Journal of Experimental Psychology: General
2001, Vol. 130, No. 3, 493–504

0096-3445/01/$5.00 DOI: 10.1037//0096-3445.130.3.493

Mental Rotation of Objects Retrieved From Memory: A Functional MRI Study of Spatial Processing

Marcel Adam Just, Patricia A. Carpenter, Mandy Maguire, Vaibhav Diwadkar, and Stephanie McMains
Carnegie Mellon University

This functional MRI study examined how people mentally rotate a 3-dimensional object (an alarm clock) that is retrieved from memory and rotated according to a sequence of auditory instructions. We manipulated the geometric properties of the rotation, such as having successive rotation steps around a single axis versus alternating between 2 axes. The latter condition produced much more activation in several areas. Also, the activation in several areas increased with the number of rotation steps. During successive rotations around a single axis, the activation was similar for rotations in the picture plane and rotations in depth. The parietal (but not extrastriate) activation was similar to mental rotation of a visually presented object. The findings indicate that a large-scale cortical network computes different types of spatial information by dynamically drawing on each of its components to a differential, situation-specific degree.

Functional neuroimaging provides a new opportunity to examine some of the fundamental characteristics of spatial thinking. Spatial thinking ability has long been studied as a major component of human intelligence, ranging from Binet's early measurements of intelligence to contemporary brain imaging studies. One of the fascinations of spatial thinking is its potential correspondence with objects and transformations in the physical world. For example, in mental rotation tasks, which require that an object be represented at a new orientation, questions arise about the three-dimensionality of the object representations and their transformations. This study examined how various types of computational demands affect the activation in a large-scale network of cortical regions that are involved in mental rotation.

Research on spatial thinking has attempted to specify the uniquely spatial representations and processes involved, focusing on their geometric properties. For example, one question about the mental rotation process has concerned the effect of the axis of rotation (Shepard & Metzler, 1971), specifically whether rotation in the picture plane is similar to rotation in depth, as it is in the physical world. If the mental representation of the object were only two-dimensional, in the plane of depiction, then the two types of rotation should be different, with the rotation in depth requiring much more processing. The Shepard and Metzler study compared rotation around the z-axis (the axis parallel to the imagined line of sight, such that the object was imagined to rotate in the picture plane) with rotation around the y-axis (a vertical axis around which the object was imagined to be rotating in depth). The behavioral measures generally showed comparable results for the two types of rotation, suggesting that the object representations maintained a three-dimensional structure rather than just a two-dimensional representation that corresponds to the picture plane. The current study also compared various types of rotations, differing in the rotation axis, and sometimes requiring that the object be rotated first around one axis and then around another. The general objective was to discover how various types of geometric properties, such as the uniformity of the rotation axis, affect brain activation in a spatial thinking task such as mental rotation.

A second question about mental rotation has concerned the effect of the input modality (e.g., visual vs. haptic). The modality issues have been investigated in behavioral studies that showed that any of several input modalities could engender mental rotation, such that response times increased monotonically with the size of the rotation angle (Carpenter & Eisenberg, 1978); however, the input modalities greatly affected the parameters of the behavioral response. If the underlying spatial processes are fundamentally geometric, then they might engender the same brain activation regardless of input modality.

A major innovation of the current study is that all of the image generation and transformation processes occurred without visual input. The participants recalled from memory the object to be rotated, an alarm clock with which they had become familiar. The experimental task was adapted from a venerable test of spatial ability, the Guilford–Zimmerman Spatial Visualization Test (Guilford & Zimmerman, 1947–56). In the original test, participants are shown a drawing of an old-fashioned alarm clock and then given visual directions for how to rotate it through a sequence of steps. In the current study, we first familiarized the participants with an old-fashioned alarm clock, with a clock face on the front, various knobs and buttons on the back, and a base at the bottom, as shown in its canonical position in Figure 1. During our task, participants were expected to retrieve their representation of the clock from memory, without visual input, in its canonical position, namely upright and with the clock face facing the viewer. This canonical

---

Marcel Adam Just, Patricia A. Carpenter, Mandy Maguire, Vaibhav Diwadkar, and Stephanie McMains, Center for Cognitive Brain Imaging, Carnegie Mellon University.

This work was supported in part by Office of Naval Research Grant N00014-96-1-0322 and National Institute of Mental Health Senior Scientist Awards MH-00661 and MH-00662.

Correspondence concerning this article should be addressed to Marcel Adam Just, Center for Cognitive Brain Imaging, Psychology Department, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213. Electronic mail may be sent to just+@cmu.edu.

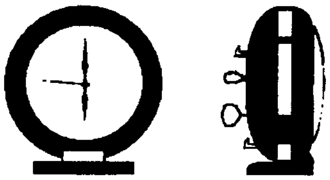

Figure 1. View of the clock in its canonical, initial position (on left) and how its orientation should look after mental rotation following the instruction "right."

position was the initial position to which the first step of any sequence of orientation transformations was to be applied. Another adaptation was that the rotations to be performed were communicated auditorily. For example, participants might be given the spoken instruction "counter," conveying that the most recently represented orientation of the clock was to be rotated counterclockwise (i.e., in the picture plane). All transformations were to be 90°, by convention. This paradigm made it possible to compare the brain activation with other studies in which the object to be mentally rotated is visually presented.

Functional Neuroimaging of Mental Rotation

Spatial thinking not only has distinctive processing characteristics in terms of the underlying geometry but is also related to several particular parts of the brain. Functional MRI (fMRI) makes it possible to address questions about spatial thinking by measuring the activation in various cortical areas while participants mentally rotate an object. The relative amount of activation in each area of a large-scale cortical network indicates not only the components of the underlying neural system but also the modulation of their relative contributions as the computational demands of the task are varied from condition to condition. For example, one can assess the effect of alternating the rotation axis on the brain activation in several areas.

The two cortical regions that might be expected to play central roles in mental rotation are the parietal and inferior temporal areas (i.e., the dorsal and ventral systems, respectively). Several neuroimaging studies indicate that the parietal areas are implicated in mental rotation, playing a large role in the geometric transformation of the represented object orientation. In a study involving the Shepard–Metzler rotation task, fMRI-measured activation was found in the left and right parietal regions (Brodmann Areas [BA] 7a and 7b and sometimes BA 40) when a rotation condition was compared with a 0° rotation condition (Cohen et al., 1996). A PET study found activation in the left posterior–superior parietal cortex during the mental rotation of letters (Alivisatos & Petrides, 1997). In two separate parametric studies involving the mental rotation of Shepard–Metzler figures, the amount of activation in the left and right intraparietal sulcal region increased linearly with rotation angle (measured in terms of the number of voxels activated above a high fixed threshold), suggesting the amount of work that this area did increased with the rotation angle (Carpenter, Just, Keller, Eddy, & Thulborn, 1999a). These data suggest that the amount of activation in the parietal areas that support visuospatial processing is related to the amount of computational demand.

Parietal involvement has also been found in several event-related potential (ERP) studies of a simpler rotation task in which participants judged whether a letter was normal or mirror imaged (Desrocher, Smith, & Taylor, 1995; Peronnet & Farah, 1989; Wijers, Otten, Feenstra, Mulder, & Mulder, 1989). All three ERP studies found that as the angular disparity increased, there was increasing negativity in the ERP waveforms in the latency range of 350–800 ms, particularly in the parietal and occipital leads, which was interpreted as reflecting rotation, and lesser but still systematic effects in central leads (Desrocher et al., 1995; Rosler, Heil, Bajric, Pauls, & Henninghausen, 1995). These studies all implicate the parietal region in some aspect of rotation of an object that is visually presented.

However, there is more to a mental rotation task than just rotation. In detailed computational models of eye-fixation studies of mental rotation, Just and Carpenter (1985) proposed that a mental rotation of a complex object that is visually presented entails figural segmentation and rotation of only one segment of a complex object at a time. Classically, the inferior temporal area is associated with visual object identification or object-part identification, so one might expect that area to play a large role in identifying salient segments of objects to be rotated and evaluated. In their fMRI study, Carpenter et al. (1999a) found that not only the parietal but also the inferior temporal areas became activated during mental rotation, regions that are primarily associated with the processes of object and object-part identification. However, in the absence of visual input in the clock task, one might expect the inferior temporal contribution to a mental rotation task to decrease. More generally, the comparison between the two paradigms permits an evaluation of the roles of the dorsal and ventral systems in various types of mental rotation.

Finally, in a task as complex as clock rotation, the executive and verbal components of the task should engage appropriate cortical areas. The executive components would likely be involved in goal management, namely developing goals and subgoals and keeping track of which ones have been satisfied, resulting in activation in dorsolateral prefrontal cortex. The verbal components might encode and rehearse representations of either the state of a clock representation or an action or goal to be pursued, resulting in activation in inferior frontal areas, posterior superior temporal areas, and inferior parietal areas, particularly in the left hemisphere.

The main goal of the study was to determine how various aspects of mental rotation affected the pattern of activation in several cortical regions, but there were three specific hypotheses. One hypothesis concerned alternation between two axes of rotation. The hypothesis was that a change in rotation axis from one rotation step to the next should be more computationally demanding than making successive rotations around the same axis; hence, there should be more brain activation when the rotations alternated axes. This hypothesis is based on a proposed model of the process, described further in the Discussion section, which assumes that the mental rotation of the clock is performed by reducing an initial three-dimensional representation to a two-dimensional representation and then applying the orientation transformation. In the case of alternating axes, a different two-dimensional representation would have to be generated after each rotation step, namely the representation in the plane that is perpendicular to the designated rotation axis, requiring extra computational and representational work relative to the same-axis rotations.

A second hypothesis was that as the number of successive rotation steps to be applied to a clock increased, the amount of brain activation should correspondingly increase. This hypothesis is based on a rotation workload hypothesis, one that has previously been supported in the fMRI study of Carpenter et al. (1999a), in which the rotation angle was systematically varied.

A third main interest was the effect of input modality, comparing the pattern of cortical activation in mental rotation of retrieved images with previous studies in which the objects to be rotated were visually presented. In addition to the comparison with previous studies, we did a within-subject comparison for three participants who also performed a mental-rotation task involving visually presented three-dimensional objects in the same testing session. The third hypothesis was that similar areas of the parietal lobule would activate in the two types of studies, but there should be dissimilarities in more sensory and perceptual areas.

Method

Participants

The participants contributing fMRI data were 12 right-handed volunteer college students (7 men and 5 women). Each participant gave signed, informed consent (approved by the University of Pittsburgh and the Carnegie Mellon Institutional Review Boards). Participants were first familiarized with an abbreviated behavioral version of the mental rotation task, and a larger group was screened for acceptable accuracy scores. Only the 12 participants who made less than 40% error responses in the most difficult condition and had fMRI data sufficiently free of head motion artifacts were entered into the analyses.

fMRI Parameters

Images were acquired on a GE Medical Systems 3.0 Tesla scanner using the following parameters: TR (repetition time) = 3,000 ms, TE (echo time) = 25 ms, flip angle = 90°, and voxel size = 3.125 × 3.125 × 5 mm with a 1-mm gap. For 10 of the 12 participants, 14 functional slice images were acquired in an oblique-axial plane and were positioned to cover as much of the parietal and temporal lobules as possible; the pitch angle of the images ranged from 9° to 21°. For 2 of the 12 participants, the 14 slice images were acquired in a coronal plane and were positioned to cover as much of the parietal and temporal lobes as possible. The latter slice prescription is limiting in that it does not cover the frontal areas. These two subjects are included in analyses concerning the parietal and temporal regions of interest (ROIs) but not in dorsolateral prefrontal cortex, frontal eye fields, or the inferior frontal gyrus.

Design

The experiment had two independent variables: the axis about which the clock was to be rotated (z-axis only, y-axis only, or z- and y-axes) and the number of successive 90° rotation steps that were to be applied to the clock. The latter variable, the number of rotation steps, is referred to as the path length variable, and it was either a short path (either two or three rotation steps) or a long path (four rotation steps). The three levels of the rotation axis variable and the two levels of the path length variable combined to produce six conditions or epoch types.

Procedure

To determine how much time to allow for performing various rotations, we ran a pilot study in which the time intervals were under the participants' control. Rotation commands were presented, and the participants pressed a button when they had completed the instructed rotation. The durations chosen for the main experiment corresponded to a point that included about 75% or more of the self-paced response times in the pilot study. These durations were 3.2 s for the first rotation in a problem, 2.2 s for any subsequent rotations in a single-axis problem, and 5.7 s for any subsequent rotations in a z and y axis problem.

The presentation of each of the various types of problems always started with a visual "ready" text lasting 1.5 s, providing the participant an interval to mentally retrieve or construct an image of the clock in its canonical position, face-forward and upright. The first rotation command was then auditorally presented, followed by silence, consuming a total of 3.2 s, during which the participant was to perform the first mental rotation. One-word commands indicated the rotation axis and the rotation direction. The command "clock" indicated a 90° clockwise rotation around the z-axis (i.e., the axis parallel to the line of sight, such that the clock was to be rotated in the picture plane). "Counter" indicated the corresponding counterclockwise rotation. The command "right" indicated a 90° rightward rotation around the y-axis (the vertical axis), and "left" indicated the corresponding leftward rotation. In problems with rotations about two axes, the rotations alternated axes from one rotation command to the next. After all of the commands had been presented, a picture of the clock in some orientation was presented for 3.5 s, during which time the participant indicated whether the picture correctly depicted the final position of the clock. Approximately 20% of the probes were "different," constructed by applying the final rotation around the specified axis but in the wrong direction. Participants were instructed to respond "same" or "different" by pressing one of two hand-held response buttons. Any participant with over 40% errors in any condition was excluded from the analysis.

The presentation duration of a problem ranged from 12.6 to 25.3 s, depending on the path length and rotation axis. The problems were presented in groups of the same problem type, composing epochs. To keep the durations of the six epoch types (conditions) similar at 25 to 30 s, the number of problems per epoch type was varied. The short path, one-axis (either z-axis only or y-axis only) epoch consisted of two 3-move problems totaling 25.2 s. The long path, one-axis (either z-axis only or y-axis only) epoch consisted of two 4-move problems totaling 29.6 s. The short path, two-axis epoch consisted of two 2-move problems totaling 27.8 s. The long, two-axis epoch consists of one 4-move problem totaling 25.3 s. Thus, the assessment of activation comparing fewer or more rotation steps per item compares such duration-equated epochs.

There were four tokens of each of the six types of epochs in a session, presented in four successive unique permutations of the six epoch types. Successive epochs were separated by 6-s rest periods during which individuals fixated an asterisk in the middle of the screen. There were also four presentations of 24-s fixation epochs distributed across the session, to provide a baseline measure of the brain activation.

Data Analysis

The image preprocessing corrected for head motion and signal drift using procedures and software developed by William Eddy and his colleagues in the Carnegie Mellon Statistics Department (Eddy, Fitzgerald, Genovese, Mockus, & Noll, 1996). The data analysis focuses on quantifying BOLD (blood oxygen level dependent)-contrast-related changes in the fMRI-measured signal using a dependent measure that takes into account both the volume of activation and the percentage change in signal, relative to a baseline level (Xiong, Rao, Gao, Woldorff, & Fox, 1998). Data from the first 6 s of each epoch were discarded to accommodate the rise of the hemodynamic response. For each voxel in the a priori defined ROIs (described below), the distribution of signal in an experimental condition was compared with that for the rest condition using a t test with a threshold of $t > 4.5$, corresponding to a high Bonferroni-corrected alpha level. We then calculated the integral of the percentage change in signal intensity for each region in each condition by summing the percentage change in signal intensity relative to the baseline estimate across voxels that reached the critical threshold in that condition. These data from each of the ROIs were then submitted to a 3 (z-axis, y-axis, both axes) × 2 (short or long path length) × 10 or 12 (participants; 10 only for regions uncovered by coronal slices) repeated measures analysis of variance.

To compare the amount of activation in a given volume across experimental conditions, we defined a priori ROIs for the cortical areas known to activate in mental rotation using the parcellation method originally described by Rademacher, Galaburda, Kennedy, Filipek, and Caviness (1992) and further refined by Caviness, Meyer, Makris, and Kennedy (1996), schematically depicted in Figure 2. The superior parietal ROI corresponded to the superior parietal lobule (SPL, BA 5 and 7). The inferior parietal lobule corresponded to the posterior supramarginal gyrus (SGp, BA 40) and the angular gyrus (AG, BA 39). The temporal ROI included both the superior and middle temporal gyrus (T1a, T1p, T2a, T2p, TO2). Inferior frontal gyrus was defined as the entire gyrus, including the pars opercularis and pars triangularis (F3t and F3o). The dorsolateral prefrontal cortex was defined as the middle frontal gyrus (F2) but excluding the anterior bank (2 voxels wide) of the precentral sulcus, which was defined as the frontal eye fields. The extrastriate/inferior temporal ROI defined for the 3 participants who also did the Shepard–Metzler task included the fusiform gyrus regions (TOF and TF; BA 36, 37, and 20) and the inferior temporal gyrus areas (TO3 and TFp; BA 37 and 20). The ROIs were defined by one of the authors (Mandy Maguire), after extensive training on the Rademacher–Caviness parcellation scheme. (The interrater reliability of this ROI-defining procedure between two trained staff members was evaluated for four ROIs in 2 participants. We obtained the reliability measure 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.)

Results

Behavioral measures of response accuracy were consistent with the predicted workload effect. The error rates were low (no more than 5%) for all the conditions except for the most demanding condition (in which the clock was to be rotated around two axes with the greater number of rotation steps), where the mean error rate was 13%.

Overview of fMRI Results

Although there were many conditions and ROIs, there was a remarkable uniformity to the findings across most of the ROIs (except for the temporal ROIs). First, successively rotating the clock around two different axes generated much more activation (about twice as much) as successively rotating the clock around a single axis (either the z- or the y-axis). Second, executing a greater number of rotation steps (i.e., a longer path length) generated more activation than executing a smaller number of rotation steps. Third, rotation around the z-axis (in the picture plane) generated a similar amount of activation to rotation around the y-axis (in depth). The results are presented in more detail below, organized by cortical area. Table 1 presents the results of the statistical analyses in each ROI, displaying the F ratio associated with the effects of the variables and their interaction, as well as the result of one of the contrast tests. The difference between z-axis and y-axis rotations (picture-plane vs. depth rotations) was generally minimal and in the same direction but always failed to reach statistical reliability, and hence we omitted the F ratio for this contrast test in Table 1 to conserve space. Table 1 also presents the centroids of the pooled activation in Talairach space for each ROI.

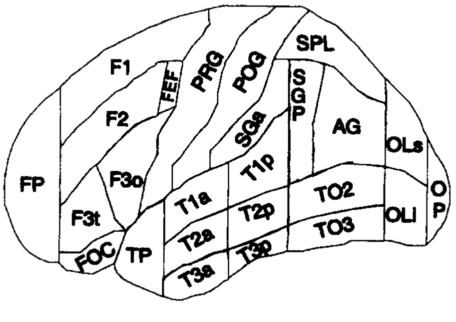

Figure 2. Cortical parcellation scheme adapted from Rademacher et al. (1992).

Parietal Areas

In all four parietal areas (left, right, inferior, and superior), there was a main effect of rotation axis, with much more activation in the dual-axis condition than in either of the single-axis conditions, as shown in Figure 3. Furthermore, the difference between single-axis rotations and dual-axis rotations was consistently large and statistically significant, as indicated by the results of a contrast test comparing the dual-axis conditions with the mean of the two single-axis conditions. The statistical results of the axis effects were the same when an analysis of variance (ANOVA) was applied to the data combined across all four parietal areas. Figure 4 illustrates the axis effects in one brain slice of 3 participants, showing the similarity in activation for the z-axis and y-axis rotations and the greater activation for the dual-axis rotations.

Activation was also significantly greater for longer versus shorter paths (operationalized in terms of the number of rotation steps), shown in Figure 5, although in the left superior parietal, the effect was marginal ($p < .06$). An interaction between axis and path length occurred only in the left superior parietal area. When an ANOVA was applied to the data combined across all four parietal areas, there was a main effect of path length and no interaction between path length and axis.

Inferior Frontal Gyrus (IFG)

The activation in both the left and right IFG was much greater in the dual-axis condition than in the single-axis conditions, as shown in Figure 6. For both the left and right IFG, there was a main effect of axis, as well as a significant contrast between dual-axis rotations and the mean of the single-axis rotations. There was no reliable difference between the activation for the y-axis and the z-axis rotations.

Left but not right IFG showed a main effect of path length, with more activation in the long path condition, as shown in Figure 7. The effect of path length was much larger in the dual-axis condition than in the single-axis conditions, resulting in a significant interaction. Generally, there was little difference between long and short conditions for the single-axis rotations, but the long, dual-axis rotation produced more activation than the short, dual-axis rotation. A possible explanation for this interaction is that IFG may

Table 1

F Ratios and MSEs from ANOVAs in Each Region of Interest (ROI)

ROI	Rotation axis effect		Single vs. dual axes: F^b	Path length effect		Axis × length interaction		Talairach centroids[a]		
	F^b	*MSE*		F^b	*MSE*	F^b	*MSE*	x	y	z
L. superior parietal	10.01***	231.96	19.97***			5.11*	43.88	−17	−60	47
R. superior parietal	9.89***	179.17	19.12***	5.88*	130.32			21	−60	48
L. inferior parietal	14.44***	101.65	28.53***	9.59*	38.77			−35	−53	40
R. inferior parietal	5.01*	54.07	10.03**	10.83**	30.06			38	−47	42
All four parietal	17.16***	1220.88	34.00***	14.36**	517.42					
L. inferior frontal	20.20***	17.02	40.19***	9.53*	13.16	5.11*	10.53	−39	17	20
R. inferior frontal	3.82*	30.68	7.63*			3.73*	4.05	40	18	19
L. dorsolateral prefrontal	15.78***	21.74	30.87***					−36	27	34
R. dorsolateral prefrontal	4.76*	24.65	9.51**	5.58*	28.17			32	24	42
L. temporal								−52	−29	3
R. temporal[c]			5.50*					53	−18	8
L. frontal eye fields	20.14***	11.44	40.18***	9.49*	3.91	4.71*	4.63	−35	1	47
R. frontal eye fields	9.13**	11.61	18.24***					35	3	46

Note. The results of the contrast tests comparing z-axis versus y-axis rotations, not shown here, were never reliable. L = left; R = right.

[a] The Talairach centroids were determined by (a) computing the mean coordinates of the activated voxels in each ROI for each of the 10 participants with axial-plane images in the dual axes, long path condition; (b) morphing these 10 participants' brains into Talairach space and converting their mean coordinates into Talairach coordinates; and (c) computing the mean of these 10 sets of Talairach coordinates. The Talairach axes are positive in the rightward (x), anterior (y), and superior (z) directions in the brain.

[b] The second degree of freedom of the F ratio in the parietal and temporal ROIs, based on 12 participants, was 11 or 22. In the remaining ROIs, based on only 10 participants (because the 2 participants with coronal slices did not have complete coverage), the second degree of freedom was 9 or 18.

[c] In the right temporal ROI, the contrast comparing single vs. dual axes, unlike in the other ROIs, the activation is greater in the single-axis condition.

$^* p < .05.$ $^{**} p < .01.$ $^{***} p < .001.$

be involved in a verbal encoding or rehearsal of the instructions while the cognitive coordinate system is generated for the required rotation. In the dual-axis rotations, the rotation axes (and hence perhaps the cognitive coordinate systems) alternate from instruction to instruction, possibly requiring more encoding or rehearsal activity in IFG, and especially so in the dual-axis, longer rotation path condition.

Dorsolateral Prefrontal Cortex (DLPFC)

For both left and right DLPFC (defined here as middle frontal gyrus, excluding the anterior bank of the precentral sulcus), there was a main effect of rotation axis. The activation was substantially less for single-axis rotations than for dual-axes rotations, as shown in Figure 8. The activation was similar for z-axis rotations and y-axis rotations. The effect of path length was not reliable in left DLPFC. For right DLPFC, the activation was significantly greater for the long condition than for the short condition.

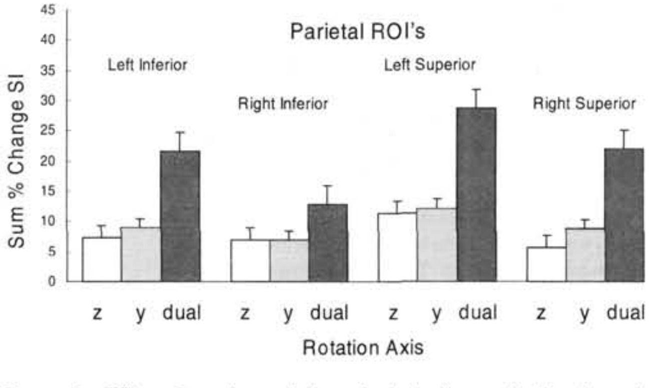

Figure 3. Effect of rotation axis in each of the four parietal regions of interest (ROIs), averaged across path length. Error bars indicate standard error of the mean based on the *MSE* for the rotation axis variable from the corresponding analysis of variance in Table 1. The measure on the y-axis is the sum of the percentage change in signal intensity (SI) in activated voxels.

Temporal Areas

Both the left and right temporal areas showed a substantially different pattern of activation than the other areas; here, the activation was *higher* in the single-axis rotations than in dual-axis rotations, as shown in Figure 9. Although the main effect of rotation axis was not reliable in either hemisphere of the temporal areas, the contrast in the right hemisphere temporal ROI between dual-axis rotations (M = 7.86) and single-axis rotations (M = 11.22) was reliable in this unusual direction, as shown in Table 1, and in the left hemisphere, it was marginally reliable, $F(1, 22) = 4.24$, $MSE = 29.51$, $p < .06$. Although the temporal ROI as defined here included the anterior portions of the superior and middle gyri, there was little activation in the anterior portions. Instead, the activation was concentrated in the posterior portions of the ROI: in Heschl's gyrus, and in the posterior portion of the superior temporal sulcus, in both the left or right hemisphere. These are the temporal areas that are activated in auditory sentence comprehension (Michael, Keller, Carpenter, & Just, 2001). There were no main effects of path length or interactions between path length and axis for the temporal areas.

This strikingly different pattern of results in the temporal lobe areas may be explained in terms of the number of items per epoch in each condition, and the impact of this variable on the processing of the auditory rotation instructions, such as "right." The main

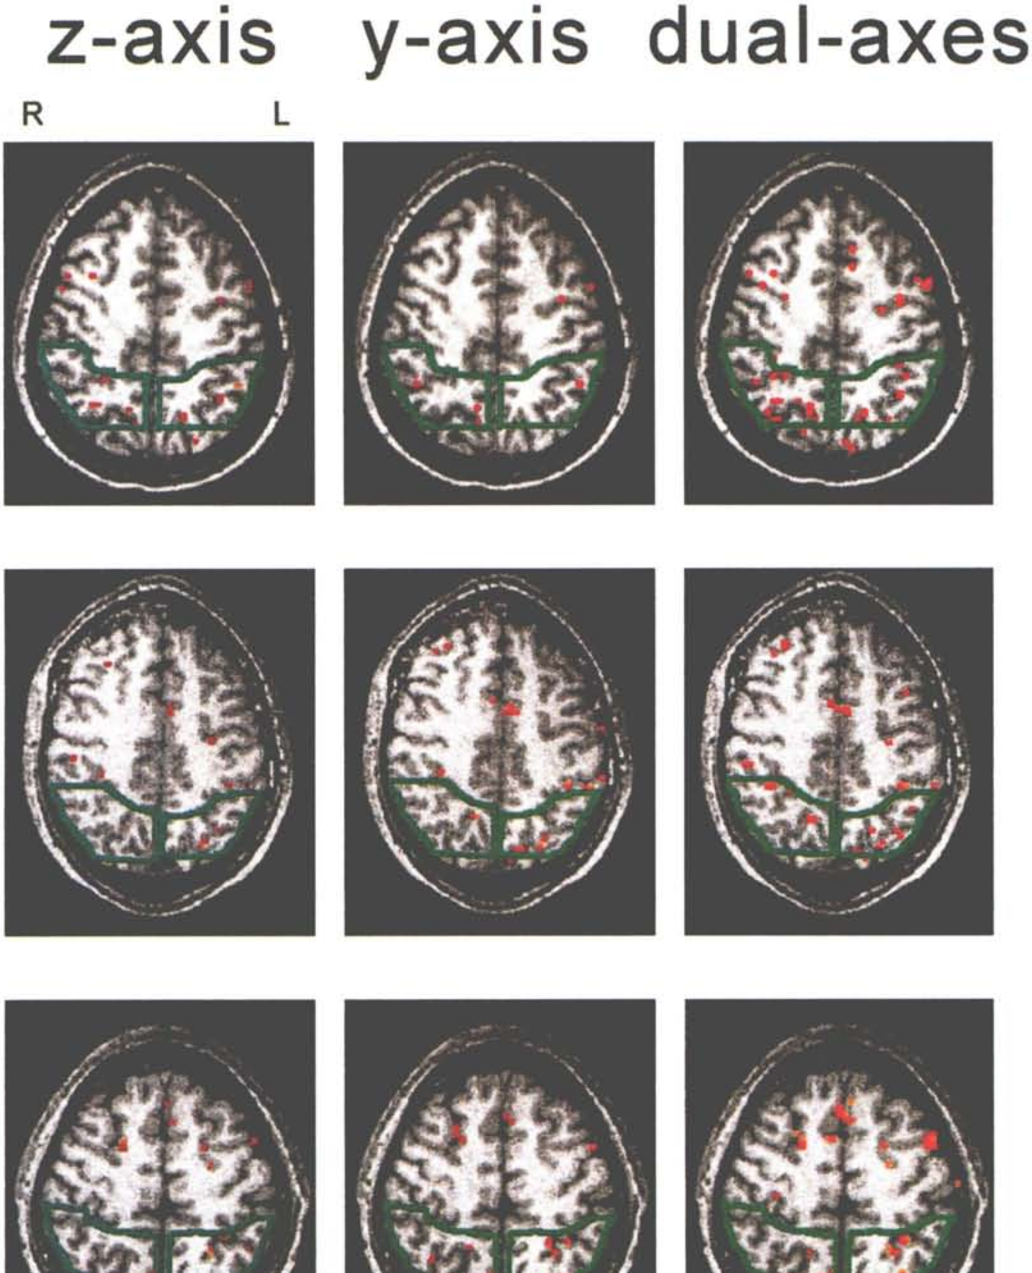

Figure 4. Activation in the most active parietal slice of 3 representative participants (one row per participant), shown for rotation around the z-axis, y-axis, and alternating between the y- and z-axes. The activated voxels are shown in red, superimposed on the structural image of each participant's own brain. The green outline shows the boundary of the superior parietal region of interest as it intersects this slice. The amount of activation is similar for z-axis and y-axis rotations but substantially greater for dual-axis rotations.

sense in which the dual-axis conditions could be construed as less demanding is that they contained only four rotation steps (hence, rotation instructions) per epoch, whereas the single-axis conditions entailed six or eight steps per epoch. It is likely that the greater activation in the temporal lobe areas in the single-axis rotation conditions reflects the perceptual and linguistic processing of the auditory rotation instructions. However, this explanation must be further elaborated to account for the greater activation in the right than in the left temporal area. During the auditory processing of language, there is typically more activation in the left than in the

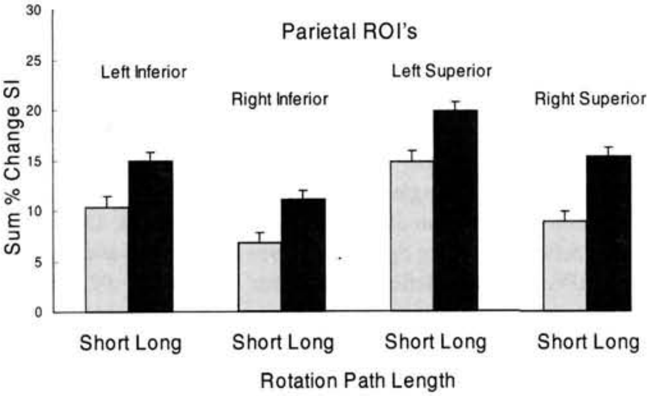

Figure 5. Path length effect in each of the four parietal regions of interest (ROIs), averaged across rotation axes. Error bars indicate standard error of the mean for this variable based on the *MSE.* The measure on the *y*-axis is the sum of the percentage change in signal intensity (SI) in activated voxels.

right temporal area (Schlosser, Aoyagi, Fulbright, Gore, & McCarthy, 1998). The right dominance of the obtained temporal activation probably entails other levels of processing besides the sensory, perceptual, and linguistic. The right dominance might reflect the spatial referent of the linguistic content. The right temporal area may be more involved when there is a mapping to be made from the language to some spatial representation. When the referent is spatial, then the relative activation in right temporal areas has been observed to be greater than otherwise, even in the case of visually presented sentences (Carpenter, Just, Keller, Eddy, & Thulborn, 1999b).

This distinction between the response of the temporal areas and the other areas reflects the relatively greater involvement of the temporal areas in the language-related processing versus the other areas' relatively greater involvement in the spatial processing. More generally, the results show the differential sensitivity of different parts of a large-scale cortical network to the demands imposed by different types of computations.

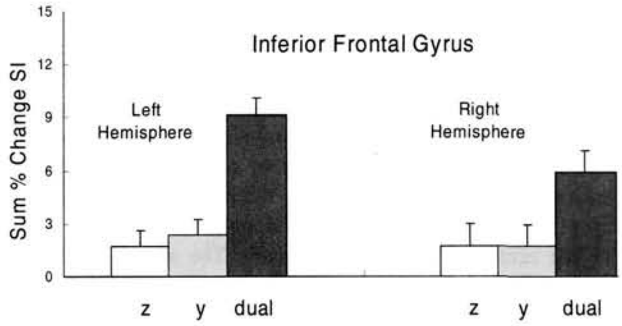

Figure 6. Effect of rotation axis in left and right inferior frontal gyrus, averaged across path length. Error bars indicate standard error of the mean for this variable based on the *MSE.* The measure on the *y*-axis is the sum of the percentage change in signal intensity (SI) in activated voxels.

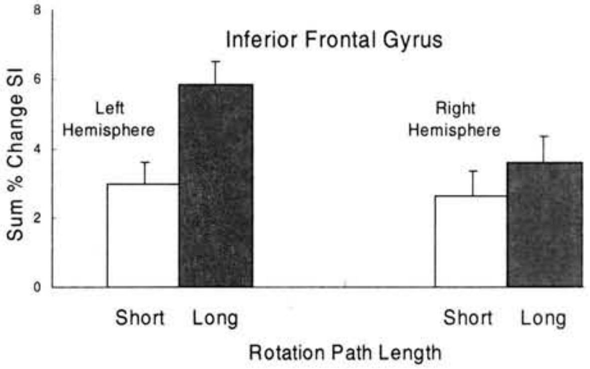

Figure 7. Path length effect in left and right inferior frontal gyrus, averaged across rotation axes. Error bars indicate standard error of the mean for this variable based on the *MSE.* The measure on the *y*-axis is the sum of the percentage change in signal intensity (SI) in activated voxels.

Frontal Eye Fields

Recall that during the mental rotation, nothing was being displayed on the screen, so the frontal eye fields are unlikely to be controlling eye movements, although they could well be involved in controlling spatial attention within an internal image space. Earlier research has shown that the control of attention within the external visual environment, even in the absence of eye movements, produces activation in the frontal eye fields (Corbetta et al., 1998). Both left and right frontal eye fields showed a main effect of rotation axis. Furthermore, the average activation for rotations in a single axis was significantly lower than dual-axis rotations. There was no reliable difference between *z*-axis and *y*-axis rotations in these areas.

In the left but not the right frontal eye fields there was a main effect of path length (with the long conditions producing more activation than the short conditions), as well as an interaction between path length and axis.

Laterality

The laterality effects in the clock task were rather systematic. In general, the amount of activation for most areas (except temporal)

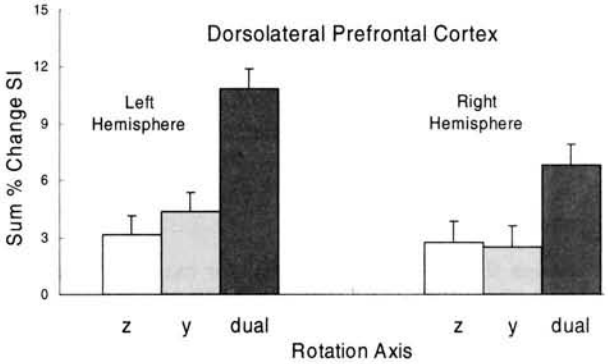

Figure 8. Effect of rotation axis in left and right dorsolateral prefrontal cortex, averaged across path length. Error bars indicate standard error of the mean for this variable based on the *MSE.* The measure on the *y*-axis is the sum of the percentage change in signal intensity (SI) in activated voxels.

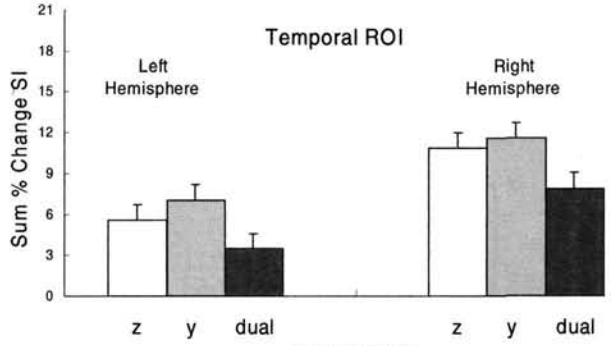

Figure 9. Effect of rotation axis in left and right superior and middle temporal area regions of interest (ROIs), averaged across path length. Error bars indicate standard error of the mean for this variable based on the *MSE*. The measure on the *y*-axis is the sum of the percentage change in signal intensity (SI) in activated voxels.

was greater on the left than on the right, particularly in some of the more demanding conditions. For example, in the inferior parietal ROI, the activation was particularly large in the left hemisphere in the dual-axis condition; there was a significant interaction between rotation axis and hemisphere for this ROI, $F(2, 22) = 3.46$, $MSE = 70.10$, $p < .05$. This effect is shown in Figure 3. A similar form of interaction occurred in the superior parietal ROI but was not reliable. The left dominance in the areas associated with spatial processing is similar to previous neuroimaging findings reported for an imagery task (D'Esposito et al., 1997). However, a previous study of mental rotation in the Shepard–Metzler task found that the right parietal area had slightly more fMRI-measured activation (Carpenter et al., 1999b). It may be that the difference in input modality between the Shepard–Metzler rotation and clock rotation tasks accounts for the small but reliable hemispheric differences being in opposite directions.

Another example of a hemisphere by complexity interaction occurred in the inferior frontal gyri. This area showed an interaction between path length and laterality, $F(1, 9) = 11.57$, $MSE = 2.36$, $p < .01$. The activation in left IFG (Broca's area) was particularly large for the long path problems, as shown in Figure 7.

The temporal ROI was unlike the others in that it showed a main effect of hemisphere with the right temporal ROI having greater activation, $F(1, 11) = 7.65$, $MSE = 105.42$, $p < .05$, as shown in Figure 9. There was also an interaction between length, axis, and hemisphere, $F(2, 22) = 5.63$, $MSE = 15.24$, $p < .05$.

Individual Differences

Although the sample size is small for examining individual differences, it nevertheless yielded results that are consistent with recent fMRI individual-difference findings (Reichle, Carpenter, & Just, 2000). In that study, participants with poorer verbal skills had larger increments in activation volume in left inferior frontal gyrus (L. IFG) when using a verbal strategy as compared with a visual strategy. Verbal ability was measured in that study and in the current one with the Daneman and Carpenter (1980) Reading Span Test. In the current clock rotation task, the scores on the Reading Span Test were also negatively related to the brain activation in L. IFG and the left inferior parietal area, namely the areas implicated in verbal rehearsal (Awh et al., 1996). The amount of brain activation was measured not as the absolute activation volume but as the number of additional voxels activated in the most demanding condition (dual-axis, long path) compared with the least demanding condition (single-axis, short path), that is, the slope of the activation as a function of the computational demand. The correlation between Reading Span Test scores and the activation slopes in L. IFG and left inferior parietal was −.70 and −.62, respectively ($p < .05$ for both correlations; $n = 10$ and 12, respectively). This pattern of correlations implies, quite plausibly, that people of lower verbal ability rehearse more. The correlations between visual–spatial ability (measured here with the Paper Folding Test; French, Ekstrom, & Price, 1963) and the activation slopes in L. IFG and left inferior parietal were negative, as expected, but they were smaller and not reliable. The individual difference findings suggest that the role of verbal rehearsal in this mental rotation task increases with decreasing verbal ability.

The Effect of Input Modality on Activation in Mental Rotation

The location of the brain activation in the clock rotation task was compared with the activation in the Shepard–Metzler task for the three participants with data of appropriate technical quality. The main goal was to determine whether the very same areas of the parietal lobule were activated for both tasks, despite the difference in the source or modality of the object to be rotated. The data analytic procedures for the two tasks were very similar, and the threshold used to define a voxel as "activated" was the same. The comparisons were between the more demanding conditions of each experiment: the dual-axis, longer-rotation conditions for clock rotation and the 80°-rotation condition in Shepard–Metzler, conditions that produced approximately similar amounts of activation in the parietal areas.

To determine whether the two tasks activated precisely the same $3.125 \times 3.125 \times 5$ mm voxels, we compared the number of such overlapping voxels to the total number of voxels activated (within each participant). The mean overlap in activated voxels in the left parietal ROI between the two tasks was .18 (i.e., the number of voxels that activated in both experiments divided by the mean of the number activated in the two tasks). In the right parietal ROI, this measure was .21. Figure 10 illustrates this overlap in detail for one slice of one participant. In addition to counting overlap in precisely the same voxels, we also counted whether at least one of the 18 voxels adjacent to an activated voxel in the clock task was also activated in the Shepard–Metzler task (*adjacent* was defined in terms of a voxel's center being less than 6 mm from the center of the activated voxel). The adjacency measure allowed for a possible slight change in the locus of activation or in the participant's head position between the two tasks. When the overlap measure additionally counted whether at least one adjacent voxel was activated, then the mean overlap measure in the left and right parietal ROIs was .62 and .58, respectively. These findings indicate a fairly close topographical similarity in the activation in the parietal areas between mental rotation of visually presented

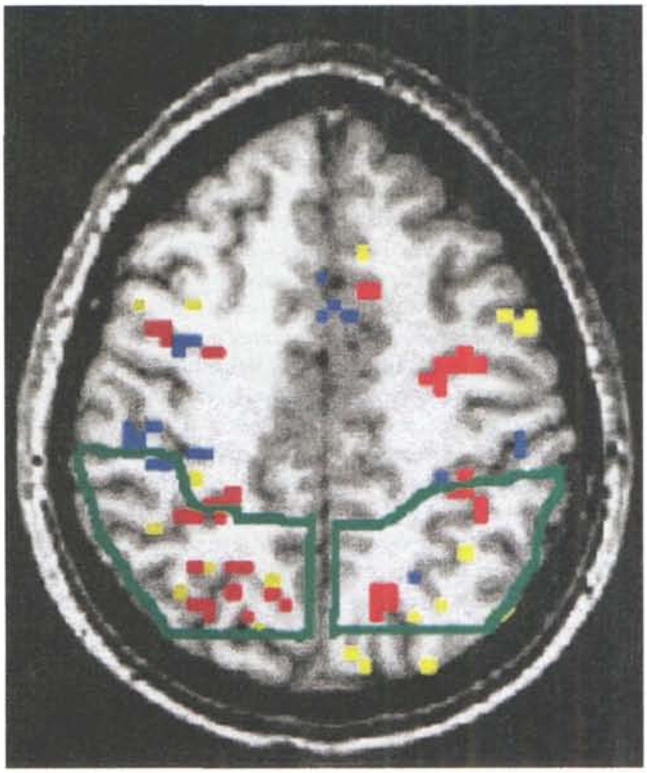

Figure 10. Comparison of activation in the clock rotation task and Shepard–Metzler rotation in one slice of 1 participant. Overlap (voxels activated to threshold in both studies) is shown in red. Voxels activated only in the clock study are shown in yellow. Voxels activated only in the Shepard–Metzler rotation are shown in blue. The green outline shows the boundary of the superior parietal region of interest as it intersects that slice.

objects and mental rotation of mentally retrieved objects, down to the level of individual voxels.[1]

Another cortical area worthy of comparison across the two tasks is the extrastriate cortex, the area between the primary visual region and the inferior temporal regions. The defined ROI includes the posterior portions of the inferior temporal lobule and the fusiform gyrus. Mellet, Petit, Mazoyer, Denis, and Tzourio (1998) cited a number of studies that reported activation in such areas when participants were performing mental imagery tasks in the absence of visual input. (It is important to note in the context of this particular comparison that the clock task did entail some visual stimulation, not as input to the rotation, but in the form of the probe stimulus.) There was some reliable activation in the extrastriate area in the clock task (a mean total of 4.7 voxels activated for left and right extrastriate combined in the hardest condition), but the activation was small in absolute volume and small relative to the volume in Shepard–Metzler (a mean total of 27.3 voxels activated in both hemispheres in the hardest condition). A larger study of the Shepard–Metzler task (Carpenter et al., 1999a) obtained a similar amount of activation in this region and found the activation volume to be modulated (but not monotonically) by rotation angle. So the retrieval and manipulation of a visual image without visual input may produce some activation in extrastriate cortex, but it is a very small amount.

Discussion

The main conclusions from this study are that the mental rotation of three-dimensional objects does not occur in a true three-dimensional space; that the rotation processing is substantially subserved by the parietal lobules bilaterally; and that the parietal contribution is similar regardless of input modality, whereas the extrastriate contribution is very small in the absence of visual input, but is substantial when the rotation operates on visual input. DLPFC, left prefrontal, and left inferior parietal may contribute to the goal management and rehearsal aspects of the task. Besides providing these broad conclusions, the results also indicated some finer grain characteristics of mental rotation that were previously unknown.

Cognitive Processing of Three-Dimensionality

Given that the world is three-dimensional, cognitive representations and transformations of physical objects might be expected to routinely incorporate their three-dimensional characteristics (Attneave, 1972; Shepard, 1981, 1984). However, several other studies indicate that there is an extra cost associated with representing three- as opposed to two-dimensional space. Imagining the movement of an object in a mentally generated array as signaled by visual or auditory cues (e.g., "up," "left") is easier in an imagined two-dimensional than three-dimensional array (Diwadkar, Carpenter, & Just, 2000; Kerr, 1987, 1993). Studies of the interpretation of visual form also indicate that the three-dimensional interpretation of drawings can be effortful, and its ease is affected by cues such as shading (Hemenway & Palmer, 1978) and the number of line segments or angles in the two-dimensional rendition of the space (Attneave & Frost, 1969; Hochberg & Brooks, 1960). Imagining the rotation of the three-dimensional object around an arbitrary axis is extremely errorful (Hinton, 1979; Pani, 1993). Such studies suggest that the cognitive interpretation of three-dimensional space requires more computation and maintenance than two-dimensional interpretations (Rock, 1983; Roth & Kosslyn, 1988). In a task requiring the mental rotation of a cube containing a single letter or numeral on each face, the verbal reports and eye fixations of participants who were low in spatial ability indicated that they operated on only one cube face at a time when the item required dual-axis rotation (Just &

---

[1] To help to interpret these values of overlap measures, we also computed the degree of activation overlap between the first and second half of the clock experiment for these 3 participants. We made the number of observations in each half comparable with the number used in the intertask comparisons by computing the overlap between the two halves of a new "condition" that consisted of the two hardest conditions in the clock study. The proportion of identical voxels activated in the left and right parietal ROIs, respectively, was .33 and .38. This degree of identical topographic overlap of activation between the two halves of the clock rotation experiment is almost twice as high as that between the clock rotation and the Shepard–Metzler task. When the overlap measure between the two halves of the clock task additionally counted whether at least one adjacent voxel was activated, overlap measures in the left and right parietal ROIs were .71 and .55, respectively. This degree of identical or adjacent topographic overlap of activation between the two halves of the clock rotation experiment is similar to that between the clock rotation and the Shepard–Metzler task.

Carpenter, 1985). The general conclusion is that the mind can do the three-dimensional work, but it requires extra computation. The mental three-dimensionality in rotation appears to be provided as it is needed and not as a default condition.

This general conclusion can be illustrated in terms of the effect of alternating rotation axes in the clock task. We briefly describe here a theoretical account of the effect. We propose that the mental rotation is performed by starting with a three-dimensional representation of the object and then reducing the representation to a two-dimensional projection in the relevant plane before rotating it. This may occur because of limits on the complexity of a representation to which a rotation transformation can be applied. A similar limitation appeared in a study of mental rotation of three-dimensional blocks with an alphabetic character on each face, where most participants reported rotating only one cube face at a time (Just & Carpenter, 1985). The limitation was manifested in the verbal protocols of the participants and was implemented in the computational model. In the case of single-axis clock rotations, this reduction to two dimensions needs to be performed only once, and then all the subsequent rotations can be applied to the two-dimensional representation. This account is consistent with the finding that in the self-paced norming study, participants took longer on the first rotation step than on subsequent rotation steps. In the condition that requires alternation of rotation axes, each subsequent rotation after the first requires that the three-dimensional representation be regenerated from the two-dimensional projection that was the outcome of the previous rotation. Then, the alternate two-dimensional projection must be generated before the next rotation is applied. The proposed algorithm indicates both a cognitive limitation on spatial processing capacity and a method for circumventing the limitation.

The fMRI results showing much greater activation (greater by at least a factor of 2) in the dual-axis condition indicate how the extra mental work engendered by alternating the rotation axes affects a network of cortical areas, including the parietal areas, but also extending to Broca's area and DLPFC. The added "spatial" work requires the added participation of language and executive areas, reflecting a systemwide demand on extra resources. The extra work performed by the nonparietal areas in the dual-axis conditions may include DLPFC's keeping track of whether various goals and subgoals are achieved, and Broca's area generating and rehearsing verballike expressions of intermediate states and goals. Any rehearsal of verbal representations could well involve phonological storage in the area of the angular gyrus, which is included in the inferior parietal ROI definition. Thus, when mental rotation is made difficult in this way, the difficulty is propagated well beyond the areas and processes that transform represented orientations.

Relation of Perception and Mental Processing

A classic issue in visual thinking concerns the degree of similarity (in terms of common neural substrate and algorithms) between perceptual processes and their purely conceptual counterparts (Finke, 1980; Kosslyn & Koenig, 1992). As the underlying processing becomes better understood, the similarities and differences become clearer. The results of this and other neuroimaging studies of mental rotation indicate that there are multiple brain areas involved, so evaluating the similarity in neural substrate entails a multidimensional comparison.

The results comparing the clock comparison task with the Shepard–Metzler task show remarkable similarity in the large amount of parietal activation, as well as in the precise overlap of activation locations within the superior parietal areas. It seems likely that the parietal area underlies the geometric processing of object orientation regardless of input modality. At the same time, the comparison between the two tasks illustrates that the perceptual processing beyond the primary sensory cortex, namely in the extrastriate cortex, is very different between the two tasks. There was enormously less activation in this area in the clock than in the Shepard–Metzler task, indicating that the activation there is not part of the geometric transformation itself. Nevertheless, the large amount of extrastriate activation in Shepard–Metzler varies with rotation angle (Carpenter et al., 1999a), so the processing there is sensitive to the two objects' relative orientations. The comparison of the two tasks contributes to an understanding of what the dorsal and ventral streams do, as well as their interplay.

The specific hypothesis that mental imagery activates primary visual cortex in the absence of visual input could not be evaluated in this study because the probe was presented visually.

Similarity of z- and y-Axis Rotations

The finding of similar amounts of activation throughout the activated cortical network for z- and y-axis rotations suggests that participants are initially representing the clock as a three-dimensional object. The alternative possibility, which is ruled out by these data, is that the clock was initially represented as a two-dimensional object as projected in the picture plane. Had this alternative been correct, then there would have been extra computations needed (and hence, more activation) for rotations in depth (around the y-axis) than in the picture plane (around the z-axis). The findings thus suggest that the clock was initially being represented as a three-dimensional object.

Consistent with this view, the rate for rotating Shepard–Metzler figures in depth in behavioral studies is similar to the rotation rate in the picture plane (Shepard & Metzler, 1971), and higher order apparent motion of complex figures is as easily realized in three-dimensional as in two-dimensional space (Shepard & Judd, 1976).

Path Length

The amount of activation in many areas increased with the number of 90° rotation steps made. The results confirm that the complexity of the mental work done during an epoch is a large determinant of the amount of activation, even when the epoch durations are equated. Having to rotate a clock through a longer series of steps imposes extra burdens of storing additional intermediate products and discriminating among them. The general finding of more activation for a greater number of rotation steps of a fixed magnitude extends the previous finding of an approximately linear increase in parietal activation obtained for single rotation steps of increasing magnitude from 0° to 120°, in 40° steps (Carpenter et al., 1999a). Mental rotation takes time proportional to the rotation angle (as behavioral studies have long indicated); the newer results indicate that not just the duration of the process but the amount of mental work done per unit time is proportional to the rotation angle.

Laterality

A common question in behavioral studies of cognition often concerned whether a given task was a right-hemisphere or a left-hemisphere task. Although the origin of such questions may have been the neuropsychological studies on split-brain surgical patients done by Sperry and his colleagues, behavioral studies finding a 40-ms advantage of sensory presentation to a given hemisphere tended to attribute the task processing to that hemisphere. Neuropsychological studies of patients with unilateral lesions similarly attributed a cognitive function to a given hemisphere (namely, the hemisphere whose lesioning produced the larger deficit in the function).

Functional neuroimaging studies cast a new light on the laterality issue. First, the laterality issue must be assessed in an entire network of brain areas, with each area presumably performing some distinguishable function, and the degree or even the direction of asymmetry between pairs of homologous areas is often different in different areas. So the laterality is multiregional and multidetermined. Second, the degree of asymmetry in a task is dependent on factors like task difficulty (Just, Carpenter, Keller, Eddy, & Thulborn, 1996) and on the input modality. For example, reading comprehension is reliably more left-lateralized than listening comprehension (Michael et al., 2001). Third, the ubiquitous empirical finding in fMRI is that although there is often a systematic asymmetry in the amount of activation of the two hemispheres, the degree of asymmetry is usually much smaller than would be expected, given that focal, unilateral lesions can produce profound impairments.[2] Thus, lateral asymmetry is just one facet of the characterization of the large-scale brain network that underpins a given task, and it is a multidimensional, dynamic facet at that.

The activation in the current study was often greater in the left hemisphere ROI than in the right hemisphere homologue. The asymmetry was not always reliable, but it always favored the left hemisphere, except for the temporal areas. The finding is consistent with Farah's (1984) meta-analysis of neuropsychological results linking left posterior hemisphere lesions with image retrieval impairment. We found that the areas associated with spatial processing had more activation on the left, whereas the language areas had more activation on the right. Such findings illustrate why it no longer makes sense to ask whether a task like clock rotation is a right-hemisphere task. The true situation can now be understood in most of its complexity, rendering a simplistic answer inaccurate and unuseful.

Summary

The characterization of the laterality results are prototypical of the more general implications of this study. The large-scale cortical network underlying mental rotation is dynamically configured as a function of the qualitative demands of the particular task, and the degree of activation of any given network component is determined by the amount of computational demand imposed on that component. For example, whether the objects to be rotated are visually presented or not determines whether the inferior temporal/extrastriate areas will play a major role in operating on the representations of complex objects. Despite the many dynamically determined properties of the network underlying mental rotation, what appears to be relatively stable is the participation of the parietal lobules. We presume that the parietal areas are involved in the geometric transformation of the represented orientation, in large part because their degree of activation increases approximately linearly with the rotation angle (Carpenter et al., 1999a). At the same time that the parietal areas may be central to mental rotation, a task as complex as clock rotation requires the coordinated action of several different types of processes, such as strategic and verbal processes that engage a number of areas in the frontal and temporal lobes. Thus, the fMRI findings permit the partial decomposition of the task performance by cortical area and allow the measurement of the workload of each area as the computational demands are manipulated from condition to condition. The resulting partial decomposition provides a more precise account of modality effects, individual differences, and laterality effects in mental rotation than was previously available, yielding a far richer explanation of how such tasks are performed.

---

[2] One inference from the combination of fMRI and neuropsychological findings on laterality is that although the participation of the "subordinate" hemisphere can be substantial in terms of brain activation, it may not be absolutely essential for task performance when the degree of difficulty is relatively undemanding. An additional inference is that although the two hemispheres may normally coperform the task, one hemisphere can exercise control over the other, and a lesion to the controlling hemisphere may disable not only the controlling hemisphere's participation but also the participation of the undamaged, controlled hemisphere. This type of control relation between the two hemispheres is not apparent from brain imaging studies of intact individuals, but a surgical intervention can provide relevant evidence. For example, surgical removal of a lesioned left hemisphere allowed a 9-year-old to acquire language for the first time (Vargha-Khadem et al., 1997).

References

Alivisatos, B., & Petrides, M. (1997). Functional activation of the human brain during mental rotation. *Neuropsychologia, 35,* 111–118.

Attneave, F. (1972). Representation of physical space. In A. W. Mellon & E. Martin (Eds.), *Coding processes in human memory* (pp. 283–306). Washington, DC: Winston.

Attneave, F., & Frost, R. (1969). The determination of perceived tridimensional orientation by minimum criteria. *Perception & Psychophysics, 6,* 391–396.

Awh, E., Jonides, J., Smith, E. E., Schumacher, E. H., Koeppe, R. A., & Katz, S. (1996). Dissociation of storage and rehearsal in verbal working memory: Evidence from positron emission tomography. *Psychological Science, 7,* 25–31.

Carpenter, P. A., & Eisenberg, P. (1978). Mental rotation and the frame of reference in blind and sighted individuals. *Perception & Psychophysics, 23,* 117–124.

Carpenter, P. A., Just, M. A., Keller, T., Eddy, W. F., & Thulborn, K. R. (1999a). Graded functional activation in the visuospatial system with the amount of task demand. *Journal of Cognitive Neuroscience, 11,* 9–24.

Carpenter, P. A., Just, M. A., Keller, T. A., Eddy, W. F., & Thulborn, K. R. (1999b). Time course of fMRI-activation in language and spatial networks during sentence comprehension. *NeuroImage, 10,* 216–224.

Caviness, V. S., Jr., 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.

Cohen, M. S., Kosslyn, S. M., Breiter, H. C., DiGirolamo, G. J., Thompson, W. L., Anderson, A. K., Bookheimer, S. Y., Rosen, B. R., & Belliveau, J. W. (1996). Changes in cortical activity during mental rotation: A mapping study using functional MRI. *Brain, 119,* 89–100.

Corbetta, M., Akbudak, E., Conturo, T. E., Snyder, A. Z., Ollinger, J. M., Drury, H. A., Linenweber, M. R., Petersen, S. E., Raichle, M. E., Van Essen, D. C., & Shulman, G. L. (1998). A common network of functional areas for attention and eye movements. *Neuron, 21,* 761–773.

Daneman, M., & Carpenter, P. A. (1980). Individual differences in working memory and reading. *Journal of Verbal Learning and Verbal Behavior, 19,* 450–466.

D'Esposito, M., Detre, J. A., Aguirre, G. K., Stallcup, M., Alsop, D. C., Tippet, L. J., & Farah, M. J. (1997). A functional MRI study of mental image generation. *Neuropsychologia, 35,* 725–730.

Desrocher, M. E., Smith, M. L., & Taylor, M. J. (1995). Stimulus and sex differences in performance of mental rotation: Evidence from event-related potentials. *Brain and Cognition, 28,* 14–38.

Diwadkar, V. A., Carpenter, P. A., & Just, M. A. (2000). Collaborative activity between parietal and dorso-lateral prefrontal cortex in dynamic spatial working memory revealed by fMRI. *NeuroImage, 12,* 85–99.

Eddy, W. F., Fitzgerald, M., Genovese, C. R., Mockus, A., & Noll, D. C. (1996). Functional imaging analysis software—computational olio. In A. Prat (Ed.), *Proceedings in computational statistics* (pp. 39–49). Heidelberg, Germany: Physica-Verlag.

Farah, M. J. (1984). The neurological basis of mental imagery: A componential analysis. *Cognition, 18,* 1–27.

Finke, R. A. (1980). Levels of equivalence in imagery and perception. *Psychological Review, 87,* 113–132.

French, J. W., Ekstrom, R. B., & Price, L. A. (1963). *Manual for a kit of reference tests of cognitive factors.* Princeton, NJ: Educational Testing Service.

Guilford, J. P., & Zimmerman, W. S. (1947). *The Guilford-Zimmerman Aptitude Survey: Manual of instructions and interpretations.* Palo Alto, CA: Sheridan Psychological Services.

Hemenway, K., & Palmer, S. E. (1978). Organizational factors in perceived dimensionality. *Journal of Experimental Psychology: Human Perception and Performance, 4,* 388–396.

Hinton, G. (1979). Some demonstrations of the effects of structural descriptions in mental imagery. *Cognitive Science, 3,* 231–250.

Hochberg, J., & Brooks, V. (1960). The psychophysics of form: Reversible-perspective drawings of spatial objects. *American Journal of Psychology, 73,* 337–354.

Just, M. A., & Carpenter, P. A. (1985). Cognitive coordinate systems: Accounts of mental rotation and individual differences in spatial ability. *Psychological Review, 92,* 137–172.

Just, M. A., Carpenter, P. A., Keller, T. A., Eddy, W. F., & Thulborn, K. R. (1996). Brain activation modulated by sentence comprehension. *Science, 274,* 114–116.

Kerr, N. H. (1987). Locational representation in imagery: The third dimension. *Memory & Cognition, 15,* 521–530.

Kerr, N. H. (1993). Rate of imagery processing in two versus three dimensions. *Memory & Cognition, 21,* 467–476.

Kosslyn, S. M., & Koenig, O. (1992). *Wet mind: The new cognitive neuroscience.* New York: Free Press.

Mellet, E., Petit, L., Mazoyer, B., Denis, M., & Tzourio, N. (1998). Reopening the mental imagery debate: Lessons from functional anatomy. *NeuroImage, 8,* 129–139.

Michael, E. B., Keller, T. A., Carpenter, P. A., & Just, M. A. (2001). An fMRI investigation of sentence comprehension by eye and by ear: Modality fingerprints on cognitive processes. *Human Brain Mapping, 13,* 239–252.

Pani, J. R. (1993). Limits on the comprehension of rotational motion: Mental imagery of rotations with oblique components. *Perception, 22,* 785–808.

Peronnet, F., & Farah, M. J. (1989). Mental rotation: An event-related potential study with a validated mental rotation task. *Brain and Cognition, 9,* 279–288.

Rademacher, J., Galaburda, A. M., Kennedy, D. N., Filipek, P. A., & Caviness, V. S., Jr. (1992). Human cerebral cortex: Localization, parcellation, and morphometry with magnetic resonance imaging. *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 skill in sentence verification. *Cognitive Psychology, 40,* 261–295.

Rock, I. (1983). *The logic of perception.* Cambridge, MA: MIT Press.

Rosler, F., Heil, M., Bajric, J., Pauls, A. C., & Henninghausen, E. (1995). Patterns of cerebral activation while mental images are rotated and changed in size. *Psychophysiology, 32,* 135–149.

Roth, J. D., & Kosslyn, S. M. (1988). Construction of the third dimension in mental imagery. *Cognitive Psychology, 20,* 344–361.

Schlosser, M. J., Aoyagi, N., Fulbright, R. K., Gore, J. C., & McCarthy, G. (1998). Functional MRI studies of auditory comprehension. *Human Brain Mapping, 6,* 1–13.

Shepard, R. N. (1981). Psychophysical complementarity. In M. Kubovy & J. R. Pomerantz (Eds.), *Perceptual organization* (pp. 279–341). Hillsdale, NJ: Erlbaum.

Shepard, R. N. (1984). Ecological constraints on internal representation: Resonant kinematics of perceiving, imagining, thinking, and dreaming. *Psychological Review, 91,* 417–447.

Shepard, R. N., & Judd, S. A. (1976). Perceptual illusion of rotation of three-dimensional objects. *Science, 191,* 952–954.

Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science, 171,* 701–703.

Vargha-Khadem, F., Carr, L. J., Isaacs, E., Brett, E., Adams, C., & Mishkin, M. (1997). Onset of speech after left hemispherectomy in a nine-year-old boy. *Brain, 120,* 159–182.

Wijers, A. A., Otten, L. J., Feenstra, S., Mulder, G., & Mulder, L. J. M. (1989). Brain potentials during selective attention, memory search, and mental rotation. *Psychophysiology, 26,* 452–467.

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.

Received March 6, 2000
Revision received August 28, 2000
Accepted October 2, 2000 ■