

Neural mechanisms of planning: A computational analysis using event-related fMRI

Jon M. Fincham^{†*}, Cameron S. Carter[§], Vincent van Veen[¶], V. Andrew Stenger^{||}, and John R. Anderson[†]

[†]Department of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213; Departments of [¶]Neuroscience and [§]Psychiatry, Psychology, and Neuroscience, University of Pittsburgh, Pittsburgh, PA 15260; and ^{||}Department of Radiology, University of Pittsburgh Medical School, Pittsburgh, PA 15213

Contributed by John R. Anderson, December 27, 2001

To investigate the neural mechanisms of planning, we used a novel adaptation of the Tower of Hanoi (TOH) task and event-related functional MRI. Participants were trained in applying a specific strategy to an isomorph of the five-disk TOH task. After training, participants solved novel problems during event-related functional MRI. A computational cognitive model of the task was used to generate a reference time series representing the expected blood oxygen level-dependent response in brain areas involved in the manipulation and planning of goals. This time series was used as one term within a general linear modeling framework to identify brain areas in which the time course of activity varied as a function of goal-processing events. Two distinct time courses of activation were identified, one in which activation varied parametrically with goal-processing operations, and the other in which activation became pronounced only during goal-processing intensive trials. Regions showing the parametric relationship comprised a frontoparietal system and include right dorsolateral prefrontal cortex [Brodmann's area (BA 9)], bilateral parietal (BA 40/7), and bilateral premotor (BA 6) areas. Regions preferentially engaged only during goal-intensive processing include left inferior frontal gyrus (BA 44). The implications of these results for the current model, as well as for our understanding of the neural mechanisms of planning and functional specialization of the prefrontal cortex, are discussed.

Planning is ubiquitous in our daily lives: we plan our workday, our child's birthday party, the most efficient route through the grocery store, and the organization of a manuscript such as this one. At a finer time scale, planning occurs when we solve multicolumn addition problems or puzzles such as the Tower of London (TOL) or Tower of Hanoi (TOH). It is at this time scale that we investigate the neural circuitry involved in the planning process.

Since its introduction as a task to study planning from the information-processing perspective by Simon in 1975 (1), the TOH has been a prototype task in the study of high-level cognition and problem-solving behavior (2–4). TOH and TOL have been widely used in studies of patient populations as well (5–6), whereas recent neuroimaging studies have used TOL to map planning behavior onto brain activity (7–9).

Recent research has investigated the role of brain areas such as the prefrontal cortex in the service of planning. However, Goel and Grafman (10) have suggested that, whereas patient populations (5, 6) and frontal patients (11, 12) show deficits in performing planning tasks, a more precise analysis in terms of information-processing mechanisms is necessary to map functionality onto explicit cognitive mechanisms. By analyzing strategy use of their patient data in terms of formally specified information-processing models as described by Simon (1), these investigators identified the focal points of difficulty in TOH in frontal patients as deficiencies in short-term memory as well as difficulties in dealing with goal-subgoal conflicts. Similarly, recent functional neuroimaging studies have begun to investigate the neural basis of planning, with several studies showing that during variants of the TOH and TOL, frontal activation is observed (7–9). However, most studies to date have used a block

design so that the functional significance of regional brain activity observed in these studies remains uncertain.

Any study using a complex task such as the TOH to activate the brain must address the issue of strategy variability in the interpretation of data. Even a seemingly simple behavioral model that captures latency and error profiles at both the aggregate and individual levels may belie the concurrent and differential use of multiple strategies when performing a task (13, 14). This issue becomes especially relevant when we consider the interpretation of neuroimaging results. Without clarification of what strategy or multiple strategies are being used in a particular task, it is difficult to assess the functionality of neural circuitry except at a coarse task-level description.

Operationalization of Planning, Cognitive Modeling, and ACT-R. We have formalized planning behavior in the TOH as a computational model within the adaptive control of thought–rational (ACT-R) cognitive modeling framework. ACT-R (15, 16) is a general cognitive modeling architecture. It is a theoretical framework that embodies important general principles of cognition as well as a simulation system that allows us to form well-specified models of performance and learning in a rich variety of complex cognitive tasks. Because it is a simulation system run on a computer, the modeler is forced to be theoretically explicit within the framework. Given this precision, such models can run and are run to simulate data that can be compared to human performance, allowing for direct evaluation of the goodness of a model, with the additional benefit of allowing for behavioral predictions to be made regarding novel variations of the modeled task. Further, such models allow us to characterize and discuss behavior across a wide variety of tasks within a formally defined set of cognitive mechanisms as specified within the architecture.

Our current effort has its focus in ACT-R's goal-processing mechanism. Goal processing has a special status in ACT-R, with complex control of behavior accomplished through the creation and manipulation of goal structures. Planning behavior in the TOH task necessitates the processing of multiple goals and exercises this mechanism to a great degree. To control the strategy used by participants, we extended the behavioral methodology described in Anderson and Douglass (4). That methodology includes training participants in what is called a “sophisticated perceptual subgoal strategy.” This training allows us to know when participants are setting goals in the performance of the task and how many goals they are setting. It is this particular strategy that we have modeled within the ACT-R framework.

Abbreviations: fMRI, functional MRI; TOL, Tower of London; TOH, Tower of Hanoi; GOP, Grid of Pittsburgh; ACT-R, adaptive control of thought–rational; DLPFC, dorsolateral prefrontal cortex; SMA, supplementary motor area; ROI, region of interest; BA, Brodmann's area; AC-PC, anterior commissure–posterior commissure.

[†]To whom reprint requests should be addressed. E-mail: fincham@cmu.edu.

The publication costs of this article were defrayed in part by page charge payment. This article must therefore be hereby marked “advertisement” in accordance with 18 U.S.C. §1734 solely to indicate this fact.




Fig. 1. Illustrative three-disk subproblem with plans and move by problem state. In this example, the initial configuration has disks 1, 2, and 3 in column a and disks 4 and 5 in column c. The goal state is to build a tower of all five disks in column c. The intermediate states shown here are the problem states in the optimal solution path, as determined by the strategy with which participants were trained. The plan and eventual move for each transition are shown below the corresponding problem states in the solution path. The succinct description of the strategy with which subjects were trained is described as follows: (i) Select the largest out-of-place disk and the destination peg. (ii) If there is no disk blocking the move, make the move and go to step i. (iii) If the largest disk blocking the move is on the destination, select it and the other peg and go to step ii. (iv) If the largest disk blocking the move is on the source, select it and the other peg and go to step ii. For example, the first move in this problem requires three planning steps resulting in the actual move of disk 1 to column c (actual moves are indicated in boldface). Each problem solved in the scanner was designed to include two of this type of three-disk subproblems. These subproblems are the focus of all analyses reported.

We exploit our well-specified cognitive model by coupling its trial-by-trial behavioral predictions with event-related functional MRI (fMRI). Event-related fMRI provides distinct advantages over the block designs that have previously been used to study planning (17, 18). Most salient of these advantages is that event-related fMRI provides a temporal component, with the magnetic resonance signal over time being associated with psychological behavior at the individual trial level. It is this rich temporal component of the data and the ability of cognitive models to capture behavior over time that suggest an approach of combining computational cognitive modeling with event-related fMRI. By training participants in the use of a particular strategy and constructing an explicit cognitive model of that strategy, we are able to frame the time courses of activity within brain areas in terms of specific cognitive mechanisms.

Methods

Subjects. Eight right-handed English-speaking subjects (four male, four female, ages 18–32 years) enrolled in this study. Institutional Review Board approval was obtained from both Carnegie Mellon University and the University of Pittsburgh. All participants were given informed consent in accordance with Carnegie Mellon and University of Pittsburgh guidelines.

Prescan Practice Procedure. Participants were explicitly trained to use the sophisticated perceptual subgoaling algorithm. This training was accomplished by having participants solve 21 TOH problems of the sort described in ref. 4. Participants explicitly indicated when each move selection was part of the planning process with which they were instructed or an immediate move to be made. They received automated immediate feedback and were kept on the optimal solution path by being forced to indicate the correct move before proceeding. To minimize eye movements in the magnet, we next transitioned to a simpler version of TOH that we call the Grid of Pittsburgh (GOP). Fig. 1 illustrates this representation and shows a sequence of such grids, each representing a state in the solution path of an example problem. The mapping is straightforward. The three grid columns correspond to the three pegs. Digits 1–5 correspond to disk 1 (the smallest) through disk 5 (the largest). Participants practiced 21 pseudorandom GOP problems of varying complexity by using a mouse to indicate plans and moves.

For the remainder of the training procedure, participants were required to use a response glove (similar to the device available in the magnet) to indicate moves. Three buttons corresponded to the three pegs in the display. Participants had to use two

button presses to indicate a move. The first press indicated the source peg (and, by convention, its topmost disk), and the second press indicated the destination peg. During the next phase, participants were trained to perform the task almost identically to how they would in the magnet. Participants were instructed to continue using the same strategy but to compute the planning steps covertly: only actual moves were to be indicated. This method was to ensure that, in the imaging data, motor responses would not be confounded with number of goal-setting operations—one move, one response, independent of the number of goal setting operations. In addition, the pace of each move was fixed at 16 sec to prepare for the task as it is conducted within the magnet. To prevent participants from planning ahead, performance of a simple secondary task was required during the last 8 sec of each 16-sec window. As in prior training, participants were forced to stay on the optimal solution path. If they made an incorrect move, the display was updated to show the correct next state.

To allow the visual interface within the magnet to be as simple as possible, participants were next instructed to memorize a single goal state. This configuration would serve as the goal state toward which they would always work during the scanning session of the study. Having them do this allowed us to display only the current state of a problem with the assumption that they would always work toward the single memorized goal state. To assist in this memorization, participants were trained by using 10 simple practice problems with initial states that varied between one and five moves from the target goal state.

Scanner Behavioral Procedure. The behavioral task in the scanner consisted of 12 blocks of ≈6 min in length. During each block, participants solved a unique TOH problem defined by a pseudorandom initial state and the goal state memorized during training. The problems were of average difficulty, requiring between 19 and 23 moves in the optimal solution path. Problems were constructed so that two seven-move sequences of the type shown in Fig. 1 were embedded in the middle of the overall solution path. Only the current state was displayed to the participant, and it subtended less than 2° of visual angle. Responses were collected through the response glove, as described in prescan practice. As during prescan training, participants were kept on the correct solution path. Incorrect trials and their immediate successors were ignored in subsequent analyses.

fMRI. Images were acquired with a conventional 1.5-T GE Signa whole-body scanner (General Electric) and a standard radio

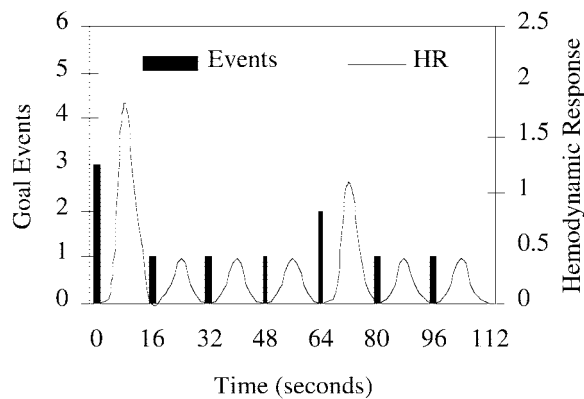


Fig. 2. Convolution of goal-setting events as predicted by the ACT-R model yields an unscaled hemodynamic response. This model serves as the reference time series for the regression model discussed in the text. The function shown here is generated from solving a three-disk subproblem common to all problems.

frequency head coil. Twenty-five oblique axial slices (3.80 mm thick, 3.75 mm² in-plane resolution) were acquired parallel to the anterior commissure–posterior commissure (AC-PC) line, with the middle of the seventh slice from the bottom through the AC-PC line (a total of 95 mm of brain coverage). Functional scans were acquired by using a two-shot T2*-weighted spiral scan pulse sequence [repetition time = 2,000 ms, echo time = 34 ms, FOV = 24 cm, flip angle = 70°] (19). Scanning was event-related, with image acquisition synchronized to stimulus onset, such that four volumes, each containing 25 slices, were acquired during each 16-sec trial.

Anatomical scans (36 slices) were acquired by using a standard T1-weighted pulse sequence, with the middle of the 15th slice from the bottom through the AC-PC line. Data from individual subjects were subjected to a between-block subtractive mean normalization. Images were then coregistered to a common reference structural MRI scan by means of a 12-parameter automatic algorithm AIR (20) and smoothed with an 8-mm full-width half-maximum three-dimensional Gaussian filter to accommodate individual differences in anatomy.

Analysis. Data analysis was driven by the underlying ACT-R model of the task. The following equation describes the voxel-wise regression model used in the analysis: $MR(t) = B_0 + B_1 \cdot \text{trial}(t) + B_2 \cdot \text{ACT-R}(t) + e(t)$ (MR , magnetic resonance). The signal for a particular voxel over time, $MR(t)$, is modeled as a linear combination of constant intercept and two reference functions. The reference function $\text{trial}(t)$ is an unscaled hemodynamic response that is identical for each trial in the experiment. This term in the model is designed to capture variation in the MR signal due to generic trial processes such as encoding and response generation, independent of what we consider to be the cognitive component of the task, captured by the term $\text{ACT-R}(t)$. In the ACT-R model of the TOH task, planning is carried out through varying numbers of goal-setting operations. To generate the $\text{ACT-R}(t)$ time series, we begin by using a model to simulate solutions for each problem encountered in the scanner. The goal-setting events and their timestamps are identified and collected. To generate the reference time series, these discrete goal-setting events are convolved to generate an unscaled hemodynamic response function (18, 21). Fig. 2 shows the discrete events and their convolution for the seven-move subproblem like the sequence shown in Fig. 1 that is common to all problems. For each of these seven-move sequences, the first move has high planning (three goals set), the fifth move medium planning (two goals set), and the rest low planning (one goal set). Spatial t maps




Fig. 3. Behavioral latency profile of three-disk subproblems that are the focus of analyses.

were generated through a between-subject voxel-wise analysis of the regressor coefficient B_2 of the ACT-R-generated time series. A voxel-clustering threshold of 8 was used to increase control for multiple comparisons (22). Only voxels significant at the $P < 0.0005$ level are presented.

Behavioral Results. All analyses are constrained to those data corresponding to moves made during three-disk subproblems of the type shown in Fig. 1. Fig. 3 shows the aggregate latency profile for correct trials over the move sequence for these subproblems. These seven moves served as the levels within a single-factor repeated-measures ANOVA. Latencies were significantly different for moves in these three-disk subproblems, $F(6,42) = 48.7$, $P < 0.0001$, $MSE = 101026$. These results are consistent with the ACT-R model of the task. As can be seen in Fig. 3, the first high-planning move in the sequence takes the longest time. The medium-planning move 5 takes the next longest. The remaining low-planning moves show the fastest latencies, reflecting the fact that each requires only a single planning step. Mean accuracy across moves was 95%. A repeated measures ANOVA of the move factor for accuracy was marginally significant, $F(6,42) = 3.07$, $P < 0.05$, $MSE = 0.003$, with most errors occurring during the high-planning state (89% accuracy). All error trials and their immediate successors were eliminated from subsequent analyses.

fMRI Results. Regions identified as differentially responsive to goal-processing operations are distributed among prefrontal cortex, parietal cortex, cingulate gyrus, and subcortical structures (Fig. 4, Table 1). Frontal regions include right dorsolateral prefrontal cortex (DLPFC) [Brodmann's area (BA 9)], left inferior frontal gyrus (BA 44), bilateral premotor cortex (BA 6), left supplementary motor area (SMA) (BA 6), and right SMA/preSMA (BA 8). Parietal regions include a large bilateral region (BA 7/40) that encompasses superior parietal cortex, precuneus (BA7), inferior parietal lobules (BA 40), and angular gyri (BA 39). Other parietal regions include bilateral cuneus (BA 18) and left precuneus (BA 31). The cingulate area identified is in the rostral anterior cingulate cortex (BA 24). Subcortical structures include right caudate nucleus and thalamus.

Examination of the time courses of activity suggests two distinct patterns of activation, one in which the blood oxygenation level-dependent (BOLD) response varied parametrically as a function of goal-processing operations as predicted by model, and another in which activity substantially increased only during the high-planning trials. Region of interest (ROI) time courses are baseline normalized to the mean activation of the ROI over the course of the experiment for each subject. Within the context of a repeated-measures ANOVA (Time $4 \times$ Planning




Fig. 4. Activation map overlay on female reference brain. Regions with eight or more contiguous voxels significant at level $P < 0.0005$ are shown. Slices are shown in radiological space (image left is brain right; image right is brain left) and are ordered from top of brain to bottom. Slice 19 is AC-PC line. See also Table 1.

Load 3), we performed a single degree of freedom contrast for each ROI, testing whether time points 2 and 3 of the medium-planning time course fall midway between the same points in the time courses of the low- and high-planning states. The significance threshold is set to $P = 0.01$; contrasts with a $P < 0.01$ indicate a rejection of the null hypothesis that there is a parametric relationship between planning load and the BOLD response. Time courses that show a clear parametric relation with planning load include those of left SMA (BA 6), $F(1,42) = 1.38$, $MSE = 0.0054$; bilateral premotor cortex (BA 6), $F(1,42) = 0.135$, $MSE = 0.0016$; bilateral parietal cortex (BA 7/40), $F(1,42) = 2.86$, $MSE = 0.0045$; right cuneus (BA 18), $F(1,42) = 0.426$, $MSE = 0.0038$; and bilateral anterior cingulate (BA 24), $F(1,42) = 0.0788$, $MSE = 0.0078$. Time courses that do not suggest a parametric relationship with planning load include left inferior frontal gyrus (BA 44), $F(1,42) = 7.345$, $MSE = 0.0022$; left precuneus (BA 31), $F(1,42) = 13.044$, $MSE = 0.0083$; and bilateral thalamus, $F(1,42) = 9.3210$, $MSE = 0.0049$. Several time courses yielded marginally significant contrasts ($P < 0.05$), suggesting an ambiguity about the parametric relationship. These areas include right SMA/preSMA (BA 8), $F(1,42) =$

4.492, $MSE = 0.0040$; right DLPFC (BA 9), $F(1,42) = 5.804$, $MSE = 0.0046$; left cuneus (BA 18), $F(1,42) = 5.537$, $MSE = 0.0040$; and caudate nucleus, $F(1,42) = 5.145$, $MSE = 0.0013$. We consider a subset of our ROIs displaying activity patterns from each category in turn in the following discussion.

Discussion

The results of our analysis of the neural substrates of planning, using GOP and a computational model of task performance, confirm the involvement of a distributed neural network that includes frontal, parietal, and subcortical elements. Within this network, an analysis of the time course of activation and its relationship to planning suggests a degree of modularity as subgoals are constructed and maintained in the service of correct task performance.

Parametric Relationship with Planning Load and the Frontoparietal Network. Fig. 5 shows the time courses displaying the parametric relationship with planning load for the frontoparietal system, including right dorsolateral prefrontal (BA 9), bilateral parietal (BA 7/40), and bilateral premotor (BA 6) regions. The right

Table 1. Foci of activation

Anatomical region	BA	Voxel count	Stereotaxic coordinates, mm			Maximum t (average t)
			x	y	z	
Left medial frontal gyrus (SMA)	6	30	-1	-2	55	9.68 (7.41)
Right medial frontal gyrus (SMA/pre-SMA)	8	8	7	18	43	8.35 (7.28)
Right precentral sulcus (premotor)	6	93	33	-8	51	13.01 (8.03)
Left precentral sulcus (premotor)	6	25	-22	-19	51	8.75 (7.25)
Right superior frontal gyrus (DLPFC)	9	8	45	37	32	8.98 (7.54)
Left inferior frontal gyrus	44	19	-42	5	27	16.71 (8.95)
Left/right parietal	7	233	-7	-59	51	23.05 (8.07)
Left/right precuneus	7					
Left/right angular gyrus	39					
Left/right inferior parietal lobule	40					
Left precuneus	31	8	-3	-71	27	7.37 (6.96)
Right cuneus	18	16	22	-67	16	9.64 (7.40)
Left cuneus	18	18	-22	-77	20	11.80 (7.94)
Right caudate nucleus	23	23	18	5	27	9.58 (7.80)
Left/right thalamus		169	1	-19	12	15.21 (7.96)
Left/right anterior cingulate	24	21	-7	33	12	9.20 (7.66)

Activation foci sensitive to number of planning steps as identified by between-subject voxel-wise t tests of the ACT-R(t) model coefficient in the linear model of the MR signal (see *Analysis*). Only regions with eight or more contiguous active voxels (significant at level $P < 0.0005$) are reported. Coordinates refer to location of maximal activation in Talairach space.

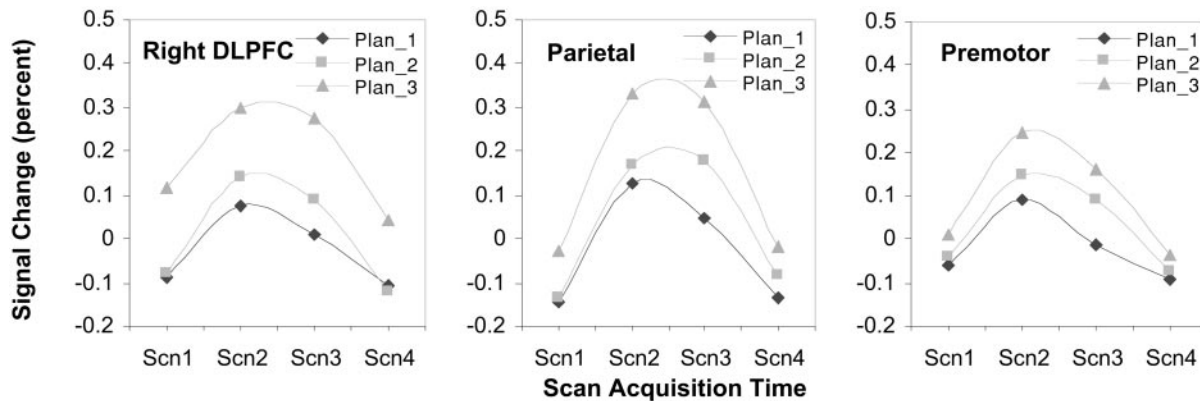


Fig. 5. Right DLPFC, premotor, and parietal time series. Activity increases parametrically with number of planning steps. Percent signal change is relative to mean signal strength over the course of the experiment.

DLPFC particle is somewhat ambiguous. Although the contrast testing whether the time courses vary strictly parametrically indicated marginally significant deviation, a pairwise t test of the difference between means of the middle time points of medium- and low-planning conditions is significant, $t(7) = 3.42$, $P = 0.005$ one-tailed, indicating that this region is differentially sensitive to planning load. Although the ACT-R time series used to identify these regions is generated by using only goal-processing events, examination of the model reveals that these events are tightly coupled with requirements for active short-term memory maintenance when formulating a plan for action. We propose that at a coarse level, these regions operate in concert and essentially serve to warehouse intended moves and resulting problem states associated with the goals.

There are rich interconnections between parietal, premotor, and dorsolateral prefrontal areas (23, 24). Numerous neuroimaging studies of spatial working memory and problem solving have shown coactivation of prefrontal and parietal areas (see ref. 25 for a review), and recent fMRI work by Diwadkar, Carpenter, and Just (26) has demonstrated a clear relationship between activity in DLPFC and parietal areas. Dorsolateral prefrontal cortex plays a crucial role in the active maintenance of information and is differentially responsive to working memory load (27). Right DLPFC, in particular, has been found active during a variety of spatial working memory tasks (28–31). Further, coactivation of right DLPFC and parietal cortex has also been demonstrated in spatial working memory tasks (32), as has coactivation of parietal and premotor areas (25).

Differentially Greater Activity During High-Planning Load. A number of areas are identified in our analysis that did not show a strictly parametric relationship between planning load and activation, instead showing pronounced activity only during the high-planning load condition. We limit our discussion to the left inferior frontal gyrus (BA 44). Fig. 6 shows the time course of activity for this area.

We see that, in contrast to the right DLPFC region, the left inferior frontal gyrus shows equal responsiveness for low- and medium-planning loads and a differentially greater response for high-planning load. The general point to be taken here is that, whereas all regions identified in the current analysis are responsive to planning load, all areas are not identically responsive. Rather, the results suggest that different brain regions are responding differently to separate cognitive demands in the task. With respect to the differences in the prefrontal regions, participants report not having to work out their plans in the medium-planning state but rather retrieving them. The plan in these states always involves moving the smallest disk out of the

way so that the next smallest can be moved. Retrieval of these plans is quite plausible, because participants were highly trained in the task and so had many experiences with this particular configuration. This fact suggests a functional dissociation between these prefrontal regions with the left inferior frontal gyrus especially active when building subgoals in the construction of a plan. This view is consistent with previous studies suggesting a general selection function for this region of the brain (33). Further, we noted earlier that the responsiveness of the right DLPFC region during the medium-planning state was slightly lower than expected with respect to the current model. To the extent that this region subserves both goal-processing and maintenance functions, the depressed time course for the medium-planning condition might reflect that fewer subgoals are created when using this strategy.

Although the parametric responsiveness of parietal, premotor, and right DLPFC regions is consistent with data reported in Baker *et al.* (8) in their direct contrast of easy and hard TOL problems, these results are not consistent with those reported in Dagher *et al.* (7). Their TOL data show that activity in lateral premotor and DLPFC areas varies parametrically with problem complexity but does not show such a relationship in parietal areas. We attribute this difference to issues inherent in the blocked design methodology (17) and to task and strategy differences between their TOL task and the GOP task.

One prefrontal cortical region previously associated with planning, anterior, or polar frontal cortex (BA 10) was not active in our

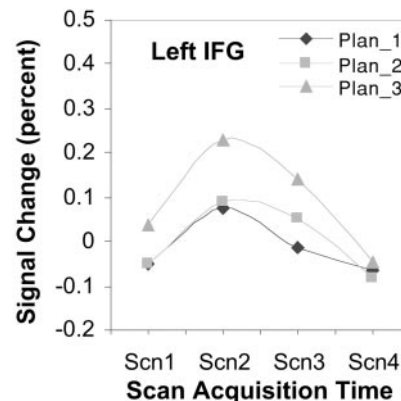


Fig. 6. Left inferior frontal gyrus time series. Activity most prevalent in high-planning states while equivalent between low and medium-planning loads. Percent signal change is relative to mean signal strength over the course of the experiment.

analysis. Current hypotheses regarding the contribution of this region to higher cognitive functions include processes related to the evaluation of self-generated representations (34) and the support of cognitive branching (35), defined as the maintenance of superordinate goals during subgoal processing. Operations such as these would be engaged during all levels of planning in the present paradigm, which may account for a lack of differential activation with the number of planning steps in our present analysis.

Conclusion

Our goals in this study were 2-fold. The first and most specific was to identify the brain circuitry that is differentially responsive to goal-processing operations in the GOP task. We have shown that the frontoparietal network and other areas are differentially activated with unique time course profiles in response to task demand, characterizing that demand in terms of goal-processing operations and spatial working-memory requirements of the task. The second and more general goal was to show that a computational cognitive architecture can be used as a framework through which we can precisely articulate the cognitive components of a task and so improve our ability to assess and describe the functionality of particular brain areas and networks of areas in terms of cognitive mechanisms. The results of our

analysis suggest a functional dissociation between a right prefrontal-parietal network whose activity reflects the working memory demands of the task and a left posterior DLPFC component associated with the active selection of task-appropriate subgoals in the service of planning. Also suggested is an extension of the ACT-R model in which multidisk configurational cues are taken into account when constructing a plan for action.

We have focused on the goal-processing mechanism of ACT-R in the current task at a still fairly coarse level. At this level, it remains challenging to isolate goal-processing proper activity from other task-related activity. However, the architecture contains many other features that we have not considered here, including retrieval mechanisms, perceptual-motor mechanisms, and subsymbolic computations. Our goal is to continue using ACT-R as a framework through which further iterations of mapping features of the theory onto neural substrates will serve to both inform the models and architecture and serve as a tool from which we can frame results across other neuroimaging studies in terms of specific cognitive mechanisms.

This work was supported by National Science Foundation Grant SBR-9873465 and Grant 5-T32-MH19983 from the National Institute of Mental Health.

1. Simon, H. A. (1975) *Cognit. Psychol.* **7**, 268–288.
2. Simon, H. A. & Hayes, J. R. (1976) *Cognit. Psychol.* **8**, 165–190.
3. Kotovsky, K., Hayes, J. R. & Simon, H. A. (1985) *Cognit. Psychol.* **17**, 248–294.
4. Anderson, J. R. & Douglass, S. (2002) *J. Exp. Psychol.* **27**, 1331–1346.
5. Shallice, T. (1982) *Philos. Trans. R. Soc. London B* **298**, 199–209.
6. Owens, A. M., Downs, J. J., Shahakian, B. J., Polkly, C. E. & Robbins, T. W. (1990) *Neuropsychologia* **28**, 1021–1034.
7. Dagher, A., Owen, A. M., Boecker, H. & Brooks, D. J. (1999) *Brain* **122**, 1973–1987.
8. Baker, S. C., Rogers, R. D., Owen, A. M., Frith, C. D., Dolan, R. J., Frackowiak, R. S. & Robbins, T. W. (1996) *Neuropsychologia* **34**, 515–526.
9. Owen, A. M., Doyon, J., Petrides, M. & Evans, A. C. (1996) *Eur. J. Neurosci.* **8**, 353–364.
10. Goel, V. & Grafman, J. (1995) *Neuropsychologia* **33**, 623–642.
11. Shallice, T. (1988) *Neuropsychology to Mental Structure* (Cambridge Univ. Press, Cambridge, U.K.).
12. Grafman, J. (1994) in *Handbook of Neuropsychology*, eds. Boller, F. & Grafman, J. (Elsevier, Amsterdam), Vol. 9, pp. 187–202.
13. Siegler, R. (1987) *J. Exp. Psychol.* **116**, 250–264.
14. Fincham, J. M. & Anderson, J. R. (1999) in *Proceedings of the 21st Annual Conference of the Cognitive Science Society* (Erlbaum, Mahwah, NJ), pp. 167–172.
15. Anderson, J. R. & Lebiere, C. (1998) *The Atomic Components of Thought* (Erlbaum, Mahwah, NJ).
16. Anderson, J. R. (1990) *Rules of the Mind* (Erlbaum, Mahwah, NJ).
17. D'Esposito, M., Zarahn, E. & Aguirre, G. K. (1999) *Psychol. Bull.* **125**, 155–164.
18. Dale, A. M. & Buckner, R. L. (1997) *Hum. Brain Mapp.* **5**, 329–340.
19. Noll, D. C., Cohen, J. D., Meyer, C. H. & Schneider, W. (1995) *J. Magn. Reson. Imaging* **5**, 49–56.
20. Woods, R. P., Cherry, S. R. & Mazziotta, J. C. (1992) *J. Comput. Assist. Tomogr.* **16**, 620–633.
21. Boynton, G. M., Engel, S. A., Glover, G. H. & Heeger, D. J. (1996) *J. Neurosci.* **16**, 4207–4221.
22. Forman, S. D., Cohen, J. D., Fitzgerald, M., Eddy, W. F., Mintun, M. A. & Noll, D. C. (1995) *Magn. Reson. Med.* **33**, 636–647.
23. Goldman-Rakic, P. S. (1987) in *Handbook of Physiology*, eds. Mountcastle, V. B., Plum, F. & Geiger, S. R. (Am. Physiol. Soc., Bethesda, MD).
24. Petrides, M. & Pandya, D. N. (1984) *J. Comp. Neurol.* **228**, 105–116.
25. Cabeza, R. & Nyberg, L. (2000) *J. Cognit. Neurosci.* **12**, 1–47.
26. Diwadkar, V. A., Carpenter, P. A. & Just, M. A. (2000) *NeuroImage* **12**, 85–99.
27. Braver, T. S., Cohen, J. D., Jonides, J., Smith, E. E. & Noll, D. C. (1997) *NeuroImage* **5**, 49–62.
28. Jonides, J., Smith, E. E., Koeppel, R. A., Awh, E., Minoshima, S. & Minton, M. A. (1993) *Nature (London)* **363**, 623–625.
29. Courtney, S. M., Petit, L., Maisog, J. M., Ungerleider, L. G. & Haxby, J. V. (1998) *Science* **279**, 1347–1351.
30. McCarthy, G., Puce, A., Constable, R. T., Krystal, J. H., Gore, J. C. & Goldman-Rakic, P. (1996) *Cereb. Cortex* **6**, 600–611.
31. McCarthy, G., Blamire, A. M., Puce, A., Nobre, A. C., Bloch, G., Hyder, F. & Goldman-Rakic, P. (1996) *Proc. Natl. Acad. Sci. USA* **91**, 8690–8694.
32. Smith, E. E., Jonides, J., Koeppel, R. A., Awh, E., Schumacher, E. H. & Minoshima, S. (1995) *J. Cognit. Neurosci.* **7**, 337–356.
33. Thompson-Schill, S. L., Swick, D., Farah, M. J., D'Esposito, M., Kan, I. P. & Knight, R. T. (1998) *Proc. Natl. Acad. Sci. USA* **95**, 15855–15860.
34. Christoff, K. & Gabrieli, J. D. E. (2000) *Psychonom. Soc.* **28**, 168–186.
35. Koechlin, E., Basso, G., Pietrini, P., Panzer, S. & Grafman, J. (1999) *Nature (London)* **399**, 148–151.