Frontal and parietal participation in problem solving in the Tower of London: fMRI and computational modeling of planning and high-level perception

Sharlene D. Newman*, Patricia A. Carpenter, Sashank Varma, Marcel Adam Just

*Department of Psychology, Center for Cognitive Brain Imaging, Carnegie Mellon University, Baker Hall,
3275000 Forbes Avenue, Pittsburgh, PA 15213, USA*

Received 7 September 2001; received in revised form 7 May 2002; accepted 25 March 2003

Abstract

This study triangulates executive planning and visuo-spatial reasoning in the context of the Tower of London (TOL) task by using a variety of methodological approaches. These approaches include functional magnetic resonance imaging (fMRI), functional connectivity analysis, individual difference analysis, and computational modeling. A graded fMRI paradigm compared the brain activation during the solution of problems with varying path lengths: easy (1 and 2 moves), moderate (3 and 4 moves) and difficult (5 and 6 moves). There were three central findings regarding the prefrontal cortex: (1) while both the left and right prefrontal cortices were equally involved during the solution of moderate and difficult problems, the activation on the right was differentially attenuated during the solution of the easy problems; (2) the activation observed in the right prefrontal cortex was highly correlated with individual differences in working memory (measured independently by the reading span task); and (3) different patterns of functional connectivity were observed in the left and right prefrontal cortices. Results obtained from the superior parietal region also revealed left/right differences; only the left superior parietal region revealed an effect of difficulty. These fMRI results converged upon two hypotheses: (1) the right prefrontal area may be more involved in the generation of a plan, whereas the left prefrontal area may be more involved in plan execution; and (2) the right superior parietal region is more involved in attention processes while the left homologue is more of a visuo-spatial workspace. A 4CAPS computational model of the cognitive processes and brain activation in the TOL task integrated these hypothesized mechanisms, and provided a reasonably good fit to the observed behavioral and brain activation data. The multiple research approaches presented here converge on a deepening understanding of the combination of perceptual and conceptual processes in this type of visual problem solving.
© 2003 Elsevier Science Ltd. All rights reserved.

Keywords: Planning; fMRI; Spatial working memory; Problem solving; Tower of London; Computational modeling; 4CAPS

1. Introduction

The ability to plan and schedule one's actions is an essential part of a functional life. Planning is a superordinate term that subsumes several subprocesses, including strategy formation, coordination and sequencing of mental functions, and holding information on-line (Morris, Miotto, Feigenbaum, Bullock, & Polkey, 1997). Both neuropsychological and neuroimaging studies have suggested that the prefrontal cortex is an important part of the cortical network involved in planning. One task that has been used to assess planning function is the Tower of London task (TOL) (Baker, Rogers, Owen, Frith, Dolan, Frackowiak, & Robbins, 1996; Dagher, Owen, Boecker, & Brooks, 1999; Morris, Ahmed,

Syed, & Toone, 1993; Owen, Downes, Sahakian, Polkey, & Robbins, 1990; Shallice, 1982). The current study uses both a functional magnetic resonance imaging (fMRI) activation study, as well as a computational model to examine plan formation and execution in the TOL task.

In the TOL paradigm used in the current study, participants were presented with two configurations (a start state and a goal state) of three balls arrayed in three bins. The easiest type of problem is shown in Fig. 1. Participants were asked to plan how to move the balls from the start configuration to match the goal configuration in the minimum number of moves. Their overt response concerned the number of moves that were required. Thus, this TOL task required participants to "look ahead" and map out a plan to solve the problem. Constraints on the possible moves are determined by the different depths of the bins (the three bins hold three, two and one ball, respectively) and being able to

* Corresponding author. Tel.: +1-412-268-6859; fax: +1-412-268-2804.
E-mail address: snewan@andrew.cmu.edu (S.D. Newman).

Fig. 1. A sample TOL problem. This display shows a 1-move problem, with the start state on the left, the goal state on the right, and a schematic diagram showing the response button assignments on the bottom. (The display was presented as a white figure on a black background.)

move only the top-most balls to a different bin. In addition to being a test of planning, the TOL task also places considerable demands on spatial working memory. Therefore, the present research examines the neural basis of some of the visuo-spatial as well as the planning processes necessary to perform the TOL task.

A computational model of the TOL task was developed in the cortical capacity-constrained collaborative activation-based production systems (4CAPS) cognitive neuroarchitecture, which is a production system architecture with several connectionist features (for more details on 4CAPS refer to Just & Varma, 2003; Just, Carpenter, & Varma, 1999). 4CAPS models function at the interface between the cortical and cognitive levels of analysis. Conventional production systems are the simulation medium of choice for high-level cognition. Some of the prominent ones, including Soar (Newell, 1990), ACT-R (Anderson, 1993), EPIC (Meyer & Kieras, 1997) and 3CAPS (Goel, Pullara, & Grafman, 2001; Just & Carpenter, 1992), account for the error patterns and processing times in problem solving, reasoning, decision making, memory and learning, language comprehension and visual thinking.

In addition to accounting for the behavioral characteristics as the systems above do, 4CAPS models also attempt to account for the amount of brain activity in a number of cortical areas as measured by functional neuroimaging. 4CAPS models consist of a number of processing centers each corresponding to a particular cortical area. Each 4CAPS center is a hybrid symbolic-connectionist processing system with its own specializations and computational resources. Productions, or if-then rules, implement these processes. The productions do their work by incrementally manipulating the activation levels of representational elements. Each center possesses a finite amount of resources paralleling the biological and informational constraints of cortical areas. The *capacity utilization* of a center is an index of its resource consumption over time. Because the actions of the productions consume resources, the total resource consumption of a center can be measured (either instantaneously or averaged over an interval of time). A key claim of 4CAPS is that the capacity utilization of a center should correspond to the relative level of fMRI-measured activation in the cor-

responding brain area, modulo the hemodynamic response function. Performing a task entails neural computation in the brain areas with relevant competencies. Because the modeled brain areas perform these theorized computations, capacity utilization can be considered an idealized measure of the activity level of the region. Information processing within each center is highly collaborative in the sense that productions fire in parallel. 4CAPS models are also collaborative in that centers interact with each other when performing a task. The pattern of interaction is dictated by the functional demands of the task modeled and the specializations attributed to each of the centers of the model.

The TOL task here requires two main types of cognitive processes: executive/planning processes and visuo-spatial processes. The 4CAPS model of TOL distributes these two types of processes across four centers: RH-EXECUTIVE, LH-EXECUTIVE, RH-SPATIAL and LH-SPATIAL. The RH- and LH-EXECUTIVE centers are proposed to correspond to right- and left-hemisphere dorsolateral prefrontal cortex (DLPFC), respectively. The model is intended to be consistent with the established literature and theory of prefrontal executive function associated with these areas. In general, the RH-EXECUTIVE center is proposed to have specialized cognitive functions for goal-management, planning, and strategy formulation. In contrast, the LH-EXECUTIVE center is proposed to have specialized cognitive functions for controlling the execution of a plan or of a sequence of mental actions. The RH- and LH-SPATIAL centers are intended to correspond to right- and left-hemisphere superior parietal cortex, respectively, including the intraparietal sulcus area. As in the case of the two EXECUTIVE centers, we propose a distinction in the functions of the two SPATIAL centers, based on emerging evidence. The RH-SPATIAL center is hypothesized to be more involved in the control of spatial attention and the geometric manipulation of spatial representations (Carpenter, Just, Keller, Eddy, & Thulborn, 1999; Chelazzi & Corbetta, 2000; Mesulam, 1990), whereas the LH-SPATIAL center is proposed to be more involved with constructing and maintaining spatial representations. The model's main predictions concern the relative amount of activation in the four areas in each of the three experimental conditions. (Further details regarding the specific processes attributed to each center are described in the discussion section and details regarding the problem-solving heuristics used are outlined in Appendix A.)

Functional neuroimaging studies typically measure cortical activation to assess the involvement of a particular region in a given task. The qualitative properties of the task are assumed to be critical to the pattern of activation. The current study goes beyond the qualitative aspects to examine the quantitative properties of the brain activation within the context of a theoretical framework. We relate the amount of cognitive demand in each experimental condition (as measured in the 4CAPS model) to the amount of cortical activity in specific areas, as measured with fMRI.

Furthermore, we examine the degree of synchronization (functional connectivity) between the activation in key cortical areas as a function of the number of moves in the problem (Diwadkar, Carpenter, & Just, 2000). Finally, we explore the impact of individual differences in working memory on patterns of brain activation, by examining the correlation between levels of brain activation and measures of working memory capacity. The major objective is to determine how the amount of neural activation in a network of brain regions is modulated by the amount and type of processing that is required in a given condition. Difficulty is manipulated here as the minimal number of moves (or solution path length) required to solve the problem. As the number of moves increases, the working memory load, as well as demands on planning processes increase. Therefore, in the current experiment we varied the path length so that it ranged from easy (1–2 moves problems) to moderate (3–4 move problems) to difficult (>5 moves). This allows for the observation of parametric changes in the activation. It also allows for substantive evaluation of the TOL model and the 4CAPS cognitive neuroarchitecture in which it is embedded.

The goal of the current study was to use the methods outlined above to further characterize the contribution of prefrontal and parietal regions in planning and visuo-spatial processing in the TOL task. One specific hypothesis examined is that there are hemispheric processing differences in TOL in both prefrontal and parietal cortex. We suggest, and there is some indication in the literature (Burgess, Veitch, de Lacy Costello, & Shallice, 2000; Feigenbaum, Polkey, & Morris, 1996; Miotto, Bullock, Polkey, & Morris, 1996; Morris et al., 1997), that the right prefrontal cortex is differentially involved in the formulation of plans, while left prefrontal cortex is involved in the execution of that plan. In addition, we suggest, and again there is some evidence in the literature, that right superior parietal cortex is differentially involved in spatial attention while left superior parietal cortex may be more involved in maintaining visuo-spatial information. These hypotheses concern differential, not absolute involvement, because the areas involved appear to closely collaborate.

2. Method

2.1. Participants

Sixteen right-handed (12 males and 4 females) Carnegie Mellon University undergraduate students participated in the experiment. All of the participants gave informed consent that was approved by the University of Pittsburgh and Carnegie Mellon Institutional Review Boards. Data from three participants were excluded due to excessive out-of-plane motion during scanning, and data from one additional participant was excluded because of excessive in-plane motion (mean displacement > 0.3 voxels).

2.2. Experimental paradigm

The experiment consisted of a practice session and a scanning session. In the practice session, participants solved two sets of Tower of London problems. During the first 19 practice problems, both the start state and goal state were presented on a computer screen, and participants used a mouse to actually move the balls from bin to bin, beginning in the start state, until they matched the goal state. This format ensured that participants were thoroughly familiarized with the Tower of London problem-solving procedures. The format of the next 18 practice problems was identical to that used in the scanner during the test session. On each trial, the start and goal states of a Tower of London problem were presented (see Fig. 1). But rather than moving the balls, the participants were asked to imagine the moves and count the total number of moves required, and then indicate (by pressing the appropriate response button) the total number of moves. As Fig. 1 indicates, the display contained a schematic diagram of the buttons and their correct response assignments. All of the presented problems required 1–2, 3, 4, or >5 moves.

Participants were familiarized with the fMRI scanner and the general scanning procedures prior to being scanned. During the test session, participants solved 12 epochs containing 3–6 Tower of London problems each. The task was self-paced, and each problem display remained visible until a response was made. Both response latencies and accuracies were collected. Epochs were separated by a 6 s rest period in which participants fixated a plus sign in the center of the screen. The images collected during these rest periods and the first 6 s of each epoch were discarded to accommodate the rise and fall of the hemodynamic response (Bandettini, Wong, Hinks, Tokofsky, & Hyde, 1992). Five 30 s fixation periods were also included to provide a baseline measure of cortical activation.

Problem difficulty (as defined by the minimum number of moves to solve a problem) was the independent variable. All of the problems had a unique solution. The epochs were constructed to give three average levels of path length, with a small amount of overlap across epochs to make the solutions within an epoch unpredictable and to require at least two different responses within each epoch: (1) *Easy* epochs contained five 1- and 2-move problems and one 3-move problem; (2) *Moderate* epochs contained three 3- and 4-move problems and one 1-move problem; and (3) *Difficult* epochs contained two 5- and 6-move problems and one 1-move problem. The durations of the epochs were approximately equated by having more items in those epochs that contained easier problems. The mean reaction times for each problem type was estimated based on data obtained during a pilot study.

2.3. Psychometric testing

The Daneman and Carpenter (1980) reading span test was administered to 10 of 12 participants. This test required

participants to simultaneously read aloud a small (1–5) set of sentences while attempting to recall the last word of each sentence at the end of the set. The number of sentences read was incremented until the participants failed to recall one or more of the sentence-final words on 50% of the trials. The number of words recalled at this point was defined as the participant's reading span. Because the reading span task is increasingly being used to assess central executive functioning in studies of the effects of frontal lobe damage and working memory (Engle & Oransky, 1999; Whitney, Arnett, Driver, & Budd, 2001), we used it here as a measure of executive functioning.

2.4. fMRI procedure

The experiment was conducted on a GE 3.0T scanner using a commercial birdcage, quadrature-drive radio-frequency whole-head coil. Fourteen oblique-axial images (5 mm thick, skipping 1 mm between slices) were chosen to maximize the coverage of each participant's cerebral cortex while minimizing coverage of the eyes and their movement. The images were collected using a gradient echo, resonant echo planar pulse sequence, with TE = 25 ms, 90° flip-angle, and a 128 × 64 acquisition matrix with 3.125 mm × 3.125 mm × 5 mm voxels, with a TR (time per repetition or sampling rate) of once of every 3000 ms.

The means of the images corresponding to each participant's 14 functional slices were registered to a high-resolution, T1-weighted structural volume scan. The volume scans were constructed from 124 3D SPGR axial images that were collected with TR = 25 ms, TE = 4 ms, 40° flip-angle, and a 24 cm × 18 cm FOV, resulting in 0.9375 mm × 0.9375 mm × 1.5 mm voxels.

Because problem solving tasks like the TOL are thought to engage a large-scale network of cortical regions (Mesulam, 1990, 1998), our fMRI analyses focused on measuring the modulation of the volume of activation from four regions of interest (ROIs): left and right superior parietal and DLPFC (see Fig. 2). However, we also examined the activation

within the inferior frontal gyrus and the inferior parietal cortex. The ROIs were anatomically defined to allow comparison of activation within some specified cortical region across conditions. The ROIs were defined according to the Rademacher, Galaburda, Kennedy, Flilipek, and Caviness (1992), and Caviness, Meyer, Makris, and Kennedy (1996) parcellation scheme, which uses limiting sulci and coronal planes defined by anatomical landmarks to segment cortical regions. The limiting sulci and anatomical landmarks were located by simultaneously viewing the structural images in the three orthogonal planes; the ROIs were then defined by manually tracing the regions onto each functional slice. The superior parietal ROI (SPL) was defined as BA 5 and 7, while the inferior parietal ROI (IPL) was defined as BA 40 and 39. The dorsolateral prefrontal cortex included the middle frontal gyrus F2, or approximately BA 10, 9, 46 and portions of 6 and 8. The inferior frontal gyrus (IFG) was defined as BA 44, 45 and 47. The inter-rater reliability of this ROI-defining procedure between two trained staff members was previously evaluated for four ROIs in two participants in another study. The reliability measure was obtained by dividing the size of the set of voxels that overlapped between the two raters by the mean of their two set sizes. The resulting eight reliability measures were in the 78–91% range, with a mean of 84%, as high as the reliability reported by the developers of the parcellation scheme.

2.5. Data analysis

Image preprocessing (including de-ghosting, mean correction, motion correction, and trend correction) was performed using FIASCO (Eddy, Fitzgerald, Genovese, Mockus, & Noll, 1996; Lazar, Eddy, Genovese, & Welling, 2001; further description and tools are available at <http://www.stat.cmu.edu/~fiasco/>). The mean of the maximum head motion per participant was 0.2 voxels, and it never exceeded 0.3 voxels. To accommodate the rise and fall time of the hemodynamic response data from the first 6 s of each epoch and the 6 s rest interval between epochs were discarded. Excluded from analysis were the 1% of all voxels showing more than 6% change in signal intensity that might have possibly arisen from blood vessels.

fMRI-measured activation was quantified by first constructing a *t*-map by computing the difference between each voxel's activation in each experimental condition and the baseline condition. Voxels whose signal change exceeded baseline by an amount determined by *t*-tests with *t* > 6 were considered "active." This high *t*-threshold provides approximately similar or more conservative correction for multiple comparisons as compared to a Bonferroni correction with an alpha level of 0.01 given approximately 5000 voxels per participant in all regions of interest. A measure called the sum of the change in signal intensity (SSI) was calculated for each ROI and each participant, to take into consideration changes in both the signal intensity and in the activation volume (see Xiong, Rao, Gao, Woldorff, & Fox, 1998 for details).

Fig. 2. A schematic depiction of the DLPFC (blue) and superior parietal (red) ROIs.

The SSI is calculated by summing the percent change in signal intensity for each active voxel within an ROI.

3. Results

3.1. Human participant results

3.1.1. Behavioral results

The behavioral data indicate that the participants followed instructions and solved the problems with a high degree of

accuracy (0.90). Fig. 3A shows the mean reaction time for those problems that were solved correctly and the proportion of problems solved. As expected, the time necessary to solve a problem increased monotonically with path length (i.e. the number of moves). A one-way ANOVA indicated that the effect of this variable was statistically reliable, $F(5, 55) = 19.46$, $P < 0.01$. A similar ANOVA on the proportion of incorrect responses also indicated a reliable effect of path length, $F(5, 55) = 4.33$, $P < 0.01$. The slight decrease in error rate for the longest length problems arose because participants sometimes responded by pressing the >5 move

Fig. 3. (A) The mean response latencies for correctly solved TOL problems and the percentage of problems solved incorrectly, both as a function of path length (i.e. the minimal number of moves needed to solve a problem). (B) The computational model's processing time for each problem. As shown, the model's processing time reveals a similar slope as the participants' reaction time, both showing a sharp increase with number of moves.

Table 1
Activation

ROI	Sum signal intensity (%)			Average coordinates		
	Easy	Moderate	Difficult	x	y	z
Left dorsolateral prefrontal cortex (BA 8, 9, 10, 46)	7.3	10.2	14.3	36.5	−26.8	35.4
Right dorsolateral prefrontal cortex	2.8	9.0	12.3	−32.5	−24.0	40.7
Left superior parietal cortex (BA 5 and 7)	19.6	29.7	32.5	18.9	67.4	46.3
Right superior parietal cortex	31.4	34.7	36.9	−19.9	65.2	47.4
Left inferior frontal gyrus (BA 44 and 45)	4.8	8.6	8.1	40.8	−14.1	27.9
Right inferior frontal gyrus	1.6	3.3	2.1	−37.4	−19.7	22.9
Left inferior parietal cortex (BA 39 and 40)	20.2	29.6	31.8	33.8	56.9	45.2
Right inferior parietal cortex	19.6	20.8	20.7	−36.9	53.3	44.3

button during very time-consuming problems before they had completed the solution. However, because the reaction time continues to increase with path length, it may be assumed that participants were engaged in the solution of the 5- and 6-move problems.

3.1.2. fMRI results

The TOL task activated several brain regions, including the dorsolateral prefrontal cortex, the inferior frontal gyrus, the superior parietal cortex, and the inferior parietal cortex, as shown in Table 1 and Fig. 4. Because no significant differences in the level of activation as a function of difficulty were observed in either the inferior frontal gyrus or the inferior parietal region, the focus of the description of the results below is on DLPFC and the superior parietal area, although all of the results are shown in Table 1. For both cortical regions a set of three ANOVAs were computed. The first was a three by two ANOVA in which difficulty (easy versus moderate versus difficult) and laterality (left versus right) as within-participant variables. The remaining two analyses were one-way ANOVAs performed for each hemisphere with difficulty as the within-participant variable.

3.1.2.1. DLPFC ROIs. As predicted, the mean change in the sum of the signal intensity in the DLPFC increased with path length, $F(2, 22) = 5.97$, $P < 0.01$, as shown in Fig. 5A. Moreover, neither the main effect of laterality nor its interaction with difficulty were reliable, $F < 1$. Further analysis of each hemisphere separately revealed that the effect of difficulty was significant in both left and right DLPFC, $F(2, 22) = 5.61$, $P < 0.05$ and $F(2, 22) = 4.35$, $P < 0.05$, respectively. These results indicate that both left and right DLPFC were involved in the TOL processing to an approximately similar extent during the performance of the moderate and difficult problems. The group averaged activation pattern in one slice is presented in Fig. 4.

3.1.2.2. Parietal ROIs. Path length also significantly affected the activation of the superior parietal region, $F(2, 22) = 4.42$, $P < 0.05$, as shown in Fig. 6A. In addition, as in the DLPFC, there was no main effect of laterality, $F(1, 22) = 2.05$, $P > 0.1$. However, unlike the DLPFC, only the left superior parietal ROI revealed a reliable effect of number of moves [$F(2, 22) = 7.6$, $P < 0.005$] while the right superior parietal region did not, $F < 1$. While these

Fig. 4. The Talairach averaged pattern of activation for each of the three levels of path length superimposed on the structural images of a single individual. The left-most column depicts the activation for the easy problems, the middle column moderate problems, and the right-most column the difficult problems. As shown, the amount of activation increases with difficulty in the prefrontal cortex as well as the superior parietal cortex.

Fig. 5. fMRI and model results for the DLPFC. (A) The sum signal intensity for both left and right DLPFC. As shown, the activation level for the right DLPFC during the easy problems is significantly less than for the left while they are quite similar for the moderate and difficult problems. (B) The capacity utilization for the model's left and right EXECUTIVE centers, which corresponds to the left and right DLPFC.

results indicate that both left and right superior parietal regions were clearly involved in the task, they also suggest that the two regions may be performing somewhat different functions.

3.1.2.3. Individual differences. To examine the relationship between cortical activation and a psychometric measure of working memory and executive functioning, the correlation between the SSI measure in each condition and the participant's reading span score was computed. Participants with higher reading span scores produced less cortical activation in the right DLPFC than individuals with a lower reading span during the solution of the longest path problems, $r = -0.67$, $P < 0.05$, as shown in Fig. 7 (left DLPFC $r = -0.12$). In addition, a t -test comparing these two correlation coefficients found that the correlation between the activation of the right prefrontal cortex and reading span is significantly different from the correlation between left prefrontal cortex activation and reading span ($t = 2.23$, $\alpha = 0.05$). When a correlation between an ability level and activation level is observed, the relation is generally like this one, with less activation with greater ability, reflecting something like cortical efficiency (Haier, Siegel, Neuchterlein, Hazlett, Wu, Paek, Browning, & Buchsbaum, 1988; Just, Carpenter,

Fig. 6. fMRI and model results for the superior parietal region. (A) The sum signal intensity for both left and right superior parietal cortex. As shown, while the left superior parietal region reveals a prominent step-wise increase in activation with difficulty, its right homologue reveals a significantly attenuated step-wise increase. (B) The capacity utilization for the model's left and right SPATIAL centers, which corresponds to the processing taking place in left and right superior parietal cortex.

& Miyake, 2003; Reichle, Carpenter, & Just, 2000). The correlation for the easy and moderately difficult problems failed to reach significance (right DLPFC $r = -0.06$ and -0.37 and left DLPFC $r = 0.07$ and 0.01 , respectively).

3.1.2.4. Planning time. Previous studies examining the Tower of London as well as the Tower of Hanoi (TOH) have examined the effect of planning time on the level of activation (Morris et al., 1993; Rowe, Owen, Johnsrude, & Passingham, 2001). It is thought that areas whose activation levels are highly correlated with the planning time are more intimately involved in planning processes. Planning time is typically defined as the time between stimulus onset and the onset of the first move. In studies in which participants actually move the balls, they are typically asked to plan the entire sequence of moves before they make their first move. In the current task, participants do not physically move the balls. It can be argued that the reaction time measure in the current study corresponds approximately to the planning time measured in previous studies. We calculated the correlation between mean response time for each participant and their corresponding activation level for each ROI. The

Fig. 7. Scatter plot indicating the correlation between the volume of activation and reading span for right DLPFC for the difficult problems. One of the 12 participants failed to undergo the reading span task and there are some participants with the same reading span score and the same activation volume resulting in the appearance of nine data points. As shown, the activation volume is significantly less for those participants with high reading spans compared to those with low span scores.

correlation between SSI during the solution of the longest path problems and response time was highly significant for both left and right superior parietal regions ($r = 0.66$, $P < 0.05$ and $r = 0.93$, $P < 0.05$ respectively), as seen in the scatter plots in Fig. 8. The correlations between response time and the activation within the prefrontal regions ($r = 0.71$, $P < 0.05$ for left and $r = 0.47$, NS, for the right DLPFC) were also relatively high.

3.1.2.5. Functional connectivity. The activation in a set of cortical areas is highly synchronized, indicating collaboration among areas. An increasingly used technique measures the correlation of the activation levels in two activated areas over some time period, and generally shows systematic synchronization between areas, modulated by a number of variables. The synchronization is taken as evidence of *functional connectivity* [or *effective connectivity* (Buchel, Coull, & Friston, 1999; Friston, 1994; Horwitz, Rumsey, & Donohue, 1998)]. The functional connectivity between ROIs was assessed by computing the correlation of the activation time-series between activated voxels in pairs of ROIs. Functionally connected areas might be collaborating or communicating, such that their activation levels are being modulated by a partially overlapping workload. The correlations were based on only those periods when the task was being performed (excluding the fixation periods), so that the time series indicates the momentary fluctuations in activation level during the task performance. The correlation between the averaged time courses from the activated voxels in each member of a pair of ROIs was computed for each participant. An ANOVA compared the functional connectivities of each ROI pair across the three levels of difficulty.

Fig. 8. Scatter plots indicating the correlation between the volume of activation and the average response time for each participant for both left and right superior parietal cortex. As shown, there is a linear increase in the activation volume with increasing response time.

The analysis revealed two major findings that support differential processing in left and right prefrontal cortex. The first is that while the right prefrontal region showed a significant linear increase in functional connectivity with path length, the left prefrontal region did not (see Fig. 9). The modulation by problem difficulty of the functional connectivity between the right DLPFC and left superior parietal cortex [$F(2, 20) = 6.13$, $P < 0.01$], and right superior parietal cortex [$F(2, 20) = 4.92$, $P < 0.05$] was found to be significant. On the other hand, the modulation of the functional connectivity between the left DLPFC and the left superior parietal cortex [$F(2, 20) = 1.24$, $P > 0.3$], and right superior parietal cortex ($F < 1$) failed to reach significance.

While the modulation of functional connectivity with path length was found to be significant only between the right prefrontal region and parietal cortices, the left prefrontal region was found to have a higher level of functional connectivity with the same regions. This is particularly true during the shorter path length problems, as shown in Fig. 9 (all t 's comparing left and right connectivity were > 2.0 , $P < 0.05$). This suggests that the left DLPFC was more involved in the processing of the shortest path length problems than was its right homologue.

Fig. 9. The functional connectivity between (A) the right DLPFC [(B) the left DLPFC] and the left and right superior parietal cortex. As shown, the connectivity increases with path length and the connectivity between left DLPFC and the superior parietal cortex is greater than that between the right DLPFC and the superior parietal cortex. Also, the connectivity of right DLPFC with left and right superior parietal cortex is much lower when solving the easy problems compared to the connectivity of left DLPFC with left and right superior parietal cortex when solving the easy problems.

3.1.3. 4CAPS model results

The TOL model corresponds well to the subset of the behavioral and neuroimaging data to which it has been compared, as shown in Figs. 3B, 5B, and 6B. Below we will discuss in more detail the relation between the model characteristics and the behavioral and neuroimaging data.

Fig. 3B plots the model's average processing time for problems of each length. As with the behavioral data in Fig. 3A, the model times (number of model macrocycles) increase in a monotonic and roughly linear fashion. The correlation between human reaction times and model processing times is 0.96 ($P < 0.01$). Because 4CAPS is deterministic, the TOL model cannot address the behavioral error rates without further assumptions about random processes that can cause errors. We chose not to make such assumptions here because this is not the current focus of the theory or modeling effort.

Fig. 5B plots the capacity utilization of the LH- and RH-EXECUTIVE centers for the three path length categories. The capacity utilization measure from the model is

the proportion of the resource pool that is being consumed, averaged over all the cycles of processing for each problem (see Haarmann, Just, & Carpenter, 1997 for a similar measure in a purely cognitive system, 3CAPS). The capacity utilizations of both centers increase linearly with path length, matching the neuroimaging data shown in Fig. 5. Specifically, the correlation between the brain imaging data and the capacity utilizations of the model is 0.96 ($P < 0.01$).

Fig. 6B shows the capacity utilizations of the LH- and RH-SPATIAL centers for the three levels of path length. Although the quantitative fit between the brain imaging data and model predictions is only marginally reliable, as reflected in the correlation of 0.79 ($0.05 < P < 0.1$), the qualitative fit is good.

4. fMRI discussion

The use of a graded, parametric design as well as functional connectivity analysis has allowed for a more detailed examination of the contribution of the cortical regions necessary to perform the Tower of London task. These new results combine with findings in the literature to suggest a theory of the function of the superior parietal and prefrontal cortical areas during TOL problem solving.

4.1. Planning

One of the major goals of this study was to examine the hemispheric differences in prefrontal cortex as it relates to planning. "Planning" includes strategy formation, coordinating and sequencing mental activity, and holding information on-line (Morris et al., 1997). In addition, Goel and Grafman (1995) have differentiated the process of planning ("charting a course from point A to B") versus executing that plan.

The results obtained in the current study suggest that although both left and right DLPFC are clearly involved in the TOL task, the two regions may be performing distinguishable functions. One account that is consistent with these data is that the right prefrontal cortex is differentially involved in constructing the plan for solving the TOL problem, whereas the left prefrontal cortex is involved in control processes (or supervising the execution of the plan). Several sources of evidence converge on this conclusion. First, only the right DLPFC revealed a significant correlation with reading span, such that participants with higher span scores elicited less activation than those with lower scores. The reading span has been used as a tool to assess the functioning of the central executive particularly in the context of working memory management (Engle & Oransky, 1999; Whitney et al., 2001). Therefore, the significant correlation with right but not left DLPFC provides a possible insight, namely that the right prefrontal cortex may be involved in the strategic planning necessary to integrate and maintain information. Second, the modulation of the functional connectivity between the left and right DLPFC with other areas was found to differ.

During the problems with shortest path length, the right DLPFC was less involved than the left DLPFC, suggesting that the right DLPFC becomes substantially involved only when the demands on planning increase. Therefore, while both left and right DLPFC may be involved in the planning required to perform the TOL task, it appears that they are involved in slightly different aspects of planning.

The hypothesis that left and right prefrontal regions may be associated with different functions is supported by other findings in both the neuropsychological and neuroimaging literature. For example, in a recent study examining verbal and spatial working memory, the right prefrontal cortex was found to elicit significantly more activation during the maintenance of integrated information (when letters to be remembered were displayed in the locations to be remembered) than unintegrated information (when the letters to be remembered were displayed centrally and separately from the locations to be remembered), despite the fact that the behavioral measures indicated that the integrated condition was easier than the unintegrated condition (Prabhakaran, Narayanan, Zhao, & Gabrieli, 2000). While the current task involves spatial working memory, it also involves holding onto several pieces of information, such as the spatial configuration of the balls as well as a dynamic count of the number of moves. While Prabhakaran et al. (2000) attributed the activation within the right prefrontal cortex to integration, it can be argued that strategic planning is necessary to integrate and maintain the information. Also, a recent fMRI study of goal management (a key facet of planning) found more activation in right prefrontal cortex than in left prefrontal cortex (Braver & Bongiolatti, 2002). In addition, it can be argued that the reading span task also taps into these planning processes such that individuals with high reading spans have an increased ability to integrate information, which enhances the efficiency of their working memory system.

There is also some evidence that strategy formation is lateralized to some degree, with spatial tasks revealing greater deficits after damage to the right prefrontal cortex. For example, in a study examining patients with either left or right frontal lesions during the performance of the Tower of Hanoi, patients with right hemisphere lesions, not left, were impaired in planning accuracy (Morris et al., 1997). Also, in the comparison between frontal and temporal lobe lesion patients, only the right frontal group revealed poor strategy formation (Feigenbaum et al., 1996; Miotto et al., 1996). Furthermore, in a study in which frontal lobe patients were asked to perform a multitasking experiment, it was found that those patients with lesions involving the lateral aspects of Brodmann's areas 8, 9, and 46 in the right frontal lobe made poor plans (Burgess et al., 2000), suggesting that it is the right DLPFC that is involved in strategy formation.

While the right prefrontal cortex appears to be more involved in strategy formation, the left prefrontal region appears to be more involved in the control processes necessary to provide top-down support. This is indicated not only by the significant involvement of the left DLPFC even during

the easy TOL problems, which require very little planning, but also in the region's lack of correlation with reading span, which is thought to measure central executive functioning. Also, in a single trial fMRI study of the Stroop task, activation within the left DLPFC was found to be modulated only by the instruction to name the color and not to read the word, which is consistent with the role of the left DLPFC in representing and maintaining task demands needed for top-down control (McDonald, Cohen, Stenger, & Carter, 2000).

4.2. Visuo-spatial processing

A secondary goal of this work was to examine the contribution of the superior parietal region to the spatial working memory and mental imagery processing taking place during the TOL task. Like the prefrontal regions, the left and right parietal regions also appear to be involved in somewhat different processes. As shown in Fig. 6, the activation within the right superior parietal region was not strongly modulated by path length, as was the activation in its left homologue. While the amount of right superior parietal activation was not reliably modulated by difficulty, its connectivity with right DLPFC was, suggesting that while the amount of processing taking place within the region does not change, its coordination with prefrontal regions does increase with path length. It has been hypothesized that the right superior parietal region is involved in attention processing (Coull & Frith, 1998). When performing any TOL problem, participants may be expected to make attentional shifts between the various imagined states (locations) of the balls, including the start state, the goal state, and the intermediate states. Therefore, the lack of a significant effect of difficulty in the right superior parietal region may be due to its constant role in attention, which was necessary in relatively high degrees in all three conditions.

While the right superior parietal region appears to be more involved in attention processes, the left homologue may be more of a visuo-spatial workspace. When solving the TOL problems, participants may be expected to make use of mental imagery processes in order to plan the moves. It seems plausible and has been suggested that imagery plays an important role in spatial reasoning and memory (Kosslyn, Behrman, & Jeannerod, 1995). In addition, it has been suggested that some aspects of mental imagery are lateralized to the left hemisphere (D'Esposito, Deter, Aguirre, Stallcup, Alsop, Tippet, & Farah, 1997; Farah, Peronnet, Weisberg, & Monheit, 1989; Ishai, Ungerleider, & Haxby, 2000). For example, in a study in which participants were to either view images of objects or generate mental images of objects, imagery elicited significantly more left superior parietal activation, particularly in and around the intraparietal sulcus, than its right homologue (Ishai et al., 2000). In the current study, the activation within the left superior parietal region was significantly modulated by task difficulty. Presumably, an increase in the number of moves planned would increase the amount of mental imagery processing required.

The activation of the superior parietal ROIs was also highly correlated with response time. Due to the design of the current study, response time here is analogous to “planning time” in previous studies (e.g. Rowe et al., 2001). Therefore, this high correlation suggests that although participants are not manually moving the balls during this “planning” period, they are mentally moving them and representing the intermediate steps. This high correlation of the response time with the superior parietal region activation suggests that this time is devoted to non-planning processes, as well as those related to planning. In addition, planning time was also correlated with the activation in both the left and right DLPFC, suggesting that in the current task, planning processes cannot be easily separated from those involved in visuo-spatial processes.

5. The TOL model

The account of executive function at the heart of the TOL model, that we have proposed in more detail elsewhere (Just & Varma, 2003), synthesizes two existing theories. The first was proposed by Shallice (1982), who suggested that executive function arises from the reconciliation of two streams of cognitive control. The “routine” stream is parallel, bottom-up, and perceptually driven. By contrast, the “non-routine” stream is serial, top-down, and goal-driven. Resolving the two streams of control is the job of the “contention scheduler.” Shallice placed the routine and non-routine streams in posterior and anterior areas of the brain, respectively. The second theory of executive function utilized in the present work is Newell’s (1990) Soar model of problem solving. Soar construes problem solving as a sequence of cycles, each consisting of four phases. During the “deliberation phase,” all problem-solving operators that can apply next are proposed in parallel. During the second phase, preferences between pairs of operators are computed, again in parallel. Each preference asserts that one operator is preferable to another. During the third phase, a decision procedure sorts through the preferences to select the most preferred operator. If possible, the selected operator is applied to the current state during the final phase, producing a new current state. However, if one or more of the preconditions of the selected operator are unsatisfied, then goals are established to resolve the impasse by satisfying the unsatisfied preconditions. The sufficiency of Soar as an account of problem solving has been demonstrated across several tasks (Rosenbloom, Laird, & Newell, 1993a, 1993b).

The account of executive function that we offer integrates Shallice’s theory and Soar (for more detail, see Just & Varma, 2003). The routine and non-routine streams are referred to as the *perceptual* and *strategic* modes of processing and are implemented using the information processing vocabulary of Soar: states, operators, goals, etc. The perceptual mode is driven by perceptual analysis, proposing operators that increase the visual similarity between the current state

and ending state. Perceptual operators are proposed without considering whether their preconditions are completely met, i.e. whether they can be applied to the current state. Note that the perceptual mode alone can solve simple TOL problems whose starting configurations can be transformed into their ending configurations by a sequence of perceptual moves. This is consistent with the success of patients with frontal lobe lesions on TOL problems requiring one or two moves (Owen et al., 1990).

If the preconditions of a proposed operator are unsatisfied and the perceptual mode falters, the strategic mode assumes control. The strategic mode is goal-driven in that it proposes goals that bring about the necessary preconditions. Goals are proposed and pursued by strategic operators. If one or more of the preconditions of a strategic operator are unsatisfied, goals will be recursively proposed to establish these preconditions, resulting in goal hierarchies. After the strategic mode establishes the outstanding preconditions of the moment, the perceptual mode resumes control and the original perceptual operator is applied. (Note that the perceptual and strategic modes correspond approximately to two artificial intelligence techniques, hill-climbing and hierarchical planning.) The adaptive transfer of control between perceptual and strategic modes as needed is a key characteristic of the proposed theory of executive function.

5.1. Cortical centers

The 4CAPS model of TOL problem solving relates the qualitative description of executive function offered above to cortical function. Executive function in this spatial problem-solving task is distributed across four centers, and emerges through their collaborative processing.

5.1.1. RH- and LH-EXECUTIVE centers

RH-EXECUTIVE is hypothesized to be deeply involved in strategic control. It is specialized for proposing new goals and proposing strategic moves based on existing goals. The re-iterative proposing of goals and strategic moves constitutes the construction of a plan. A plan can have the structure of a goal hierarchy if embedded goals are required. RH-EXECUTIVE becomes engaged when a selected move cannot be applied because one or more of its preconditions is unsatisfied (i.e. there is an impediment to making the move). RH-EXECUTIVE then proposes new goals that lead to satisfying each unsatisfied precondition.

LH-EXECUTIVE is hypothesized to be involved in the control of the execution of moves according to an existing plan. At any point in time, multiple moves may be available to be performed next. LH-EXECUTIVE selects among possible moves by asserting a preference between pairs of possible proposed moves (based on heuristic, algorithmic, and/or frequency principles) and choosing the next move by consolidating over these pair-wise move preferences. LH-EXECUTIVE functions like the contention scheduler of Shallice’s theory.

The capacity utilizations of both executive centers increase linearly with path length, matching the neuroimaging data reported (as shown in Fig. 5A). The LH-EXECUTIVE center displays this capacity utilization profile because harder problems require the establishment of a greater number of goals for longer durations of problem solving, often in a nested fashion, which imposes a high capacity utilization load. The RH-EXECUTIVE center shows the same pattern, but for a subtly different reason. Harder problems require more goal-driven processing. These goals generate proposed (indirect) moves that compete with the direct moves proposed by the RH-SPATIAL center. The RH-EXECUTIVE center must select between these competing proposals. The harder a problem, the more proposals to consider, and the more proposals to consider, the more resource-consuming is the selection.

5.1.2. RH- and LH-SPATIAL centers

The RH-SPATIAL center is hypothesized to be deeply involved in the perceptual control mode, at least during spatial problem solving, and thus specialized for the generation of perceptual moves. Perceptual moves are generated by comparing the current and ending configurations and proposing moves that increase the surface similarity between the two. The perceptual moves that are proposed attempt to place balls in their bin position in the ending configuration, regardless of whether the preconditions of the moves are satisfied. For example, a perceptual move could propose moving a ball into an occupied bin position.

The LH-SPATIAL center serves as the spatial workspace of the TOL model. The starting and ending configurations become available to the LH-SPATIAL center via perceptual processing that is outside the scope of the model. This center is responsible for defining the spatial characteristics of the move selected by the LH-EXECUTIVE center to produce a new configuration. This requires copying the current configuration and spatially transforming it by imagining the moving of the specified ball to a new bin position. The LH-SPATIAL center is also responsible for the maintenance of the intermediate configurations generated during problem solving. These representations implicitly record the number of moves required to solve the problem, which is the required participant response in the current administration of the TOL task.

The LH-SPATIAL center displays increasing capacity utilization with path length, resembling the increased brain activation displayed in Fig. 6A. The model's performance corresponds to the hypothesis that the LH-SPATIAL center serves as the visuo-spatial workspace, subserving the imaginal representation of various states of the problem. In particular, it encodes the starting and ending puzzle configurations as well as every intermediate configuration generated during problem solving. These encodings are necessary for the model to perform the task and be able to report the number of moves required for solution. By definition, the number of configurations to be imagined increases

with path length, and consequently so does the capacity utilization of the LH-SPATIAL center. The RH-SPATIAL center shows no systematic change in capacity utilization with trial difficulty. This is because it is specialized for activating perceptual moves during problem solving. (Recall that perceptual moves are those that directly place a ball in its ending position.) The number of direct moves does not increase with path length in TOL in this range of path lengths. Rather, as path length increases, there is an increase in the number of strategic moves driven by goals established in the EXECUTIVE centers. Thus, the RH-SPATIAL center exhibits the same lack of an effect of trial difficulty as does the right-hemisphere superior parietal cortex.

5.2. Relation to other accounts of problem solving

The theory of TOL problem solving presented above and instantiated in the 4CAPS model is both consistent with the literature and the data reported here. It is also consistent with, and to some degree synthesizes a prominent connectionist and prominent symbolic account of problem solving. The connectionist account is Dehaene and Changeux's (1997) model of TOL problem solving. Their "gesture" and "operation" levels perform the perceptually driven problem solving implemented by the SPATIAL centers of the 4CAPS model. Their "plan" level establishes goals that modulate problem solving at these two lower levels. These goals effectively seize control of problem solving away from perception. Their plan level maps neatly to the EXECUTIVE centers of the 4CAPS TOL model. The 4CAPS model is also consistent with the account of problem solving given by the Soar cognitive architecture (Newell, 1990). Soar has modeled a breadth of problem solving data using the same basic elements—problem states, moves, goals, and move selection logic—that populate the TOL model. This bodes well for the generality of the TOL model to other problem solving domains. The TOL model also differs from Soar in important ways. For example, it is not as rigorously top-down and goal-driven as Soar. On simple problems that require only perceptual moves, the SPATIAL centers perform the bulk of problem solving relatively free of executive supervision. In Soar, goals structure all problem solving from above. The present TOL model supersedes both the connectionist and Soar accounts in its ability to model both the brain activation and behavioral dimensions of cognition. This ability is inherited through the 4CAPS cognitive neuroarchitecture.

We have not modeled the relationship between individual differences and patterns of brain activation. The finding of less right DLPFC activation in higher-span participants could be modeled by assuming that higher span participants have more resources in right DLPFC. Because capacity utilization is defined as the proportion of available resources consumed by a component, this would result in lower capacity utilizations for the higher-span participants in the RH-EXECUTIVE center. However, because reading span does not measure a single processing ability (Whitney

et al., 2001), it is not clear that the observed differences in brain activation are due solely to resource differences. Also, working memory capacity differences are probably not localizable to a single brain region (Carpenter et al., 2000). Therefore, we chose not to model the individual difference results at this time.

6. Conclusions

The current study provides for a clearer characterization of the prefrontal and superior parietal cortices. With the use of a graded fMRI paradigm and functional connectivity analysis, we have shown that while both the left and right prefrontal and parietal regions are intimately involved in the TOL problem solving, the processing taking place in the two hemispheres may have differentiable characteristics. The TOL computational model is consistent with this and past data and it demonstrates the sufficiency of the proposed account.

We note that the hypothesized distinction between the left and right prefrontal processing is a preliminary one. Furthermore, the proposed specializations of the various cortical centers are proposed to be relative, not absolute, with the centers functioning in a highly collaborative style. While previous neuropsychological studies have provided some hints for such a dissociation in processing (Burgess et al., 2000; Feigenbaum et al., 1996; Miotto et al., 1996; Morris et al., 1997), the current study clearly articulates the hypothesis and provides further support for it. Further studies are necessary to determine the accuracy of the hypothesis. However, we have tried to provide an integrated account that will facilitate the evaluation and possible modification of this and other hypotheses concerning prefrontal and parietal processing in visual problem solving.

Acknowledgements

This work was supported in part by the Office of Naval Research, Grant N00014-96-1-0322, Air Force Office of Scientific Research, Grant F-49620-97-1-0368, National Institute of Mental Health Senior Scientist Awards MH-00661 and MH-00662.

Appendix A. Heuristics of the computational model

Several algorithms exist for solving Tower of Hanoi problems in the least number of moves, such as the goal recursion strategy and the perceptual strategy (Simon, 1975). In contrast, there appears to be no such algorithm for the Tower of London task. For example, empirical investigations of this task have employed operational definitions of perceptual and strategic problem solving (e.g. Ward & Allport, 1997). Perceptual moves are typically defined as those that place a ball

in its bin position in the goal configuration. All other moves are defined as strategic, and sequences of strategic moves are interpreted as implying the existence of goals. Computational models have also side-stepped postulating definitive algorithms, choosing instead to supply heuristics—either explicitly (e.g. Cooper & Waldau, 1999) or implicitly (e.g. Dehaene & Changeux, 1997)—that work on a subset of the problem space.

For these reasons, the computational model reported in this paper is heuristic rather than algorithmic. The heuristics it employs are sufficient for solving, in a minimum number of moves, the problems that participants solved in the present study (and in other experiments we are conducting). They govern the proposal of moves, selection among competing moves, and establishment of new goals in response to impasses. A number of them are described below using the following format.

- *Heuristic-name: production that implements the heuristic (center).*

In the interest of space, some heuristics are collapsed together and others omitted. A listing of the computational model, including all of its heuristics, is available from the authors upon request.

A.1. Heuristics of perceptual mode problem solving

Two heuristics are employed during perceptual mode problem solving. The first heuristic is used to propose perceptual moves.

- **Propose-perceptual-move:** If a ball is not currently in the bin position it occupies in the goal configuration, then proposing moving it there directly (RH-SPATIAL).
The second heuristic is used to select the current move among competing perceptual moves.
- **Binary-preference:** Among perceptual moves, prefer the one that makes the current configuration more visually similar to the ending configuration (as measured by the number of balls in their bin positions in the ending configuration) (LH-EXECUTIVE).

A.2. Heuristics of strategic problem solving

Three sets of heuristics are employed during strategic problem solving. The first set of heuristics establishes new goals when the current move cannot be performed.

- **Propose-unblock-ball-goal:** If the current move cannot be performed because the ball to be moved is blocked from above by another ball, then establish a new goal to unblock the ball to be moved (RH-EXECUTIVE).
- **Propose-unblock-bin-goal:** If the current move cannot be performed because the destination bin position of the ball to be moved is blocked by another ball, then establish a new goal to unblock the occupied bin position (RH-EXECUTIVE).

The second set of heuristics is used to propose strategic moves in response to the current goal.

- Propose-unblock-ball-move: If the goal is to unblock a blocked ball and a “buffer bin” exists with room, then move the blocking ball to the buffer bin (RH-EXECUTIVE).
- Propose-unblock-bin-move: If the goal is to unblock a blocked bin position and a “buffer bin” with extra room exists, then move the blocking ball to the buffer bin (RH-EXECUTIVE).
- Propose-unblocking-buffer-bin: If a “buffer bin” does not have enough room to accommodate a blocking ball, then move the top ball from the buffer bin to create room (RH-EXECUTIVE).

The final set of heuristics is used to select the current move from among the competing perceptual and strategic moves.

- Prefer-congruency-with-the-end: Among strategic moves, prefer the one that makes the current configuration more visually similar to the ending configuration (LH-EXECUTIVE).
- Prefer-reversals: Among strategic moves, prefer the one that is in the midst of “popping” the balls of one bin and “pushing” them (in reverse order) into another bin because this is a particularly efficient “macro move” (LH-EXECUTIVE).
- Buffer-preference: Prefer a strategic move making use of a “buffer bin” to a perceptual move that will lead to a configuration that will require use of a buffer bin on the next move (LH-EXECUTIVE).

The model contains additional heuristics for proposing strategic moves and for choosing among competing perceptual and strategic moves that are elided here in the interests of space.

References

- Anderson, J. R. (1993). *Rules of the mind*. Hillsdale, NJ: Erlbaum.
- Baker, S., Rogers, R. D., Owen, A. M., Frith, C. D., Dolan, R. J., Frackowiak, R. S. J., & Robbins, T. W. (1996). Neural systems engaged by planning: A PET study of the Tower of London task. *Neuropsychologia*, 34, 515–526.
- Bandettini, P. A., Wong, E. C., Hinks, R. S., Tokofsky, R. S., & Hyde, J. S. (1992). Time course EPI of human brain function during task activation. *Magnetic Resonance in Medicine*, 25, 390–397.
- Braver, T. S., & Bongiolatti, S. R. (2002). The role of frontopolar cortex in subgoal processing during working memory. *NeuroImage*, 15, 523–536.
- Buchel, C., Coull, J. T., & Friston, K. J. (1999). The predictive value of changes in effective connectivity for human learning. *Science*, 283, 1538–1540.
- Burgess, P. W., Veitch, E., de Lacy Costello, A., & Shallice, T. (2000). The cognitive and neuroanatomical correlates of multitasking. *Neuropsychologia*, 38, 848–863.
- Carpenter, P. A., Just, M. A., & Reichle, E. D. (2000). Working memory and executive function: Evidence from neuroimaging. *Current Opinion in Neurobiology*, 10, 195–199.
- Carpenter, P. A., Just, M. A., Keller, T. A., Eddy, W. F., & Thulborn, K. R. (1999). Graded functional activation in visuo-spatial system with the amount of task demand. *Journal of Cognitive Neuroscience*, 11, 9–24.
- Caviness Jr., V. S., Meyer, J., Makris, N., & Kennedy, D. N. (1996). MRI-based topographic parcellation of human neocortex: An anatomically specified method with estimate of reliability. *Journal of Cognitive Neuroscience*, 8, 566–587.
- Cooper, R., & Waldau, R. (1999). Development of problem solving strategies: Strategy shift in the Tower of London task. <http://cogent.psyc.bbk.ac.uk/publications/cooperwaldau99.html>.
- Coull, J. T., & Frith, C. D. (1998). Differential activation of right superior parietal cortex and intraparietal sulcus by spatial and nonspatial attention. *NeuroImage*, 8, 176–187.
- D’Esposito, M., Deter, J. A., Aguirre, G. K., Stallcup, M., Alsop, D. C., Tippet, L. J., & Farah, M. J. (1997). A functional MRI study of mental imagery generation. *Neuropsychologia*, 35, 715–730.
- Dagher, A., Owen, A., Boecker, H., & Brooks, D. (1999). Mapping the network for planning: A correlational PET activation study with the Tower of London. *Brain*, 122, 1973–1987.
- Daneman, M., & Carpenter, P. A. (1980). Individual differences in working memory and reading. *Journal of Verbal Learning and Verbal Behavior*, 19, 450–466.
- Dehaene, S., & Changeux, J. -P. (1997). A hierarchical neuronal network for planning behavior. *Proceedings of the National Academy Science of the United States of America*, 94, 13293–13298.
- Divadkar, V. A., Carpenter, P. A., & Just, M. A. (2000). Collaborative activity between parietal and dorsolateral prefrontal cortex in dynamic spatial working memory revealed by fMRI. *NeuroImage*, 12, 85–99.
- Eddy, W., Fitzgerald, M., Genovese, C., Mockus, A., & Noll, D. (1996). Functional imaging analysis software—computational oliv. In *Proceedings in computational statistics* (pp. 39–49). Heidelberg: Physica-Verlag.
- Engle, R. W., & Oransky, N. (1999). Multi-store versus dynamic models of temporary storage in memory. In R. J. Sternberg (Ed.), *The nature of cognition* (pp. 515–555). Cambridge, MA: MIT Press.
- Farah, M. J., Peronnet, F., Weisberg, L. L., & Monheit, M. A. (1989). Brain activity underlying mental imagery: Event-related potentials during image generation. *Journal of Cognitive Neuroscience*, 1, 302–316.
- Feigenbaum, J. D., Polkey, C. E., & Morris, R. G. (1996). Deficits in spatial working memory after unilateral temporal lobectomy in man. *Neuropsychologia*, 34, 163–176.
- Friston, K. J. (1994). Functional and effective connectivity: A synthesis. *Human Brain Mapping*, 2, 56–78.
- Goel, V., & Grafman, J. (1995). Are the frontal lobes implicated in “planning” functions? Interpreting data from the Tower of Hanoi. *Neuropsychologia*, 33, 623–642.
- Goel, V., Pullara, S. D., & Grafman, J. (2001). A computational model of frontal lobe dysfunction: Working memory and the Tower of Hanoi task. *Cognitive Science*, 25, 287–313.
- Haarmann, H., Just, M. A., & Carpenter, P. A. (1997). Aphasic sentence comprehension as a resource deficit: A computational approach. *Brain and Language*, 59, 76–120.
- Haier, R. J., Siegel, B. V., Neuchterlein, K. H., Hazlett, E., Wu, J. C., Paek, J., Browning, H. L., & Buchsbaum, M. S. (1988). Cortical glucose metabolic rate correlates of abstract reasoning and attention studied with Positron Emission Tomography. *Intelligence*, 12, 199–217.
- Horwitz, B., Rumsey, J. M., & Donohue, B. C. (1998). Functional connectivity of the angular gyrus in normal reading and dyslexia. *Proceedings of the National Academy of Sciences of the United States of America*, 95, 8939–8944.
- Ishai, A., Ungerleider, & Haxby, J. V. (2000). Distributed neural systems for the generation of visual images. *Neuron*, 28, 979–990.
- Just, M. A., & Carpenter, P. A. (1992). A capacity theory of comprehension: Individual differences in working memory. *Psychological Review*, 99, 122–149.
- Just, M. A., & Varma, S. (2003). The organization of thinking: What functional brain imaging reveals about the neuroarchitecture of cognition, manuscript in preparation.

- Just, M. A., Carpenter, P. A., & Varma, S. (1999). Computational modeling of high-level cognition and brain function. *Human Brain Mapping*, 8, 128–136.
- Just, M. A., Carpenter, P. A., & Miyake, A. (2003). Neuroindices of cognitive workload: Neuroimaging, pupillometric, and event-related potential studies of brain work [Special Issue]. *Theoretical Issues in Ergonomics*, 4, 56–88.
- Kosslyn, S. M., Behrman, M., & Jeannerod, M. (1995). The cognitive neuroscience of mental imagery. *Neuropsychologia*, 33, 1335–1344.
- Lazar, N. A., Eddy, W. F., Genovese, C. R., & Welling, J. S. (2001). Statistical issues in fMRI for brain imaging. *International Statistical Review*, 69, 105–127.
- McDonald, 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*, 288, 1835–1838.
- Mesulam, M. -M. (1990). Large-scale neurocognitive networks and distributed processing for attention, language and memory. *Annals of Neurology*, 28, 597–613.
- Mesulam, M. -M. (1998). From sensation to cognition. *Brain*, 121, 1013–1052.
- Meyer, D. E., & Kieras, D. E. (1997). *Precis to a practical unified theory of cognition and action: Some lesions from EPIC computational models of human multiple-task performance (TR-97/ONR-EPIC-08)*. Ann Arbor, MI: Department of Psychology, University of Michigan, Electrical Engineering and Computer Science Department.
- Miotto, E. C., Bullock, P., Polkey, C. E., & Morris, R. G. (1996). Spatial working memory in patients with frontal lobe lesions. *Cortex*, 32, 613–630.
- Morris, R. G., Ahmed, S., Syed, G. M., & Toone, B. K. (1993). Neural correlates of planning ability: Frontal lobe activation during the Tower of London test. *Neuropsychologia*, 31, 1367–1378.
- Morris, R. G., Miotto, E. C., Feigenbaum, J. D., Bullock, P., & Polkey, C. E. (1997). Planning ability after frontal and temporal lobe lesions in humans: The effects of selection equivocation and working memory load. *Cognitive Neuropsychology*, 14, 1007–1027.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Owen, A. M., Downes, J. J., Sahakian, B. J., Polkey, C. E., & Robbins, T. W. (1990). Planning and spatial working memory following frontal lobe lesions in man. *Neuropsychologia*, 28, 1021–1034.
- Prabhakaran, V., Narayanan, K., Zhao, Z., & Gabrieli, J. D. E. (2000). Integration of diverse information in working memory within the frontal lobe. *Nature Neuroscience*, 3, 85–90.
- Rademacher, J., Galaburda, A. M., Kennedy, D. N., Filipek, P. A., & Caviness, V. S. (1992). Human cerebral cortex: Localization, parcellation, and morphometry with magnetic resonance imaging. *The Journal of Cognitive Neuroscience*, 4, 352–374.
- Reichle, E. D., Carpenter, P. A., & Just, M. A. (2000). The neural basis of strategy and skill in sentence–picture verification. *Cognitive Psychology*, 40, 261–295.
- Rowe, J. B., Owen, A. M., Johnsrude, I. S., & Passingham, R. E. (2001). Imaging the mental components of a planning task. *Neuropsychologia*, 39, 315–327.
- Rosenbloom, P. S., Laird, J. E., & Newell, A. (1993a). *The Soarpapers: Research in integrated intelligence: Vol. 1*. Cambridge, MA: MIT Press.
- Rosenbloom, P. S., Laird, J. E., & Newell, A. (1993b). *The Soarpapers: Research in integrated intelligence: Vol. 2*. Cambridge, MA: MIT Press.
- Shallice, T. (1982). Specific impairments of planning. *Philosophical Transactions of the Royal Society of London B*, 298, 199–209.
- Simon, H. A. (1975). The functional equivalence of problem solving skills. *Cognitive Psychology*, 8, 86–97.
- Ward, G., & Allport, A. (1997). Planning and problem-solving using the five-disc Tower of London task. *The Quarterly Journal of Experimental Psychology*, 50A, 49–78.
- Whitney, P., Arnett, P. A., Driver, A., & Budd, D. (2001). Measuring central executive functioning: What's in a reading span? *Brain and Cognition*, 45, 1–14.
- Xiong, J., Rao, S., Gao, J. -H., Woldorff, M., & Fox, P. T. (1998). Evaluation of hemispheric dominance for language using functional MRI: A comparison with positron emission tomography. *Human Brain Mapping*, 6, 42–58.