain regions.

The above thesis suggests a distinction between two kinds of short term storage of information: those that are maintained actively in the working memory and increase the net activity in fronto-parietal regions, and others that are stored implicitly and decrease the activity in these regions. The information that creates the organisational structure of the task episode seems to belong to the latter category. Plausibly, task related information that cannot be maintained in the implicit representational structure is maintained in working memory by some other process like the phonological loop (Baddeley, 2007).

Some critical findings that were hitherto unexplainable seem to fit this dichotomy. Consider the case of the experiment described in chapter 4. At the cue stage, trial type 1 having the greatest working memory load elicited the greatest activity compared to other trial types. However, by the beginning of the subtrial sequence, the cue information has been incorporated into the structural framework of the trial episode, and the overall activity at the initial subtrials did not differ across the trial types differing in their working memory load. Similarly in chapter 3, the hard cues with greater working memory load elicited greater activity, but in the following interval, by the time of which the most of the cue information might have got incorporated into stable structure representations, the intervals preceding the hard task blocks did not have greater activity than those preceding easy blocks. On the

contrary, in some regions like IPS, the intervals preceding the easy blocks had greater activity (Fig 3.5).

6.4 Implications

Task Episodes

The results of this thesis show that mental organisation is an important determinant of activity in widespread regions including many within the fronto-parietal cortices, and that information about the structure of mental (or task) episode is represented widely across these regions.

How control arises in cognition, is a central question in the neurosciences. Earlier formulations have suggested that representations of task-relevant information could act as the source of control by biasing the relevant processing across various brain systems (Miller and Cohen, 2001; Duncan, 2006). Earlier studies in this regard have shown fronto-parietal representation of information related to task stimuli (Hampshire et al., 2007; Woolgar et al., 2011), rules (Bunge and Wallis, 2008), decisions (Haynes, 2007; Reverberi et al., 2011) and responses (Woolgar et al., 2011; Soon et al., 2008). The representation of the structure of the task episode in the context of which such information was represented was not looked into by earlier studies. It is possible that this representation, which has been shown by the current results to be present in widespread brain regions, is the source of control required for organising extended sequential behaviour. Perhaps disruption of such representations after frontal lesions is the cause of behavioural disorganisation seen in them (Duncan, 1986; Shallice and Burgess, 1991).

Interestingly, current results also show that the magnitude of representation of attended events (chapter 2) and cognitive set related to task events (chapter 5), are dependent on their position in the episode. The current results, thus, highlight an important issue, hitherto, ignored by many studies. Many previous studies have tried

to correlate a cognitive process with activity in the fronto-parietal regions e.g. conflict monitoring, (MacDonald et al., 2000), response inhibition (Aron, 2004), attentional shift (Tamber-Rosenau et al., 2011). However, as the results in chapter 2 illustrate in the context of attentional detection, such an attempt requires a further qualification – where does that cognitive process occur in the extended mental episode, and what role does it fulfil in the organisation of that episode?

It's also worth pointing out that an experimental trial is also an organised task episode, and so the issues highlighted in this thesis pertain to it as well. Notably, in many studies the kind of trials that elicit greater prefrontal activity, require a more extended and complicated organisation of the mental episode. Consider, for example, the two trials from Koechlin et al. (2003), (described in detail in chapter 1). In the low prefrontal activity eliciting trials, subjects decided about the response directly based on the colour of the stimuli; whereas in the high activity eliciting trials, they had to decide first about the episode, then about the currently relevant rule and, finally, about the colour-to-response mapping based on that rule. Clearly the number of sequential steps required to reach the solution is greater in the latter trials. While the authors interpreted the results as showing that greater frontal activity is elicited when the context of the episode has to be used to select for the relevant context or rule, the results in chapter 2 suggest that the same can be elicited in any extended, multistep task.

Conceptually similar issues can be speculated about the results of Dosenbach et al. (2006). They looked at the activity related to different aspects of task blocks in ten different experiments. However, the structure of the task blocks across these conditions was the same - extended blocks consisting of numerous individual trials with jittered inter-trial intervals. It is therefore possible that the qualitatively similar activity elicited by these blocks across the different experiments stemmed from their identical pattern of organisation.

Default Mode Regions

Another noteworthy result has been the presence of episode structure information in regions that have been thought of as task negative or *default* mode regions (Fox et al., 2005). The current results show that beyond representing episode structure information, these regions even show an effect of this information on their representation of task rule, since their content of task rule information increased across the sequential phases of the trial episode (chapter 5). Thus, these regions not only have information about the gross structure of the episode, but they also seem to represent the specific content of the individual steps of the episode. Further, the pattern of activity in these regions can be correlated with that in multiple demand regions. Together, these results call into question the assumption that these regions do not actively participate in tasks requiring cognitive control and attention to the external world.

Differentiation within MD Regions

The results discussed in the thesis also point at some functional distinctions amongst the different MD regions. It was noted in chapter 2 that the behaviour of posterior prefrontal regions like IFS was in distinction to the anterior prefrontal regions. While the anterior regions were only active at the completion of higher level episodes, posterior regions were active at even the lowest level episode completions, suggesting that it is only the anterior prefrontal regions that specifically represent hierarchically higher level episode. The behaviour of left IFS was specifically notable. This region was active in response to the completion of lowest level subtasks, but did not show additional activity at the completion of higher level task episodes. Similarly, in chapter 4, it was seen that in all trial types left IFS represented lower chunk level organisation, but not the higher sequence level organisation. In contrast the right APFC only represented the higher sequence level organisation but not the lower chunk level organisation. Right APFC and left IFS thus seem to be at the opposite

ends of the spectrum. The former only represents higher level aspects of task organisation, while the latter exclusively represents the lower level aspects of task hierarchy.

The results also pointed to a lateralisation in the representation of task episode structure. It was seen in chapter 2 that right prefrontal regions – right IFS and right APFC - showed greater activity in response to the completion of higher level task episodes than left prefrontal regions. Conceptually similar results were obtained in chapters 3 and 4, wherein the right prefrontal regions showed greater evidence of higher level task structure representation. Left prefrontal regions, on the other hand, showed greater modulation of activity with task difficulty (chapter 3). Taken together, these results suggest that the higher level structure of task episodes is preferentially represented in the right prefrontal regions, in contrast to lower level task structure which seems to have greater representation in the left prefrontal regions.

Primary Sensory and Motor Regions

One of the most puzzling results of this thesis has been the evidence that episode structure information is present in the primary sensory and motor regions of the brain. Some earlier studies have shown that task related information is present in the primary sensory region of the relevant modality e.g. primary visual cortex in visual tasks (Mirabella et al., 2007; Woolgar et al., 2011). However, such effects can be accounted for by the fact that the pattern of top down control exerted on these regions will be different when different rules are relevant. However, our results show that not only is the relevant rule coded in these regions but that this coding is also affected by the structure of task episode in which the individual rules are to be applied.

Finally, the demonstration of episode structure information in sensory regions of a different modality or in the task irrelevant primary motor cortex is, to my knowledge, novel. This and other issues raised by the current thesis will require further investigation.

Appendix

A. Tables of peak coordinates

Table A2.1. Chapter 2, Exp 1: T3 vs T2 & T1

BA	Label	Z	XYZ(mm)
46	R Mid Frontal	3.92	34,48,34
46	R Mid Frontal	3.53	26,50,18
46	L Mid Frontal	3.65	-40,28,38

10	L Mid Frontal	4.1	-26,54,10
6	R Sup Frontal	4.15	26,4,58
32	L Sup Med	4.84	-10,28,36
32	L Ant Cingulate	4.63	-10,28,30
32	R SMA	4.04	8,12,48
32	L SMA	3.81	2,14,46
32	R Mid Cingulate	3.98	6,14,46
32	L Mid Cingulate	4.69	-10,26,34
7	R Sup Parietal Lob	4.02	12,-68,54
7	L Sup Parietal Lob	3.47	-14,-68,44
40	R Inf Parietal Lob	3.85	34,-44,44
40	L Inf Parietal Lob	3.46	-60,-52,38
7	R Precuneus	4.1	10,-68,56
7	L Precuneus	3.52	-14,-64,38
7	R Sup Occipital	4.01	20,-60,40
7	R Cuneus	3.88	18,-60,40
18	L Calcarine	3.8	-4,-92,-12
	R Caudate Nucleus	3.75	8,22,6
	L Caudate Nucleus	4.02	-4,16,8

Table A2.2 . Chapter 2, Exp 1: X vs T2 & T1

BA	Label	Z	XYZ(mm)
8	R Sup Frontal	5.59	26,12,58
6	L Sup Frontal	3.8	-20,4,68
8	R Mid Frontal	5.41	28,10,54
46	R Mid Frontal	4.93	30,44,30
46	L Mid Frontal	4.77	-30,44,32
45	R Inf Frontal (p. Tri)	5.15	54,22,18
45	R Inf Frontal (p. Op)	5.13	56,22,14
6	L Inf Frontal (p. Op)	4.11	-56,8,12
6	R Precentral	4.89	44,6,52
47	R Insula Lobe	4.06	28,22,-10
47	L Insula Lobe	3.96	-28,26,-6
24	R Ant Cingulate	4.95	4,26,30
24	L Ant Cingulate	5.14	-2,26,32
8	R SMA	5.09	8,26,54
32	L SMA	4.32	2,16,46
8	R Sup Med	4.7	8,28,54
24	L Sup Med	5.02	-2,26,38
24	R Mid Cingulate	5.41	2,26,32
24	L Mid Cingulate	5.43	-6,18,36

40	R Inf Parietal Lob	4.51	40,-56,44
40	L Inf Parietal Lob	4.29	-54,-44,38
40	R Supramarginal	4.66	60,-42,32
23	R Post Cingulate	4.62	4,-34,30
23	L Post Cingulate	4.66	-4,-36,26
7	R Precuneus	4.88	2,-56,52
7	L Precuneus	5.25	-2,-58,54
18	R Lingual	4.7	12,-74,2
17	L Lingual	4.41	-8,-76,6
6	R Rol Operculum	4.55	54,8,10
6	L Rol Operculum	4.81	-46,-6,10
18	R Cuneus	5.23	16,-66,36
17	L Cuneus	4.42	-12,-70,32
17	R Calcarine	4.69	10,-74,4
17	L Calcarine	4.41	-8,-78,6

Table A3.1: Chapter 3, Hard block – Easy Block

BA	Label	Z	XYZ(mm)
8	R Sup Frontal	5.12	24,8,52
8	L Sup Frontal	4.99	-20,6,50
6	R Mid Frontal	4.57	32,4,40
8	R Mid Frontal	5.93	30,20,54
10	R Mid Frontal	5.27	36,62,10
46	R Mid Frontal	4.39	40,38,36

10	L Mid Frontal	5.56	-34,54,12
6	L Mid Frontal	4.72	-40,4,54
44	L Inf Frontal (p. Op)	4.82	-52,16,24
45	L Inf Frontal (p. Tri)	5.26	-46,32,24
10	R Mid Orbital	4.56	34,58,0
32	L Sup Med	4.18	-8,26,40
32	R SMA	4.4	8,22,46
8	L SMA	4.65	-10,20,58
7	R Sup Parietal Lob	5.23	12,-68,54
7	L Sup Parietal Lob	4.87	-22,-68,42
40	R Inf Parietal Lob	4.31	38,-50,48
40	L Inf Parietal Lob	5.05	-38,-46,50
7	L Inf Parietal Lob	4.32	-24,-66,44
7	R Precuneus	5.66	4,-60,48
7	L Precuneus	5.42	2,-62,46
19	R Mid Occipital	4.65	30,-72,34
	R Cerebelum (Crus 1)	5.75	32,-64,-28
	L Cerebelum (Crus 1)	5.09	-32,-64,-30
	Vermis (7)	5.02	0,-78,-26

Table A3.2. Chapter 3, Final block – Initial Block

BA	Label	Z	XYZ(mm)
46	R Mid Orbital	6.49	42,54,-6
47	L Insula Lobe	6.08	-30,24,-8
8	R SMA	5.77	6,24,54
45	R Inf Frontal (p. Op)	5.42	54,18,8
8	R Sup Med	5.3	6,32,44
8	R Mid Frontal	4.91	30,16,54
40	L Inf Parietal Lob	4.9	-54,-48,38
47	R Insula Lobe	4.88	46,22,-2
46	R Mid Frontal	4.87	44,52,2
47	L Mid Orbital	4.8	-36,44,-6
6	L Sup Frontal	4.73	-24,-4,70
40	R Supramarginal	4.72	52,-42,40
8	R Sup Frontal	4.69	12,26,54
44	R Precentral	4.49	48,12,40
3	L Postcentral	4.44	-40,-32,48
6	L Precentral	4.41	-32,-12,66
24	L Mid Cingulate	4.3	0,30,36
	L Cerebelum (VI)	4.26	-16,-72,-24
45	L Inf Frontal (p. Tri)	4.25	-56,20,8

37	R Mid Temporal	4.22	46,-54,8
3	R Postcentral	4.13	56,-22,44
46	L Mid Frontal	4.13	-44,48,2
2	L Sup Parietal Lob	4.09	-40,-42,58
2	R Sup Parietal Lob	4.05	44,-40,58
38	L Temporal Pole	3.96	-36,18,-28
38	R Temporal Pole	3.94	48,20,-14
6	R Rol Operculum	3.9	52,8,12
44	L Mid Frontal	3.89	-40,22,40
20	L Mid Temporal	3.64	-52,-34,-14
	R Cerebelum (Crus 2)	3.63	14,-76,-34
32	L Ant Cingulate	3.62	-2,36,32
37	L Mid Occipital	3.57	-48,-68,6
32	R Ant Cingulate	3.54	4,34,30
	R Hippocampus	3.48	38,-12,-14
22	L Sup Temporal	3.42	-62,-50,18
42	R Sup Temporal	3.4	64,-36,16
20	R Inf Temporal	3.31	58,-32,-18
	R Caudate Nucleus	3.3	12,10,6
7	R Cuneus	3.27	20,-66,34
18	R Sup Occipital	3.15	22,-66,32

Table A4.1. Chapter 4, Trial type 1: Subtrial 8 – intermediate subtrials

BA	Label	Z	XYZ(mm)
37	L Fusiform	6.05	-32,-44,-14
4	R Postcentral	5.84	18,-36,70
47	L Insula Lobe	5.6	-32,18,-12
	L Cerebelum (IV-V)	5.53	-8,-44,-6
18	L Cuneus	5.42	-10,-74,20
21	L Mid Temporal	5.36	-56,4,-14

18	R Lingual	5.28	10,-78,-2
23	R Mid Cingulate	5.25	6,-28,42
47	R Inf Frontal (p. Orb)	5.25	46,24,-2
47	R Inf Frontal (p. Tri)	5.24	50,24,0
47	R Insula Lobe	5.17	32,18,-10
23	L Mid Cingulate	5.16	-6,-36,46
11	R Mid Orbital	5.15	12,58,-8
20	R Mid Temporal	5.15	50,-16,-12
	R ParaHippocampal	5.15	18,-36,-8
3	L Postcentral	5.14	-42,-18,36
37	R Fusiform	5.11	26,-40,-12
45	R Inf Frontal (p. Op)	5.09	50,22,14
39	L Inf Parietal Lob	5.06	-56,-54,40
18	R Sup Occipital	5.02	12,-94,20
17	R Calcarine	5.01	8,-80,2
22	R Angular	4.99	62,-54,26
38	L Inf Frontal (p. Orb)	4.99	-36,18,-14
	L ParaHippocampal	4.99	-20,-40,-4
	R Amygdala	4.99	32,4,-20
5	R Precuneus	4.95	2,-40,48
10	R Sup Frontal	4.95	28,64,10
18	L Sup Occipital	4.95	-14,-90,18
11	L Sup Frontal	4.91	-22,58,2
8	R SMA	4.85	10,26,52
	L Hippocampus	4.85	-36,-16,-16
23	L Precuneus	4.83	-12,-38,46
	R Caudate Nucleus	4.82	8,10,2
9	R Mid Frontal	4.8	32,34,42
23	R Post Cingulate	4.77	8,-44,32
9	L Mid Frontal	4.75	-28,38,34
40	R Inf Parietal Lob	4.67	46,-54,44
20	L Inf Temporal	4.65	-40,-18,-18
32	L Ant Cingulate	4.62	-2,36,32
11	R Olfactory	4.57	20,8,-14
39	R Mid Occipital	4.55	50,-74,26
	R Cerebelum (IV-V)	4.55	22,-38,-18
	R Putamen	4.54	28,16,-6
3	R Precentral	4.53	36,-16,40
32	R Ant Cingulate	4.53	4,38,30
4	R Paracentral Lob	4.51	14,-40,50
4	L Paracentral Lob	4.48	-10,-32,52
6	L Precentral	4.48	-20,-24,70
44	R Inf Frontal (p. Tri)	4.47	44,24,32
	L Heschl	4.46	-38,-18,8

	R Hippocampus	4.45	20,-30,-6
23	L Post Cingulate	4.39	-8,-34,34
38	R Rol Operculum	4.38	58,16,0
5	R Sup Parietal Lob	4.28	14,-42,66

Table A4.2. Chapter 4, Trial type 2: Subtrial 4 – intermediate subtrials

BA label	AAL label	Z	XYZ(mm)
6	L SMA	4.6	-12,10,50
18	R Cuneus	4.31	8,-82,28
	Vermis (3)	4.3	0,-40,0
27	R Lingual	4.22	2,-38,0
47	R Insula Lobe	4.2	32,18,4
18	L Cuneus	4.16	-8,-84,20
	L Cerebelum (IV-V)	4.16	-4,-40,0
	R Putamen	4.12	30,18,4
47	L Mid Frontal	4.1	-28,48,4
11	L Mid Frontal	4.1	-26,48,2
7	R Sup Parietal Lob	3.97	16,-66,58
17	L Calcarine	3.94	-6,-90,0
6	R Sup Frontal	3.9	18,4,54
47	R Inf Frontal (p. Tri)	3.88	32,32,4
19	L Sup Occipital	3.86	-14,-86,18
7	L Sup Parietal Lob	3.85	-16,-70,52
3	R Postcentral	3.83	56,-24,46
47	L Insula Lobe	3.81	-36,26,4
6	R SMA	3.79	2,-4,74
32	L Mid Cingulate	3.78	-12,8,46
2	R Supramarginal	3.75	56,-26,46
18	L Mid Occipital	3.66	-18,-88,16
45	R Inf Frontal (p. Op)	3.64	46,16,2
22	R Mid Temporal	3.61	58,-46,12
27	R Post Cingulate	3.53	4,-40,6
21	L Mid Temporal	3.53	-58,-54,6
45	L Inf Frontal (p. Op)	3.49	-46,12,2
40	L Inf Parietal Lob	3.48	-28,-46,38
19	R Mid Occipital	3.43	36,-72,30
32	L Sup Med	3.42	-6,18,44
27	R ParaHippocampal	3.42	16,-40,-4
8	R Mid Frontal	3.41	26,12,50
6	L Precentral	3.39	-34,2,56
46	R Mid Frontal	3.33	24,42,24
6	R Precentral	3.31	42,4,36
41	R Angular	3.28	48,-46,26

32	R Mid Cingulate	3.21	10,18,38
29	L Post Cingulate	3.21	-6,-42,8
7	R Sup Occipital	3.21	32,-72,42

Table A4.3. Chapter 4, Trial type 2: Subtrial 8 – intermediate subtrials

BA	Label	Z	XYZ(mm)
18	L Calcarine	5.63	-2,-92,14
17	R Cuneus	5.5	12,-96,16
18	R Calcarine	5.45	8,-90,10
17	R Linual	5.07	6,-82,0
18	L Sup Occipital	4.86	-12,-86,22
19	L Linual	4.61	-26,-54,-4
47	R Inf Frontal (p. Orb)	4.6	38,26,-8
47	R Insula Lobe	4.57	30,18,-8
7	L Precuneus	4.49	-10,-74,38
45	R Inf Frontal (p. Op)	4.46	52,20,12
40	R Angular	4.43	60,-48,32
3	L Postcentral	4.43	-46,-18,34
2	R Postcentral	4.41	28,-40,62
21	R Mid Temporal	4.41	62,-26,-4
19	L Fusiform	4.4	-28,-56,-6
7	R Precuneus	4.38	6,-80,46
27	R ParaHippocampal	4.38	18,-40,-6
18	R Mid Occipital	4.38	26,-90,18
18	L Cerebelum (VI)	4.3	-20,-70,-14
18	Vermis (6)	4.26	2,-70,-4
32	R Ant Cingulate	4.26	4,40,30
41	R Sup Temporal	4.23	50,-42,24
5	R Sup Parietal Lob	4.17	14,-46,64
18	L Mid Occipital	4.14	-18,-88,16
19	R Fusiform	4.14	28,-72,-10
44	R Mid Frontal	4.14	40,20,40
23	R Post Cingulate	4.13	4,-38,32
6	R Precentral	4.13	34,-16,46
3	R Postcentral	4.09	44,-18,34
8	R SMA	4.07	6,26,56
	L Cerebelum (Crus 2)	4.01	-22,-78,-34
40	R Inf Parietal Lob	3.98	60,-48,38
	R Putamen	3.93	28,16,-6
	L Cerebelum (Crus 1)	3.91	-16,-84,-26
27	L ParaHippocampal	3.86	-20,-40,-6
32	L Ant Cingulate	3.86	-2,36,32
18	R Cerebelum (VI)	3.85	18,-64,-12

	R Putamen	3.8	34,-8,4
41	R Mid Temporal	3.77	46,-46,20
23	L Post Cingulate	3.73	-4,-34,34

Table A4.4. Chapter 4, Trial type 3: Subtrial 4 – intermediate subtrials

BA	Label	Z	XYZ(mm)
45	L Inf Frontal (p. Tri)	4.53	-36,22,10
18	R Sup Occipital	4.34	20,-92,20
24	L Ant Cingulate	4.29	-10,18,26
19	R Mid Occipital	4.21	42,-86,2
18	L Lingual	4.05	-10,-60,4
17	L Calcarine	4.04	-10,-60,6
19	R Inf Occipital	3.91	48,-74,-4
6	L Sup Frontal	3.88	-30,-6,64
6	R Mid Frontal	3.74	42,-2,60
38	L Temporal Pole	3.72	-52,6,2
47	R Inf Frontal (p. Tri)	3.68	38,26,8
44	L Mid Frontal	3.67	-48,28,34
47	R Insula Lobe	3.63	38,24,4
32	L Sup Med	3.63	-12,32,32
40	L Inf Parietal Lob	3.62	-38,-36,48
37	L Mid Occipital	3.6	-42,-70,0
24	L Mid Cingulate	3.53	-8,8,32
19	L Sup Occipital	3.4	-24,-70,30
27	R Lingual	3.36	8,-32,-6
	L Cerebellum (IV-V)	3.26	-8,-56,-2
24	R Mid Cingulate	3.25	6,14,34
7	R Angular	3.24	28,-52,44
40	R Inf Parietal Lob	3.23	30,-48,42

Table A4.5. Chapter 4, Trial type 3: Subtrial 8 – intermediate subtrials

BA	Label	Z	XYZ(mm)
23	R Precuneus	4.98	12,-50,36
30	L Post Cingulate	4.97	-6,-48,22
23	R Mid Cingulate	4.89	10,-48,34
23	L Precuneus	4.83	-6,-54,20
39	R Angular	4.71	40,-54,28
22	R Sup Temporal	4.68	58,-54,24

	L Cerebelum (Crus 2)	4.61	-16,-80,-38
39	R Mid Temporal	4.6	50,-60,20
23	L Cuneus	4.55	-8,-62,28
30	R Calcarine	4.45	14,-48,12
39	R Angular	4.36	38,-56,28
11	R Mid Orbital	4.32	10,58,-6
10	L Mid Orbital	4.32	-8,62,-4
37	L Fusiform	4.22	-26,-42,-12
37	L Lingual	4.03	-24,-44,-8
23	L Mid Cingulate	3.97	-8,-40,34
	L Cerebelum (VII)	3.95	-20,-74,-40
39	R Mid Occipital	3.88	44,-64,30

Table A5.1. Chapter 5, Regions where discriminability of A vs B increased across trial type 1

BA	Label	Z	XYZ(mm)
47	R Insula Lobe	4.16	27,26,-11
11	R Sup Orbital	3.8	21,29,-14
47	R Putamen	3.64	24,23,-5
47	R Inf Frontal (p. Tri)	3.49	39,38,1
11	R Mid Orbital	3.35	27,38,-11
11	R Rectal	3.12	18,20,-14
11	R Caudate Nucleus	3.11	18,26,-5
37	R Fusiform	3.81	27,-52,-17
37	R Linual	3.13	24,-49,-8
5	R Precuneus	3.93	9,-40,58
2	R Paracentral Lob	3.7	15,-40,55
5	L Precuneus	3.59	-12,-43,49
5	R Sup Parietal Lob	3.54	15,-43,61
	R Insula Lobe	3.68	36,-13,22
41	R Sup Temporal	3.63	42,-28,16
3	R Postcentral	3.37	39,-19,37
	R Putamen	3.29	30,-10,13
18	R Cuneus	3.7	15,-79,31
19	R Sup Occipital	3.69	27,-82,40
23	R Precuneus	3.27	12,-67,28
19	R Mid Occipital	3.22	36,-88,28
8	R Sup Med	3.69	6,26,46
32	R Ant Cingulate	3.66	12,35,19
32	R Mid Cingulate	3.58	6,29,40
8	R SMA	3.57	9,26,49
32	L Sup Med	3.32	0,29,40
9	R Sup Frontal	3.23	12,38,52

Table A5.2. Chapter 5, Regions where discriminability of A vs B increased across trial type 2

BA	Label	Z	XYZ(mm)
19	R Cuneus	4.62	12,-82,40
17	R Calcarine	4.56	3,-64,16
23	R Precuneus	4.07	6,-67,25
19	R Sup Occipital	3.81	15,-85,37
23	L Precuneus	3.81	0,-61,19
18	L Cuneus	3.79	0,-70,22
	R Cerebelum (IV-V)	3.21	15,-55,-11
19	R Fusiform	3.12	27,-58,-2
22	R Sup Temporal	4.18	66,-31,7
22	R Mid Temporal	4.12	54,-49,16
22	R Sup Temporal	3.15	66,-49,25
32	L Sup Med	3.82	-9,47,31
32	R Ant Cingulate	3.77	9,47,10
32	L Ant Cingulate	3.58	0,47,13
10	L Mid Orbital	3.3	-3,53,-5
10	R Mid Orbital	3.22	0,53,-2

B. Task Instructions

B.2 Chapter 2

At the beginning of each trial, you will see a cue three-letter word, followed by a series of letters presented in the middle of the computer screen. **Two tasks** are to be done on each trial: **First Task** is to detect the appearance of the letters of the word shown, in the correct order, amongst the other letters. **The second task** is to then search for the letter 'X' amongst the remaining letters. You only need to respond at the end of the trial when the cue "X?" appears: If you have seen the letters of the cue word and the letter 'X' in the correct order, then press the first button; if the two-word sequence was not completed in the correct order, then press the second button. You have two seconds to make this decision.

Examples:

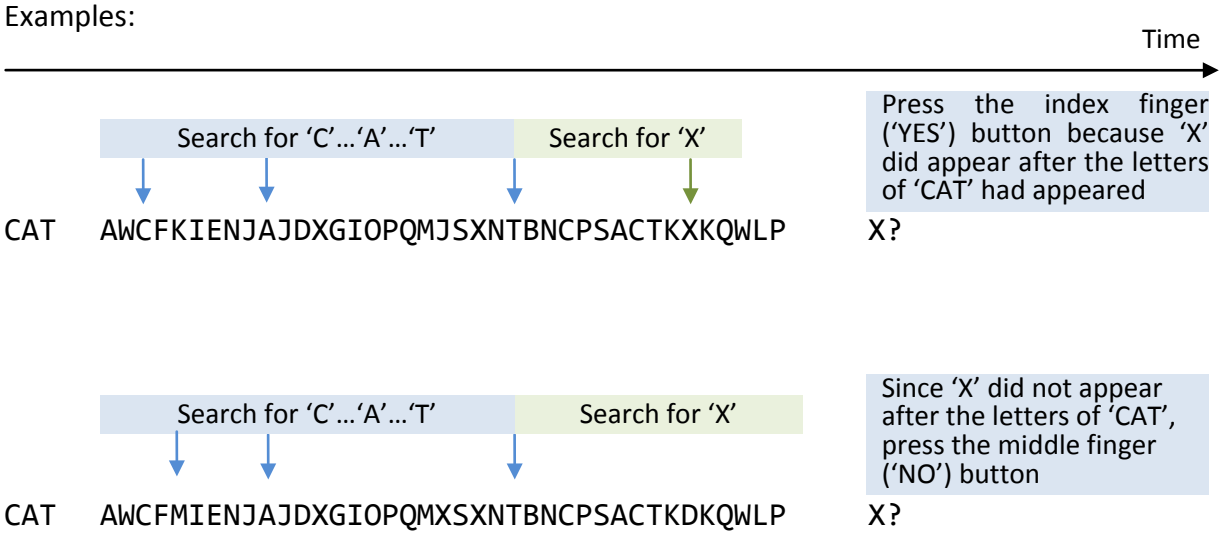

(Note that 'X' did occur earlier in the sequence, but was irrelevant. Only the appearance of 'X' after the appearance of the letters of the cued word is to be considered relevant)

Correctly detecting 'X' will increment your score by one point. Any error will decrease your score by one point. Your total score will be presented briefly, followed by a short pause before the next trial starts automatically.

Please try to remain still and look at the centre of the screen at all times.

This sequence will be repeated throughout three blocks of about 20 minutes each, with short breaks in between.

B.3 Chapter 3

At the beginning of each trial you will be presented with some arithmetic information (consisting of either a number e.g. '3' or two numbers along with a '+' or '-', in the format 'x+ y' e.g. '6+ 3') and one or two words. These are the cues that you need to remember. In rest of the trial, you will do an arithmetic task block using the arithmetic information, and a word task block using the cued word(s). The two kinds of cued information will be displayed in different lines. The information related to the task block that will be done first will always be displayed on top.

Consider the two cue displays below:

In case of the left display, word task will be done first, followed by the arithmetic task. Reverse will be the case for cue display on the right.

The cue screen will then disappear, and will be followed by an empty box whose margin colour will depend upon the nature of the 1st task block: Blue – word task; Green- arithmetic task.

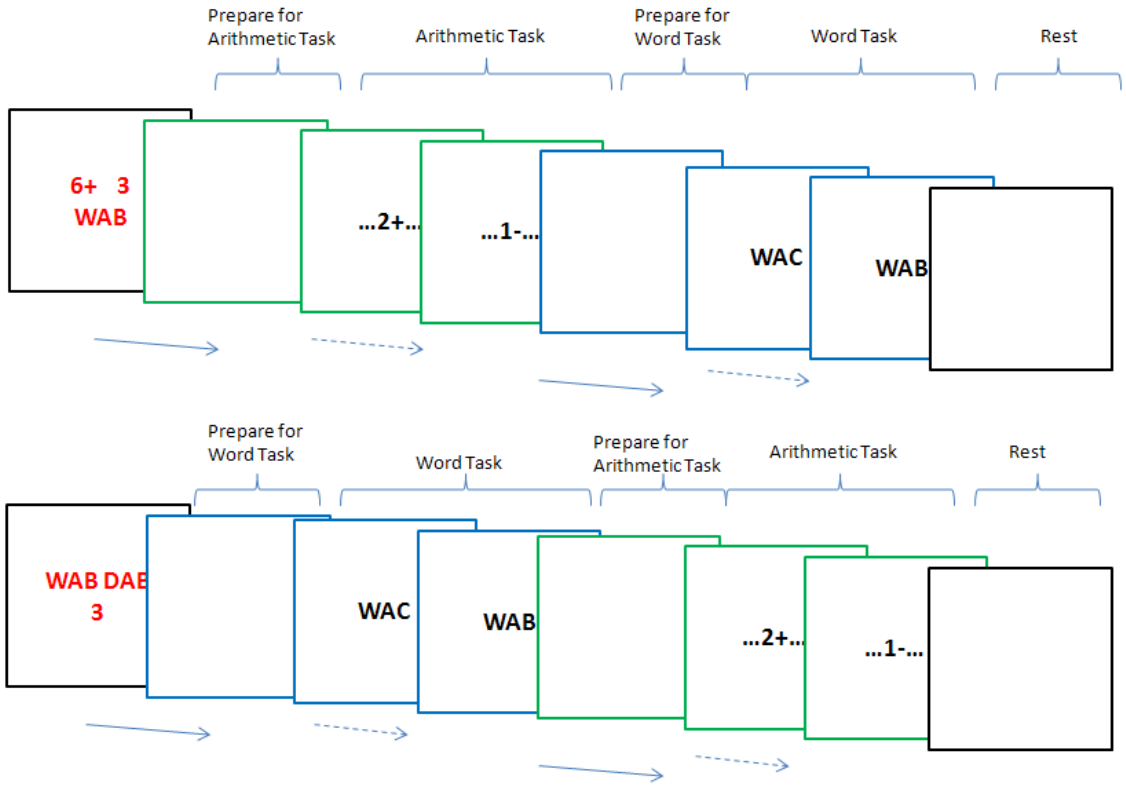

Subsequently, the first task block will start. This will involve a series of stimulus presentations on each of which you will execute the relevant task (details below).

The first task block will end with a blank interval, at which point the colour of the box margins will change, and will now signal the nature of the ensuing task block. Note that the two task blocks in a trial are always different i.e. the same task block is never repeated within a trial, thus if the first block was arithmetic, the second will be word, and vice-versa.

You will then execute the second task block, at the end of which, the colour of the margins will become black. You can now relax till the next trial begins.

Arithmetic Task:

In arithmetic blocks, the presented stimulus consists of two blanks with a number followed by '+' or '-' in between (e.g. '...4+...'). If the arithmetic cue was a single number (e.g. '3') then mentally put this number in the second blank and then solve the resulting expression (e.g. '4+3'). If the solution is less than 10, press the index finger button, else press the middle finger button.

In trials where arithmetic cue was in the pattern 'x+ y' e.g. '6+ 3', mentally put the first part of this ('6+') in the first blank, and the second part ('3') in the second blank, and then solve the resulting expression ('6+4+3). Again, if the solution is less than 10, press the index finger button, else press the middle finger button.

Word Task:

In word blocks, you will see a number of word presentations, at each respond with index finger button press if the presented word is the same as the word(s) cued at the beginning of the trial.

B.4 Chapter 4

Pre-Scan Session Instructions

We will start with a practise session.

Before each trial of this session, you will be presented with a four letter string consisting of As and Bs. This represents the sequence of four sequential tasks that you will execute on the ensuing four sequential stimulus presentations e.g. if the string is 'ABAB', you execute task A then task B then task A then task B. It is hence important that you keep this presented string in your mind for the rest of the trial. The trial will end after the fourth stimulus presentation.

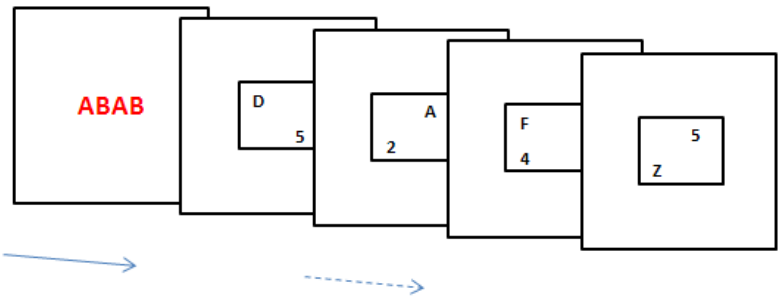

The stimuli consist of a letter and a number presented within a box. Task A requires you to categorise the letter as vowel or consonant. Press the index finger button if the letter is a vowel, else press the middle finger button. Task B is to categorise the number as even or odd; press the ring finger button if the number is even, else press the little finger button.

Now, you will do a 20 minute practise session.

Scan Session Instruction

In the main session, you will do a very similar experiment. You will now see an eight letter string at the beginning of each trial. The trial will now consist of eight sequential stimulus presentations instead of four, at each of which you will execute the task corresponding to the letters of the string.

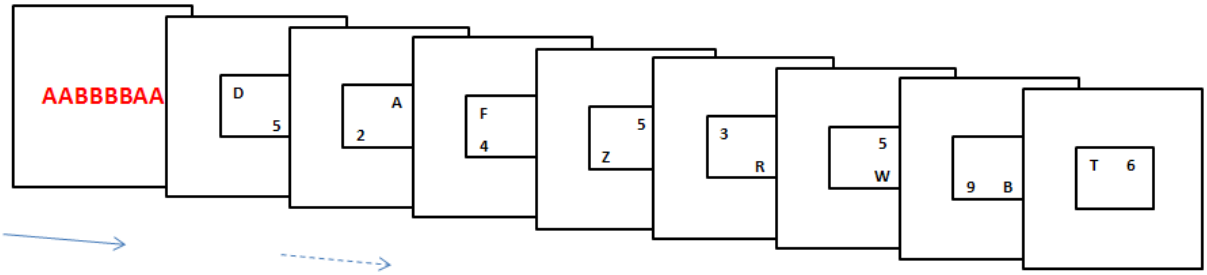

In some trials, you will not be presented with the initial eight letter string; instead you will see a blank. In such trials, the colour of the margins of the enclosing box will tell you the task to be done at each step - blue: letter task; green: number task).

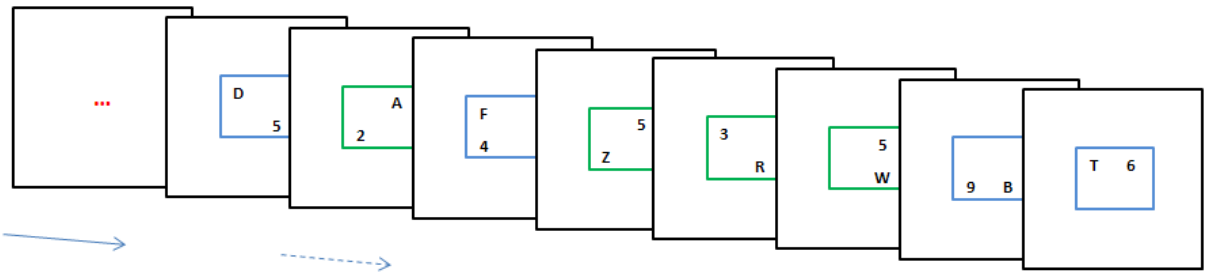

After the response to the eighth stimulus presentation, the trial will end. You can now relax till the beginning of the next trial.

References

Allport, A., & Wylie, G. (2000). Task- switching, stimulus-response bindings, and negative priming. In S. Monsell & J. Driver (Eds.), *Control of cognitive processes: attention and performance XVIII*. Cambridge, MA: MIT Press.

Alvarez, J. A., & Emory, E. (2006). Executive Function and the Frontal Lobes: A Meta-Analytic Review. *Neuropsychology Review*, *16*(1), 17-42. doi:10.1007/s11065-006-9002-x

Anderson, J. R., & Matessa, M. (1997). A Production System Theory of Serial Memory. *Psychological Review*, *104*(4), 728-748.

Anderson, J. R., Anderson, J. F., Ferris, J. L., Fincham, J. M., & Jung, K.-J. (2009). Lateral inferior prefrontal cortex and anterior cingulate cortex are engaged at different stages in the solution of insight problems. *Proceedings of the National Academy of Sciences*, *106*(26), 10799- 10804.

Anderson, J. R., Bothell, D., Lebiere, C., & Matessa, M. (1998). An Integrated Theory of List Memory, *Journal of Memory and Language*, *38*(4), 341-380.

Anderson, S. W., Damasio, H., Jones, R. D., & Tranel, D. (1991). Wisconsin Card Sorting Test performance as a measure of frontal lobe damage. *Journal of Clinical and Experimental Neuropsychology*, *13*(6), 909-922.

Aron, A. R., Robbins, T. W., & Poldrack, R. A. (2004). Inhibition and the right inferior frontal cortex. *Trends in Cognitive Sciences*, *8*(4), 170-177.

Asaad, W. F., Rainer, Gregor, & Miller, Earl K. (1998). Neural Activity in the Primate Prefrontal Cortex during Associative Learning. *Neuron*, *21*(6), 1399-1407.

Asaad, W. F., Rainer, Gregor, & Miller, Earl K. (2000). Task-Specific Neural Activity in the Primate Prefrontal Cortex. *Journal of Neurophysiology*, *84*(1), 451 -459.

Azuma, M., & Suzuki, H. (1984). Properties and distribution of auditory neurons in the dorsolateral prefrontal cortex of the alert monkey. *Brain Research*, *298*(2), 343-346.

Baddeley, A. (2007). *Working memory*. Oxford, England: Oxford University Press.

Badre, D., & D'Esposito, M. (2007). Functional Magnetic Resonance Imaging Evidence for a Hierarchical Organization of the Prefrontal Cortex. *Journal of Cognitive Neuroscience*, *19*(12), 2082-2099.

Badre, D., & D'Esposito, M. (2009). Is the rostro-caudal axis of the frontal lobe hierarchical? *Nat Rev Neurosci*, *10*(9), 659-669.

Badre, D., Kayser, A. S., & D'Esposito, M. (2010). Frontal cortex and the discovery of abstract action rules. *Neuron*, *66*(2), 315-326.

Bahlmann, J., Schubotz, R. I., Mueller, J. L., Koester, D., & Friederici, A. D. (2009). Neural circuits of hierarchical visuo-spatial sequence processing. *Brain Research*, *1298*, 161-170.

Berdyyeva, T. K., & Olson, C. R. (2009). Monkey Supplementary Eye Field Neurons Signal the Ordinal Position of Both Actions and Objects. *The Journal of Neuroscience*, *29*(3), 591 - 599.

Bianchi, L., & Macdonald, J. H. (1992). *The Mechanism of the Brain: And the Function of the Frontal Lobes*. E. & S. Livingstone.

Bode, S., & Haynes, J.-D. (2009). Decoding sequential stages of task preparation in the human brain. *NeuroImage*, *45*(2), 606-613.

Bode, S., He, A. H., Soon, C.S., Trampel, R., Turner, R., & Haynes, J.-D. (2011). Tracking the Unconscious Generation of Free Decisions Using Uitra-High Field fMRI. *PLoS ONE*, *6*(6).

Bode, Stefan, & Haynes, John-Dylan. (2009). Decoding sequential stages of task preparation in the human brain. *NeuroImage*, *45*(2), 606-613.

Botvinick, M. M. (2008). Hierarchical models of behavior and prefrontal function. *Trends in Cognitive Sciences*, *12*(5), 201-208.

Botvinick, M., & Plaut, D. C. (2004). Doing without schema hierarchies: a recurrent connectionist approach to normal and impaired routine sequential action. *Psychological Review*, *111*(2), 395-429.

Bowman, E. M., Aigner, T. G., & Richmond, B. J. (1996). Neural signals in the monkey ventral striatum related to motivation for juice and cocaine rewards. *Journal of Neurophysiology*, *75*(3), 1061 -1073.

Boy, F., Husain, M., & Sumner, P. (2010). Unconscious inhibition separates two forms of cognitive control. *Proceedings of the National Academy of Sciences*, *107*(24), 11134 -11139.

Braver, T. S., Reynolds, J. R., & Donaldson, D. I. (2003). Neural Mechanisms of Transient and Sustained Cognitive Control during Task Switching. *Neuron*, *39*(4), 713-726.

Braver, T. S., Paxton, J. L., Locke, H. S., & Barch, D. M. (2009). Flexible neural mechanisms of cognitive control within human prefrontal cortex. *Proceedings of the National Academy of Sciences*, *106*(18), 7351-7356.

Brett, M., Johnsrude, I. S., & Owen, A. M. (2002). The problem of functional localization in the human brain. *Nat Rev Neurosci*, *3*(3), 243-249.

Buckner, Randy L, Andrews-Hanna, J. R., & Schacter, D. L. (2008). The brain's default network: anatomy, function, and relevance to disease. *Annals of the New York Academy of Sciences*, *1124*, 1-38.

Bunge, S. A., Helskog, E. H., & Wendelken, C. (2009). Left, but not right, rostrolateral prefrontal cortex meets a stringent test of the relational integration hypothesis. *NeuroImage*, *46*(1), 338-342.

Burgess, P. W., Dumontheil, I., & Gilbert, Sam J. (2007). The gateway hypothesis of rostral prefrontal cortex (area 10) function. *Trends in Cognitive Sciences*, *11*(7), 290-298.

Buschman, T. J., & Miller, E. K. (2007). Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science*, *315*(5820), 1860.

Cabeza, R., & Nyberg, L. (2000). Imaging cognition II: An empirical review of 275 PET and fMRI studies. *Journal of cognitive neuroscience*, *12*(1), 1–47.

Campos, M., Breznen, B., & Andersen, R. A. (2010). A Neural Representation of Sequential States Within an Instructed Task. *Journal of Neurophysiology*, *104*(5), 2831 -2849.

Chikazoe, J., Jimura, K., Asari, T., Yamashita, K.-ichiro, Morimoto, H., Hirose, S., Miyashita, Y., et al. (2009). Functional Dissociation in Right Inferior Frontal Cortex during Performance of Go/No-Go Task. *Cereb. Cortex*, *19*(1), 146-152.

Christoff, K., Keramatian, K., Gordon, A. M., Smith, R., & Mädler, B. (2009). Prefrontal organization of cognitive control according to levels of abstraction. *Brain Research*, *1286*, 94-105.

Clark, V. P., Fannon, S., Lai, S., Benson, R., & Bauer, L. (2000). Responses to Rare Visual Target and Distractor Stimuli Using Event-Related fMRI. *Journal of Neurophysiology*, *83*(5), 3133 - 3139.

Clegg, B. A., Digirolamo, G. J., & Keele, S. W. (1998). Sequence learning. *Trends in Cognitive Sciences*, *2*(8), 275-281.

Cole, M. W., & Schneider, W. (2007). The cognitive control network: Integrated cortical regions with dissociable functions. *NeuroImage*, *37*(1), 343–360.

Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews. Neuroscience*, *3*(3), 201-215.

Corbetta, M., Patel, G., & Shulman, G. L. (2008). The Reorienting System of the Human Brain: From Environment to Theory of Mind. *Neuron*, *58*(3), 306-324.

Cusack, R., & Papadakis, N. (2002). New Robust 3-D Phase Unwrapping Algorithms: Application to Magnetic Field Mapping and Undistorting Echoplanar Images. *NeuroImage*, *16*(3, Part 1), 754-764.

De Renzi, E., & Barbieri, C. (1992). The incidence of the grasp reflex following hemispheric lesion and its relation to frontal damage. *Brain: A Journal of Neurology*, *115 Pt 1*, 293-313.

Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in Cognitive Sciences*, *10*(5), 204-211.

Dehaene, S., & Changeux, J.-P., Experimental and theoretical approaches to conscious processing. *Neuron*, 70 (2011), 200–227

Desimone, Robert, & Duncan, John. (1995). Neural Mechanisms of Selective Visual Attention. *Annual Review of Neuroscience*, *18*(1), 193-222.

Desmurget, M., Reilly, K. T., Richard, N., Szathmari, A., Mottolese, C., & Sirigu, Angela. (2009). Movement Intention After Parietal Cortex Stimulation in Humans. *Science*, *324*(5928), 811 -813.

Dosenbach, N. U. F., Fair, D. A., Miezin, F. M., Cohen, A. L., Wenger, K. K., Dosenbach, R. A. T., Fox, M. D., et al. (2007). Distinct brain networks for adaptive and stable task control in humans. *Proceedings of the National Academy of Sciences*, *104*(26), 11073-11078.

Dosenbach, Nico U.F., Visscher, K. M., Palmer, E. D., Miezin, Francis M., Wenger, Kristin K., Kang, H. C., Burgund, E. D., et al. (2006). A Core System for the Implementation of Task Sets. *Neuron*, *50*(5), 799-812.

Dove, A., Pollmann, S., Schubert, T., Wiggins, C. J., & von Cramon, D. Y. (2000). Prefrontal cortex activation in task switching: an event-related fMRI study. *Cognitive Brain Research*, *9*(1), 103-109.

Downar, J., Crawley, A. P., Mikulis, D. J., & Davis, K. D. (2001). The Effect of Task Relevance on the Cortical Response to Changes in Visual and Auditory Stimuli: An Event-Related fMRI Study. *NeuroImage*, *14*(6), 1256-1267.

Downar, J., Crawley, A. P., Mikulis, D. J., & Davis, K. D. (2002). A Cortical Network Sensitive to Stimulus Salience in a Neutral Behavioral Context Across Multiple Sensory Modalities. *Journal of Neurophysiology*, *87*(1), 615 -620.

Driver, J., & Mattingley, J. B. (1998). Parietal neglect and visual awareness. *Nat Neurosci*, *1*(1), 17-22.

Dumontheil, I., Thompson, Russell, & Duncan, John. (2010). Assembly and Use of New Task Rules in Fronto-parietal Cortex. *Journal of Cognitive Neuroscience*.

Duncan, John. (1986). Disorganisation of behaviour after frontal lobe damage. *Cognitive Neuropsychology*, *3*(3), 271-290.

Duncan, John. (2001). An adaptive coding model of neural function in prefrontal cortex. *Nat Rev Neurosci*, *2*(11), 820-829.

Duncan, John. (2006). EPS - Mid-Career Award 2004: Brain mechanisms of attention. *The Quarterly Journal of Experimental Psychology*, *59*(1), 2.

Duncan, John. (2010). The multiple-demand (MD) system of the primate brain: mental programs for intelligent behaviour. *Trends in Cognitive Sciences*, *14*(4), 172-179.

Duncan, John, Bundesen, C., Olson, A., Humphreys, G., Chavda, S., & Shibuya, H. (1999). Systematic analysis of deficits in visual attention. *Journal of Experimental Psychology: General*, *128*(4), 450-478.

Duncan, John, Burgess, P., & Emslie, H. (1995). Fluid intelligence after frontal lobe lesions. *Neuropsychologia*, *33*(3), 261-268.

Duncan, John, Humphreys, G., & Ward, R. (1997). Competitive brain activity in visual attention. *Current Opinion in Neurobiology*, *7*(2), 255-261.

Duncan, John, Parr, A., Woolgar, A., Thompson, Russell, Bright, P., Cox, S., Bishop, S., et al. (2008). Goal neglect and Spearman's g: competing parts of a complex task. *Journal of Experimental Psychology. General*, *137*(1), 131-148. :

Dux, P. E., Tombu, M. N., Harrison, S., Rogers, B. P., Tong, F., & Marois, R. (2009). Training Improves Multitasking Performance by Increasing the Speed of Information Processing in Human Prefrontal Cortex. *Neuron*, *63*(1), 127-138.

Eger, E., Michel, V., Thirion, B., Amadon, A., Dehaene, S., & Kleinschmidt, A. (2009). Deciphering Cortical Number Coding from Human Brain Activity Patterns. *Current Biology*, *19*(19), 1608-1615.

Esterman, M., Chiu, Y.-C., Tamber-Rosenau, B. J., & Yantis, S. (2009). Decoding cognitive control in human parietal cortex. *Proceedings of the National Academy of Sciences*, *106*(42), 17974 -17979.

Farag, C., Troiani, V., Bonner, M., Powers, C., Avants, B., Gee, J., & Grossman, M. (2010). Hierarchical Organization of Scripts: Converging Evidence from fMRI and Frontotemporal Degeneration. *Cerebral Cortex*, *20*(10), 2453 -2463.

Forstmann, B. U., Brass, Marcel, Koch, I., & Cramon, D. Y. von. (2005). Internally generated and directly cued task sets: an investigation with fMRI. *Neuropsychologia*, *43*(6), 943-952.

Fox, Michael D., Snyder, Abraham Z., Vincent, Justin L., Corbetta, M., Van Essen, D. C., & Raichle, Marcus E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences of the United States of America*, *102*(27), 9673 -9678.

Freedman, D. J., Riesenhuber, M., Poggio, T., & Miller, E K. (2001). Categorical representation of visual stimuli in the primate prefrontal cortex. *Science (New York, N.Y.)*, *291*(5502), 312-316.

Fuster, Joaquín M. (2008). *The prefrontal cortex*. Elsevier.

van Gaal, S., Ridderinkhof, K. R., Fahrenfort, J. J., Scholte, H. S., & Lamme, V. A. F. (2008). Frontal Cortex Mediates Unconsciously Triggered Inhibitory Control. *The Journal of Neuroscience*, *28*(32), 8053 -8062.

Gat, E. (1998). On Three-Layer Architectures. *Artificial Intelligence and Mobile Robots*. MIT Press.

Georgopoulos, A., Schwartz, A., & Kettner, R. (1986). Neuronal population coding of movement direction. *Science*, *233*(4771), 1416 -1419.

Gilbert, S.J. (2011). Decoding the content of delayed intentions. *Journal of Neuroscience*, *31*(8), 2888-2894.

Grafman, J, Salazar, A., Weingartner, H., Vance, S., & Amin, D. (1986). The relationship of brain-tissue loss volume and lesion location to cognitive deficit. *The Journal of Neuroscience*, *6*(2), 301 -307.

Grant, D. A., & Berg, E. (1948). A behavioral analysis of degree of reinforcement and ease of shifting to new responses in a Weigl-type card-sorting problem. *Journal of Experimental Psychology*, *38*(4), 404-411.

Greenberg, A. S., Esterman, M., Wilson, D., Serences, J. T., & Yantis, S. (2010). Control of Spatial and Feature-Based Attention in Frontoparietal Cortex. *The Journal of Neuroscience*, *30*(43), 14330 -14339.

Gregoriou, G. G., Gotts, S. J., Zhou, H., & Desimone, R. (2009). High-frequency, long-range coupling between prefrontal and visual cortex during attention. *Science*, *324*(5931), 1207.

Greicius, M. D., & Menon, V. (2011). Default-Mode Activity during a Passive Sensory Task: Uncoupled from Deactivation but Impacting Activation. *Journal of Cognitive Neuroscience*, *16*(9), 1484-1492.

Gruber, O., Karch, S., Schlueter, E. K., Falkai, P., & Goschke, T. (2006). Neural mechanisms of advance preparation in task switching. *NeuroImage*, *31*(2), 887-895.

Hampshire, A., Duncan, John, & Owen, A. M. (2007). Selective Tuning of the Blood Oxygenation Level-Dependent Response during Simple Target Detection Dissociates Human Frontoparietal Subregions. *J. Neurosci.*, *27*(23), 6219-6223.

Harrison, S. A., & Tong, F. (2009). Decoding reveals the contents of visual working memory in early visual areas. *Nature*, *458*(7238), 632-635.

Hasegawa, R. P., Blitz, A. M., & Goldberg, M. E. (2004). Neurons in Monkey Prefrontal Cortex Whose Activity Tracks the Progress of a Three-Step Self-Ordered Task. *Journal of Neurophysiology*, *92*(3), 1524 -1535.

Haynes, John-Dylan, Sakai, K., Rees, Geraint, Gilbert, S., Frith, C., & Passingham, R. E. (2007). Reading Hidden Intentions in the Human Brain. *Current Biology*, *17*(4), 323-328.

Henry, H. M., & Rogers, D. E., Increased response latency for complicated movements and a "memory drum" theory of neuromotor reaction, *Research Quarterly* **31** (1960), pp. 448–458.

Hon, N., Epstein, R. A., Owen, A. M., & Duncan, John. (2006). Frontoparietal Activity with Minimal Decision and Control. *J. Neurosci.*, *26*(38), 9805-9809.

Hoshi, E., Shima, K., & Tanji, J. (1998). Task-Dependent Selectivity of Movement-Related Neuronal Activity in the Primate Prefrontal Cortex. *Journal of Neurophysiology*, *80*(6), 3392 -3397.

Husain, M., & Nachev, P. (2007). Space and the parietal cortex. *Trends in Cognitive Sciences*, *11*(1), 30-36.

Husain, M., & Rorden, C. (2003). Non-spatially lateralized mechanisms in hemispatial neglect. *Nature Reviews. Neuroscience*, *4*(1), 26-36.

Husain, M., Mannan, S., Hodgson, T., Wojciulik, E., Driver, J., & Kennard, C. (2001). Impaired spatial working memory across saccades contributes to abnormal search in parietal neglect. *Brain*, *124*(5), 941 -952.

Jenkins, L. J., & Ranganath, C. (2010). Prefrontal and Medial Temporal Lobe Activity at Encoding Predicts Temporal Context Memory. *The Journal of Neuroscience*, *30*(46), 15558 -15565.

Jiang, Y, & Kanwisher, N. (2003a). Common neural mechanisms for response selection and perceptual processing. Journal of Cognitive Neuroscience , 15: 1095–1110

Jiang, Y, & Kanwisher, N. (2003b). Common neural substrates for response selection across modalities and mapping paradigms. Journal of Cognitive Neuroscience , 15: 1080–1094.

De Jong, R. (1995). The role of preparation in overlapping-task performance. *The Quarterly Journal of Experimental Psychology*, *48*(1), 2-25.

Kesner, R. P., Hopkins, R. O., & Fineman, B. (1994). Item and order dissociation in humans with prefrontal cortex damage. *Neuropsychologia*, *32*(8), 881-891.

Kiesel, A., Steinhauser, M., Wendt, M., Falkenstein, M., Jost, K., Philipp, A. M., & Koch, I. (2010). Control and interference in task switching--a review. *Psychological Bulletin*, *136*(5), 849-874.

Kim, C., Johnson, N. F., Cilles, S. E., & Gold, B. T. (2011). Common and Distinct Mechanisms of Cognitive Flexibility in Prefrontal Cortex. *The Journal of Neuroscience*, *31*(13), 4771 -4779.

Klapp, S. T., Anderson, W. G., & Berrian, R. W. (1973). Implicit speech in reading: Reconsidered. *Journal of Experimental Psychology*, *100*(2), 368-374.

Knapen, T., Brascamp, J., Pearson, J., van Ee, R., & Blake, R. (2011). The Role of Frontal and Parietal Brain Areas in Bistable Perception. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *31*(28), 10293-10301.

Knight, R. T. (2007). Neural Networks Debunk Phrenology. *Science*, *316*(5831), 1578 -1579.

Koechlin, Etienne, & Jubault, T. (2006). Broca's Area and the Hierarchical Organization of Human Behavior. *Neuron*, *50*(6), 963-974.

Koechlin, Etienne, & Summerfield, Christopher. (2007). An information theoretical approach to prefrontal executive function. *Trends in Cognitive Sciences*, *11*(6), 229-235.

Koechlin, Etienne, Basso, G., Pietrini, P., Panzer, S., & Grafman, Jordan. (1999). The role of the anterior prefrontal cortex in human cognition. *Nature*, *399*(6732), 148-151.

Koechlin, Etienne, Corrado, G., Pietrini, P., & Grafman, Jordan. (2000). Dissociating the role of the medial and lateral anterior prefrontal cortex in human planning. *Proceedings of the National Academy of Sciences of the United States of America*, *97*(13), 7651-7656.

Koechlin, Etienne, Ody, C., & Kouneiher, F. (2003). The Architecture of Cognitive Control in the Human Prefrontal Cortex. *Science*, *302*(5648), 1181-1185.

Koenigs, M., Barbey, A. K., Postle, B. R., & Grafman, Jordan. (2009). Superior Parietal Cortex Is Critical for the Manipulation of Information in Working Memory. *The Journal of Neuroscience*, *29*(47), 14980 -14986.

Kühn, S., Bodammer, N. C., & Brass, M. (2010). Dissociating mental states related to doing nothing by means of fMRI pattern classification. *NeuroImage*, *53*(4), 1294-1300.

Lashley, K. (1951). The problem of serial order in behaviour. *Cerebral Mechanisms in Behaviour* (pp. 112-136).

Levy, R., & Goldman-Rakic, P. S. (2000). Segregation of working memory functions within the dorsolateral prefrontal cortex. *Experimental Brain Research. Experimentelle Hirnforschung. Expérimentation Cérébrale*, *133*(1), 23-32.

Lhermitte, f. (1983). "Utilization behaviour" and its relation to lesions of the frontal lobes. *Brain*, *106*(2), 237 -255.

Li, S., Ostwald, D., Giese, M., & Kourtzi, Z. (2007). Flexible Coding for Categorical Decisions in the Human Brain. *J. Neurosci.*, *27*(45), 12321-12330.

Libet, B., Gleason, C. A., Wright, E. W., & Pearl, D. K. (1983). Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential). The unconscious initiation of a freely voluntary act. *Brain: A Journal of Neurology*, *106 (Pt 3)*, 623-642.

Lien, M.-C., & Ruthruff, E. (2004). Task Switching in a Hierarchical Task Structure: Evidence for the Fragility of the Task Repetition Benefit. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*(3), 697-713.

Logan, G. D., & Gordon, R. D. (2001). Executive control of visual attention in dual-task situations. *Psychological Review*, *108*(2), 393-434.

Louro, M.J.S., Pieters, R., & Zeelenberg, M. (2007). Dynamics of multiple goal pursuit. *Journal of Personality and Social Psychology*, 93(2), 174-193.

Lumer, E. D., & Rees, G. (1999). Covariation of activity in visual and prefrontal cortex associated with subjective visual perception. *Proceedings of the National Academy of Sciences of the United States of America*, *96*(4), 1669-1673.

Luria, A. R. (1966). *Higher cortical functions in man*. Oxford, England: Basic Books.

MacDonald, A. W., Cohen, J D, Stenger, V. A., & Carter, C. S. (2000). Dissociating the role of the dorsolateral prefrontal and anterior cingulate cortex in cognitive control. *Science (New York, N.Y.)*, *288*(5472), 1835-1838.

Marklund, P., Fransson, P., Cabeza, Roberto, Petersson, K. M., Ingvar, M., & Nyberg, Lars. (2007). Sustained and Transient Neural Modulations in Prefrontal Cortex Related to Declarative Long-Term Memory, Working Memory, and Attention. *Cortex*, *43*(1), 22-37.

Mason, M. F., Norton, M. I., Van Horn, J. D., Wegner, D. M., Grafton, S. T., & Macrae, C. N. (2007). Wandering Minds: The Default Network and Stimulus-Independent Thought. Science, 315(5810), 393–395.

Mayer, J. S., Roebroeck, A., Maurer, K., & Linden, D. E. J. (2009). Specialization in the default mode: Task-induced brain deactivations dissociate between visual working memory and attention. *Human Brain Mapping*, NA-NA.

Mesulam, M. M. (1999). Spatial attention and neglect: parietal, frontal and cingulate contributions to the mental representation and attentional targeting of salient extrapersonal events. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *354*(1387), 1325-1346.

Mikami, A., Ito, S., & Kubota, K. (1982). Visual response properties of dorsolateral prefrontal neurons during visual fixation task. *Journal of Neurophysiology*, *47*(4), 593-605.

Milea, D., Lehéricy, S., Rivaud-Péchoux, S., Duffau, H., Lobel, E., Capelle, L., Marsault, C., et al. (2003). Antisaccade deficit after anterior cingulate cortex resection. *Neuroreport*, *14*(2), 283-287.

Miller, E K, & Cohen, J D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, *24*, 167-202.

Miller, G. (1960). *Plans and the structure of behavior*. New York: Holt.

B Milner (1964). Some effects of frontal lobectomy in man. - *The frontal granular cortex and behavior*, 1964

Mirabella, G., Bertini, G., Samengo, I., Kilavik, B. E., Frilli, D., Della Libera, C., & Chelazzi, L. (2007). Neurons in area V4 of the macaque translate attended visual features into behaviorally relevant categories. *Neuron*, *54*(2), 303-318.

Mizobuchi, K., Takahasi, N., & Ajima, A. (2011). Utilization and imitation behavior following right parietotemporal lesions. *Rinshō Shinkeigaku = Clinical Neurology*, *51*(5), 350-353.

Monti, M. M., Vanhaudenhuyse, A., Coleman, M. R., Boly, M., Pickard, J. D., Tshibanda, L., Owen, A. M., et al. (2010). Willful Modulation of Brain Activity in Disorders of Consciousness. *New England Journal of Medicine*, *362*(7), 579-589.

Nachev, P., Kennard, C., & Husain, M. (2008). Functional role of the supplementary and pre-supplementary motor areas. *Nature Reviews Neuroscience*, *9*(11), 856-869.

Norman, D., & Shallice, T. (1986). Attention to Action: Willed and Automatic Control of Behavior. In R. J. Davidson (Ed.), *Consciousness and Self-regulation: Advances in Research and Theory* (Vol. 4, pp. 1-18). Plenum.

O'Reilly, R. C. (2010). The What and How of prefrontal cortical organization. *Trends in Neurosciences*, *33*(8), 355-361.

Owen, A. M., Coleman, M. R., Boly, M., Davis, M. H., Laureys, S., & Pickard, J. D. (2006). Detecting Awareness in the Vegetative State. *Science*, *313*(5792), 1402.

Pashler, H. E. (1999). *The Psychology of Attention*. MIT Press.

di Pellegrino, G., & Wise, S P. (1991). A neurophysiological comparison of three distinct regions of the primate frontal lobe. *Brain: A Journal of Neurology*, *114 (Pt 2)*, 951-978.

Perret, E. (1974). The left frontal lobe of man and the suppression of habitual responses in verbal categorical behaviour. *Neuropsychologia*, *12*(3), 323-330.

De Pisapia, N., Slomski, J. A., & Braver, T. S. (2007). Functional Specializations in Lateral Prefrontal Cortex Associated with the Integration and Segregation of Information in Working Memory. *Cereb. Cortex*, *17*(5), 993-1006.

Platt, M. L., & Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. *Nature*, *400*(6741), 233-238.

Posner, M., Walker, J., Friedrich, F., & Rafal, R. (1984). Effects of parietal injury on covert orienting of attention. *The Journal of Neuroscience*, *4*(7), 1863 -1874.

Povel, D.-J., & Collard, R. (1982). Structural factors in patterned finger tapping. *Acta Psychologica*, *52*(1-2), 107-123.

Procyk, E., & Joseph, J. P. (2001). Characterization of serial order encoding in the monkey anterior cingulate sulcus. *European Journal of Neuroscience*, *14*(6), 1041-1046.

Pujol, J., Vendrell, P., Deus, J., Junqué, C., Bello, J., MartI[combining acute accent]-Vilalta, J. L., & Capdevila, A. (2001). The Effect of Medial Frontal and Posterior Parietal Demyelinating Lesions on Stroop Interference. *NeuroImage*, *13*(1), 68-75.

Quintana, J., & Fuster, J M. (1992). Mnemonic and predictive functions of cortical neurons in a memory task. *Neuroreport*, *3*(8), 721-724.

Race, E. A., Shanker, S., & Wagner, A. D. (2009). Neural Priming in Human Frontal Cortex: Multiple Forms of Learning Reduce Demands on the Prefrontal Executive System. *Journal of Cognitive Neuroscience*, *21*(9), 1766-1781.

Raichle, Marcus E. (2003). Functional brain imaging and human brain function. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *23*(10), 3959-3962.

Rao, S. C., Rainer, G, & Miller, E K. (1997). Integration of what and where in the primate prefrontal cortex. *Science (New York, N.Y.)*, *276*(5313), 821-824.

Reed, E. S., Montgomery, M., Palmer, C., & Pittenger, J. (1995). Method for Studying the Invariant Knowledge Structure of Action: Conceptual Organization of an Everyday Action. *The American Journal of Psychology*, *108*(1), 37-65.

Restle, F., & Burnside, B. L. (1972). Tracking of serial patterns. *Journal of Experimental Psychology*, *95*(2), 299-307. doi:37/h0033619

Reuss, H., Kiesel, A., Kunde, W., & Hommel, B. (n.d.). Unconscious activation of task sets. *Consciousness and Cognition*, *In Press, Corrected Proof*.

Reverberi, C., Görgen, K., & Haynes, John-Dylan. (2011). Compositionality of Rule Representations in Human Prefrontal Cortex. *Cerebral Cortex*.

Reynolds, J. R., West, R., & Braver, T. (2009). Distinct Neural Circuits Support Transient and Sustained Processes in Prospective Memory and Working Memory. *Cerebral Cortex*, *19*(5), 1208 -1221.

Rigotti, M., Ben Dayan Rubin, D., Morrison, S. E., Salzman, C. D., & Fusi, S. (2010). Attractor concretion as a mechanism for the formation of context representations. *NeuroImage*, *52*(3), 833-847.

Robertson, I. H., Manly, T., Beschin, N., Daini, R., Haeske-Dewick, H., Hömberg, V., Jehkonen, M., et al. (1997). Auditory sustained attention is a marker of unilateral spatial neglect. *Neuropsychologia*, *35*(12), 1527-1532.

Roca, M., Parr, A., Thompson, Russell, Woolgar, A., Torralva, T., Antoun, N., Manes, F., et al. (2010). Executive function and fluid intelligence after frontal lobe lesions. *Brain*, *133*(1), 234 -247.

Rogers, B. P., Avery, S. N., & Heckers, S. (2010). Internal representation of hierarchical sequences involves the default network. *BMC Neuroscience*, *11*, 54.

Romo, R., Brody, C. D., Hernández, A., & Lemus, L. (1999). Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature*, *399*(6735), 470-473.

Rosenbaum, D. A., Kenny, S. B., & Derr, M. A. (1983). Hierarchical control of rapid movement sequences. *Journal of Experimental Psychology: Human Perception and Performance*, *9*(1), 86-102.

Rowe, J. B., Eckstein, D., Braver, T., & Owen, A. M. (2008). How Does Reward Expectation Influence Cognition in the Human Brain? *Journal of Cognitive Neuroscience*, *20*(11), 1980- 1992.

Russell, S., & Norvig, P. (2002). *Artificial Intelligence: A Modern Approach (2nd Edition)*. Prentice Hall.

Rylander, G., (1939). *Personality changes after operations on the frontal lobes: a clinical study of 32 cases*. - E. Munksgaard

Saalmann, Y. B., Pigarev, I. N., & Vidyasagar, T. R. (2007). Neural mechanisms of visual attention: how top-down feedback highlights relevant locations. *Science*, *316*(5831), 1612.

Saga, Y., Iba, M., Tanji, J., & Hoshi, E. (2011). Development of Multidimensional Representations of Task Phases in the Lateral Prefrontal Cortex. *The Journal of Neuroscience*, *31*(29), 10648 - 10665.

Sakai, K. (2008). Task Set and Prefrontal Cortex. *Annual Review of Neuroscience*, *31*(1), 219-245.

Sakai, K., & Passingham, R. E. (2006). Prefrontal Set Activity Predicts Rule-Specific Neural Processing during Subsequent Cognitive Performance. *The Journal of Neuroscience*, *26*(4), 1211 -1218.

Savine, A. C., & Braver, T. S. (2010). Motivated cognitive control: reward incentives modulate preparatory neural activity during task-switching. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *30*(31), 10294-10305.

Schall, U., Johnston, P., Todd, J., Ward, P. B., & Michie, P. T. (2003). Functional neuroanatomy of auditory mismatch processing: an event-related fMRI study of duration-deviant oddballs. *NeuroImage*, *20*(2), 729-736.

Schank, R., & Abelson, R. (1977). *Script, Plans, Goals and Understanding: An Inquiry into Human Knowledge Structures*. Lawrence Erlbaum Associates.

Schneider, D. W., & Logan, G. D. (2007). Retrieving information from a hierarchical plan. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *33*(6), 1076-1091.

Schneider, D. W., & Logan, G. D. (2006). Hierarchical Control of Cognitive Processes: Switching Tasks in Sequences. *Journal of Experimental Psychology: General*, *135*(4), 623-640.

Seeley, W. W., Menon, V., Schatzberg, A. F., Keller, J., Glover, G. H., Kenna, H., Reiss, A. L., et al. (2007). Dissociable Intrinsic Connectivity Networks for Salience Processing and Executive Control. *J. Neurosci.*, *27*(9), 2349–2356.

Serences, J. T., & Boynton, G. M. (2007). Feature-Based Attentional Modulations in the Absence of Direct Visual Stimulation. *Neuron*, *55*(2), 301-312.

Shaw, G. L., & McGaugh, J. L. (1990). *Neurobiology of learning and memory*. World Scientific.

Shidara, M., & Richmond, Barry J. (2002). Anterior cingulate: single neuronal signals related to degree of reward expectancy. *Science (New York, N.Y.)*, *296*(5573), 1709-1711.

Shidara, M., Mizuhiki, T., & Richmond, Barry J. (2005). Neuronal firing in anterior cingulate neurons changes modes across trials in single states of multitrial reward schedules. *Experimental Brain Research*, *163*(2), 242-245.

Shimamura, Arthur P., Janowsky, Jeri S., Squire, Larry R. (1991). What is the role of frontal lobe damage in memory disorders? Frontal lobe function and dysfunction. Levin, Harvey S. (Ed); Eisenberg, Howard M. (Ed); Benton, Arthur L. (Ed), (1991). *Frontal lobe function and dysfunction,* (pp. 173-195). New York, NY, US: Oxford University Press, xv, 427 pp.

Sigala, N., Kusunoki, M., Nimmo-Smith, I., Gaffan, D., & Duncan, John. (2008). Hierarchical coding for sequential task events in the monkey prefrontal cortex. *Proceedings of the National Academy of Sciences*, *105*(33), 11969 -11974.

Singh, K. D., & Fawcett, I. P. (2008). Transient and linearly graded deactivation of the human default-mode network by a visual detection task. *NeuroImage*, *41*(1), 100-112.

Sirigu, A, Cohen, L., Zalla, T., Pradat-Diehl, P., Van Eeckhout, P., Grafman, J, & Agid, Y. (1998). Distinct frontal regions for processing sentence syntax and story grammar. *Cortex; a Journal Devoted to the Study of the Nervous System and Behavior*, *34*(5), 771-778.

Sirotin, Y. B., & Das, A. (2009). Anticipatory haemodynamic signals in sensory cortex not predicted by local neuronal activity. *Nature*, *457*(7228), 475-479.

Smania, N., Martini, M. C., Gambina, G., Tomelleri, G., Palamara, A., Natale, E., & Marzi, C. A. (1998). The spatial distribution of visual attention in hemineglect and extinction patients. *Brain*, *121*(9), 1759 -1770.

Smith, R., Keramatian, K., & Christoff, K. (2007). Localizing the rostrolateral prefrontal cortex at the individual level. *NeuroImage*, *36*(4), 1387-1396.

Sohn, J.-W., & Lee, D. (2007). Order-Dependent Modulation of Directional Signals in the Supplementary and Presupplementary Motor Areas. *The Journal of Neuroscience*, *27*(50), 13655 -13666.

Soon, Chun Siong, Brass, Marcel, Heinze, H.-J., & Haynes, John-Dylan. (2008). Unconscious determinants of free decisions in the human brain. *Nat Neurosci*, *11*(5), 543-545.

Spearman, C. E. (1927). *The abilities of man: their nature and measurement*. The Macmillan Company.

Sridharan, D., Levitin, D. J., & Menon, V. (2008). A critical role for the right fronto-insular cortex in switching between central-executive and default-mode networks. *Proceedings of the National Academy of Sciences*, *105*(34), 12569 -12574.

Sridharan, D., Levitin, D. J., Chafe, C. H., Berger, J., & Menon, V. (2007). Neural Dynamics of Event Segmentation in Music: Converging Evidence for Dissociable Ventral and Dorsal Networks. *Neuron*, *55*(3), 521-532.

Stevens, A. A., Skudlarski, P., Gatenby, J. C., & Gore, J. C. (2000). Event-related fMRI of auditory and visual oddball tasks. *Magnetic Resonance Imaging*, *18*(5), 495-502.

Stiers, P., Mennes, M., & Sunaert, S. (2010). Distributed task coding throughout the multiple demand network of the human frontal-insular cortex. *NeuroImage*, *52*(1), 252-262.

Stokes, M., Thompson, R., Cusack, R., & Duncan, J. (2009). Top-down activation of shape-specific population codes in visual cortex during mental imagery. *Journal of Neuroscience*, *29*(5), 1565-1572.

Stroop, J. R. (1935). The Basis of Ligon's Theory. *The American Journal of Psychology*, *47*(3), 499-504.

Stuphorn, V., & Schall, J. D. (2006). Executive control of countermanding saccades by the supplementary eye field. *Nat Neurosci*, *9*(7), 925-931.

Summerfield, C., Egner, T., Greene, M., Koechlin, E., Mangels, J., & Hirsch, J. (2006). Predictive codes for forthcoming perception in the frontal cortex. *Science*, *314*(5803), 1311.

Swick, D., & Turken, A. U. (2002). Dissociation between conflict detection and error monitoring in the human anterior cingulate cortex. *Proceedings of the National Academy of Sciences*, *99*(25), 16354 -16359.

Tamber-Rosenau, B. J., Esterman, M., Chiu, Y.-C., & Yantis, S. (2011). Cortical mechanisms of cognitive control for shifting attention in vision and working memory. *Journal of Cognitive Neuroscience*, *23*(10), 2905-2919.

Taren, A. A., Venkatraman, V., & Huettel, S. A. (2011). A Parallel Functional Topography between Medial and Lateral Prefrontal Cortex: Evidence and Implications for Cognitive Control. *The Journal of Neuroscience*, *31*(13), 5026 -5031.

Thompson-Schill, S.L., D'Esposito, M.D., Aguirre, G.K. & Farah, M.J. (1997). Role of left inferior prefrontal cortex in retrieval of semantic knowledge: A reevaluation. Proceedings of the National Academy of Sciences USA , 94: 14792–14797

Thompson, Russell, & Duncan, John. (2009). Attentional modulation of stimulus representation in human fronto-parietal cortex. *NeuroImage*, *48*(2), 436-448.

Todd, J. J., Fougnie, D., & Marois, R. (2005). Visual Short-Term Memory Load Suppresses Temporo-Parietal Junction Activity and Induces Inattentional Blindness. *Psychological Science*, *16*(12), 965 -972.

Tsujimoto, S., Genovesio, A., & Wise, Steven P. (2010). Evaluating self-generated decisions in frontal pole cortex of monkeys. *Nat Neurosci*, *13*(1), 120-126.

Uddin, L. Q., Clare Kelly, A. M., Biswal, B. B., Xavier Castellanos, F., & Milham, M. P. (2009). Functional connectivity of default mode network components: Correlation, anticorrelation, and causality. *Human Brain Mapping*, *30*(2), 625-637.

Uretzky, S., & Gilboa, A. (2010). Knowing your lines but missing your cue: rostral prefrontal lesions impair prospective memory cue detection, but not action-intention superiority. *Journal of Cognitive Neuroscience*, *22*(12), 2745-2757.

Velanova, K., Jacoby, L. L., Wheeler, M. E., McAvoy, M. P., Petersen, S. E, & Buckner, R. L. (2003). Functional–anatomic correlates of sustained and transient processing components engaged during controlled retrieval. *The Journal of neuroscience*, *23*(24), 8460.

Vendrell, P., Junqué, C., Pujol, J., Jurado, M. A., Molet, J., & Grafman, Jordan. (1995). The role of prefrontal regions in the Stroop task. *Neuropsychologia*, *33*(3), 341-352.

Volle, E., Gonen-Yaacovi, G., de Lacy Costello, A., Gilbert, Sam J, & Burgess, P. W. (2011). The role of rostral prefrontal cortex in prospective memory: A voxel-based lesion study. *Neuropsychologia*, *49*(8), 2185-2198.

Wallis, J D, Anderson, K. C., & Miller, E K. (2001). Single neurons in prefrontal cortex encode abstract rules. *Nature*, *411*(6840), 953-956.

Wallis, Jonathan D, & Miller, Earl K. (2003). From rule to response: neuronal processes in the premotor and prefrontal cortex. *Journal of Neurophysiology*, *90*(3), 1790-1806.

Watanabe, M. (1990). Prefrontal unit activity during associative learning in the monkey. *Experimental Brain Research. Experimentelle Hirnforschung. Expérimentation Cérébrale*, *80*(2), 296-309.

Watanabe, M. (1992). Frontal units of the monkey coding the associative significance of visual and auditory stimuli. *Experimental Brain Research. Experimentelle Hirnforschung. Expérimentation Cérébrale*, *89*(2), 233-247.

Watanabe, M. (1996). Reward expectancy in primate prefrontal neurons. *Nature*, *382*(6592), 629-632.

Weissman, D. H., Roberts, K. C., Visscher, K. M., & Woldorff, M. G. (2006). The neural bases of momentary lapses in attention. *Nat Neurosci*, *9*(7), 971–978.

White, I. M., & Wise, S P. (1999). Rule-dependent neuronal activity in the prefrontal cortex. *Experimental Brain Research. Experimentelle Hirnforschung. Expérimentation Cérébrale*, *126*(3), 315-335.

Womelsdorf, T., Schoffelen, J. M., Oostenveld, R., Singer, W., Desimone, R., Engel, A. K., & Fries, P. (2007). Modulation of neuronal interactions through neuronal synchronization. *Science*, *316*(5831), 1609.

Woolgar, A., Parr, A., Cusack, R., Thompson, R., Nimmo-Smith, I., Torralva, T., Roca, M., et al. (2010). Fluid intelligence loss linked to restricted regions of damage within frontal and parietal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *107*(33), 14899–14902.

Woolgar, A., Thompson, Russell, Bor, D., & Duncan, John. (2011). Multi-voxel coding of stimuli, rules, and responses in human frontoparietal cortex. *NeuroImage*, *56*(2), 744-752.

Yeung, N., Nystrom, L. E., Aronson, J. A., & Cohen, Jonathan D. (2006). Between-task competition and cognitive control in task switching. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *26*(5), 1429-1438.

Yoshida, W., Funakoshi, H., & Ishii, S. (2010). Hierarchical rule switching in prefrontal cortex. *NeuroImage*, *50*(1), 314-322.

Young, M., & Yamane, S. (1992). Sparse population coding of faces in the inferotemporal cortex. *Science*, *256*(5061), 1327 -1331.

Zacks, J. M., Braver, T. S., Sheridan, M. A., Donaldson, D. I., Snyder, Abraham Z., Ollinger, J. M., Buckner, Randy L., et al. (2001). Human brain activity time-locked to perceptual event boundaries. *Nat Neurosci*, *4*(6), 651-655.