BOLD Activity during Mental Rotation and Viewpoint-Dependent Object Recognition

Isabel Gauthier,[1,5] William G. Hayward,[2]
Michael J. Tarr,[3] Adam W. Anderson,[4]
Pawel Skudlarski,[4] and John C. Gore[4]
[1]Vanderbilt Vision Research Center
Department of Psychology
Vanderbilt University
502 Wilson Hall
Nashville, Tennessee 37203
[2]Department of Psychology
Chinese University of Hong Kong
Shatin, New Territories
Hong Kong
[3]Department of Cognitive and Linguistic Sciences
Brown University
Providence, Rhode Island 02912
[4]Department of Radiology
Yale School of Medicine
New Haven, Connecticut 06510

Summary

We measured brain activity during mental rotation and object recognition with objects rotated around three different axes. Activity in the superior parietal lobe (SPL) increased proportionally to viewpoint disparity during mental rotation, but not during object recognition. In contrast, the fusiform gyrus was preferentially recruited in a viewpoint-dependent manner in recognition as compared to mental rotation. In addition, independent of the effect of viewpoint, object recognition was associated with ventral areas and mental rotation with dorsal areas. These results indicate that the similar behavioral effects of viewpoint obtained in these two tasks are based on different neural substrates. Such findings call into question the hypothesis that mental rotation is used to compensate for changes in viewpoint during object recognition.

Introduction

Since Shepard and Metzler's (1971) (see also Shepard and Cooper, 1982) groundbreaking discovery of mental rotation, the noninvasive technique of functional magnetic resonance imaging (fMRI) has been used to investigate the neural substrates of this cognitive task. The most consistent fMRI finding relates activation in the superior parietal lobe (SPL) to performance in mental rotation tasks (Alivisatos and Petrides, 1997; Carpenter et al., 1999; Cohen et al., 1996; Tagaris et al., 1996, 1997; Richter et al., 1997).

Critically, investigations of mental rotation generally rely on a "handedness" judgment in which participants discriminate between mirror reflections of either a 2D or a 3D object. The canonical finding is that response times increase monotonically with increasing rotation distance (Shepard and Cooper, 1982). At the same time,

numerous studies of visual object recognition reveal a similar pattern of viewpoint-dependent performance (Corballis and McLaren, 1984; Humphrey and Khan, 1992; Jolicoeur, 1985; Lawson et al., 1994). When these two tasks are compared directly using the same novel objects as stimuli, there is remarkable similarity in both the qualitative monotonic shape of the response function and in the quantitative rate of rotation (Tarr and Pinker, 1989; Tarr, 1995). This has led to the hypothesis that mental transformations akin to mental rotation are integral to object recognition (Jolicoeur, 1990; Tarr, 1995). Although the behavioral data are compelling, they cannot distinguish between a common neural mechanism and different neural substrates that happen to execute viewpoint-dependent transformations producing similar patterns of performance. Indeed, visual object recognition is thought to be mediated primarily by occipito-temporal areas, in which cells are found that respond to objects in a viewpoint-selective manner (Perrett et al., 1987; Logothetis et al., 1995). This stands in contrast to the association of mental rotation with activation in SPL and motor areas in numerous fMRI studies. Here, we specifically ask whether we find differences in the neural substrates supporting mental rotation (MR) and object recognition (OR) when similar viewpoint-dependent behavior is observed for both tasks.

Evidence for separable neural bases for the viewpoint-dependent subcomponents of mental rotation and object recognition would oblige us to reconsider some accounts of viewpoint-dependent object recognition (Tarr, 1995). Note that we are not attempting to address the ongoing "view-dependent/view-independent" debate in the object recognition literature (Tarr and Bülthoff, 1995; Biederman and Gerhardstein, 1995). We are simply assuming that there are conditions under which viewpoint-dependent mechanisms are used in object recognition. An argument for a reconsideration of the "rotate-to-recognize" hypothesis has already been made on the basis of indirect comparisons between the two tasks (Perrett et al., 1998) and computational considerations (Poggio and Edelman, 1990; Yuille and Steiger, 1982). Tarr and Bülthoff (1998) explicitly develop an account in which the normalization procedures used for object recognition tasks are not the same as those used for mental rotation tasks. These critiques of the rotation-for-recognition approach have taken two distinct directions, positing either: (1) separate mechanisms for OR and MR tasks; or (2) a single mechanism, but one that is not isomorphic with continuous transformations in the physical world as originally posited by Shepard and Cooper (1982).

Complicating matters is the fact that effects of viewpoint may not reflect a single mechanism, but rather a set of transformation mechanisms that can be applied to object representations depending on stimulus configuration (e.g., similar versus dissimilar objects; see Hayward and Williams, 2000) and task demands (e.g., MR versus OR). In particular, rotations around different axes place far different demands on putative transformation mechanisms. Rotations in the picture plane preserve

[5]Correspondence: isabel.gauthier@vanderbilt.edu
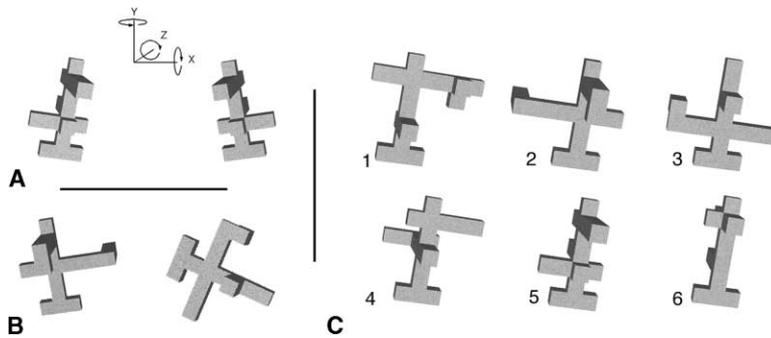
Figure 1. Examples of Objects Used in the MR and OR Tasks

(A) Mirror reflections of the same object.
(B) Two different objects. On each trial, subjects made a judgment about two stimuli presented sequentially (S1 for 500 ms and S2 for 3000 ms). In the MR condition, subjects determined whether S1 and S2 were the same or different handedness. In the OR condition, the task was to determine whether S1 and S2 showed the same object or two different objects.
(C) The complete set of objects in one handedness from the canonical view. Objects 1, 2, and 3 were used exclusively in the OR task, while objects 4, 5, and 6 were used exclusively in the MR task.

the visibility of all of the features of an object, but perturb the top-bottom relations between those features. Rotations in depth around the vertical axis alter the visibility of features, some coming into view and others becoming occluded (unless the object is completely nonselfoccluding, as with paperclips; see Bülthoff and Edelman, 1992), but do not change the top-bottom relations of those features. Although such rotations alter the left-right relations between features, almost all theories of object representation assume that left and right are not explicitly represented (see for instance, Tarr and Pinker, 1989; Biederman, 1987). Finally, rotations in depth around a nonvertical axis alter both top-bottom relations between features and which features are visible. Thus, the transformation mechanisms used to compensate for changes in viewpoint under each of these cases may vary (Hummel and Biederman, 1992; Tarr and Bülthoff, 1998) and yet lead to similar effects on performance.

Our design was based on the one used by Tarr (1995), including the same novel objects used in that study, themselves based on the objects created by Shepard and Metzler (1971) (see Figure 1). Fifteen subjects performed two different judgments about pairs of objects presented sequentially: an *MR task*—whether the two objects were identical or mirror reflections of each other; and an *OR task*—whether the two objects had the same shape (in which case two mirror images of the same object should be considered the same; see Tarr, 1995). Object pairs for both tasks were either separated by small or large viewpoint differences generated by rotations in the picture plane or in depth around the vertical or horizontal axis. Prior fMRI studies of mental rotation have typically used simultaneous presentation of object pairs and often for very long durations (e.g., up to 8 s in Tagaris et al., 1996), possibly encouraging eye movements or shifts of attention back and forth between objects presented together—both are behaviors that can activate the SPL (Corbetta et al., 1998; Coull and Nobre, 1998; Luna et al., 1998). We used a sequential-matching paradigm with only a single object shown at any given time. Although subjects may still scan an object more thoroughly for larger degrees of rotation in our paradigm, the sequential matching task provides a way to ensure that the role of the SPL is not related to a strategy specifically tied to simultaneous displays rather than to mental rotation generally.

Results

Behavioral Results

We attempted to equate mental rotation and object recognition by using a sequential-matching paradigm with a common trial structure in both conditions (Figure 1). These manipulations produced the expected effects of viewpoint on both response times (RT) and accuracy in each task and for all axes of rotation (Figure 2). However, there are also some differences between tasks. First, accuracy was higher for OR than MR trials, an effect that was more pronounced at large angular disparities. Indicating that there was no speed-accuracy tradeoff between tasks, OR trials also led to faster responses. This result is not entirely unexpected in that MR trials require a more specific judgment, discriminating the handedness of a particular object, as compared to OR judgments, which only necessitate discriminating object identity.

These observations are supported by three-way, within-subjects ANOVAs, with task (MR/OR), magnitude of rotation (small/large), and rotation axis (x/y/z) as factors. For accuracy, all three main effects were statistically significant, revealing better performance for OR than MR, $F(1,12) = 24.26$, $p < 0.001$; for small versus large rotations, $F(1,12) = 31.16$, $p < 0.01$; and for the axis of rotation, $F(2,24) = 21.48$, $p < 0.001$. For this last result, Scheffé ($p < 0.05$) tests revealed better performance for picture plane (z) rotations than either type of rotation in depth (x/y). The interaction between task and magnitude of rotation was also significant, $F(1,12) = 9.97$, $p < 0.01$, suggesting that the magnitude of rotation had a larger effect on performance in the MR condition than OR condition. The interaction between task and axis of rotation was likewise significant, $F(2,24) = 7.93$, $p < 0.01$, as there appear to be larger differences between the axes in OR as compared to MR. Finally, the three-way interaction between task, rotation magnitude, and rotation axis was marginally significant, $F(2,24) = 3.24$, $p = 0.057$. The two-way interaction between rotation magnitude and rotation axis was not significant, $F = 1.69$.

For mean RT, the results were very similar. Again, all three main effects were reliable: responses were faster in OR than MR, $F(1,12) = 9.69$, $p < 0.01$; responses were faster to small rotations as compared to large rotations,
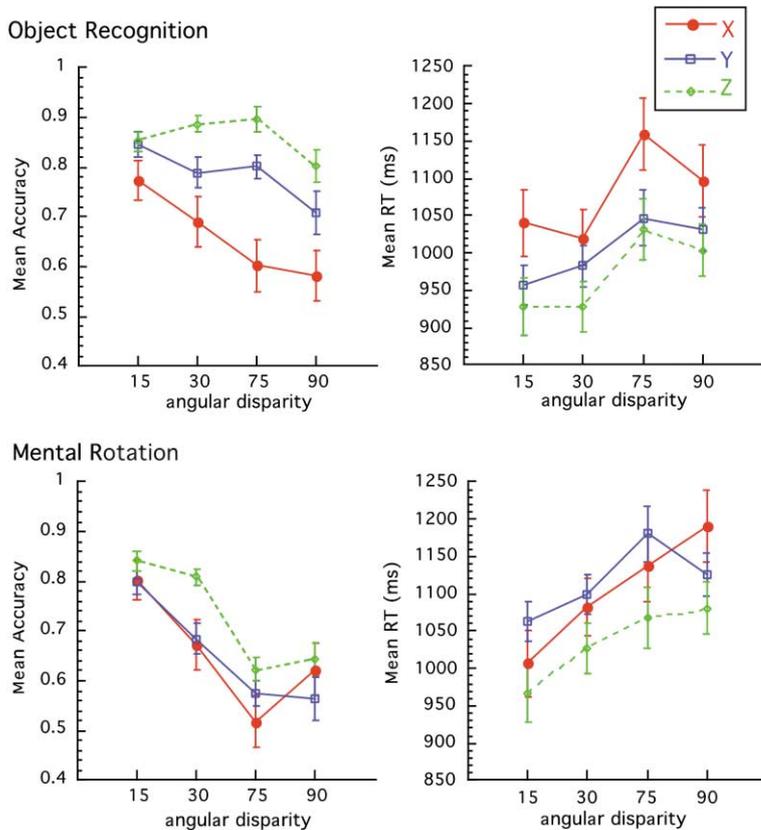
Figure 2. Mean Accuracy (Left) and Response Times (Right) in the OR (Top) and MR (Bottom) Tasks for Each Axis of Rotation, Collapsed Across Subjects

Error bars show the SEM. Results are plotted separately for each angle for comparison with behavioral studies. However, in the fMRI experiment, trials were blocked by small disparity (15° and 30°) and large disparity (75° and 90°).

$F(1,12) = 26.95$, $p < 0.001$; and there were differences across the three axes of rotation, $F(2,24) = 12.96$, $p < 0.001$. Scheffé ($p < 0.05$) tests again revealed differences between picture plane (z) and depth rotations (x/y). The only interaction that was significant was the interaction of MR/OR and axis of rotation, $F(2,24) = 7.31$, $p < 0.01$, as there were larger differences between axes in OR than MR. The interactions for task × rotation magnitude, $F < 1$, axis of rotation × rotation magnitude, $F = 1.77$, and the three-way interaction, $F < 1$, were all nonsignificant. Because of the large number of errors made in both tasks, accuracy is a more reliable measure of performance than is RT for correct responses. This may explain the failure of interactions in RT to reach significance even though both accuracy and RT exhibit similar patterns, with no indication of a speed-accuracy trade-off. Reinforcing the pattern of behaviors which we obtained in the scanner, examination of the main effects shown in Tarr (1995), restricted to orientation differences comparable to our manipulations (less than 90°), shows a pattern of results similar to ours.

Role of the SPL in Mental Rotation versus Object Recognition

In several neuroimaging studies of mental rotation, activation in the SPL was correlated with accuracy and/or RTs (Tagaris et al., 1996, 1997; Richter et al., 1997). We isolated a comparable region of interest (ROI) and examined its activity (as measured by fMRI using the blood oxygen level dependent [BOLD] contrast) across axes and tasks. For this purpose, we used six of the nine scan series conducted for each subject (the first

and last run for each axis) to identify the region in the area surrounding the intraparietal sulcus (including both SPL and some inferior parietal lobe, as in prior MR studies) that responded more during large rotation than small rotation MR trials (see Figure 3). We then used the data from the remaining three series to compare activity in this region in the MR versus OR task as a function of axis. This "internal localizer" method allows us to use a subset of the data to define the ROI and the other subset, not used in the localizer, to evaluate the activation level.

As shown in Figure 3, the BOLD effect in the SPL ROI revealed an interaction between task and the effect of viewpoint. Specifically, SPL activity increased with larger rotations for MR, but did not increase with viewpoint for OR (it even decreased with viewpoint for z rotations). A three-way ANOVA with task (MR and OR), axis (x, y, and z), and magnitude of rotation (small, large) confirmed this pattern. The interaction between task and magnitude of rotation was significant, $F(1,14) = 7.80$, $p = 0.014$. There was also a significant interaction of magnitude of rotation with axis, $F(2,14) = 5.94$, $p = 0.007$. This was due to the activity for both tasks combined increasing with larger rotations for axes x and y, but decreasing with larger rotations for the z axis. The three-way interaction was not significant ($p = 0.45$).

The interaction between task (MR and OR) and magnitude of rotation in the SPL is particularly interesting because this dissociation is not observed in the behavioral results; MR led to larger viewpoint effects than OR, but in both tasks, the effect of viewpoint was significant and in the same direction. Although our fMRI design does not include a parametric manipulation of view-
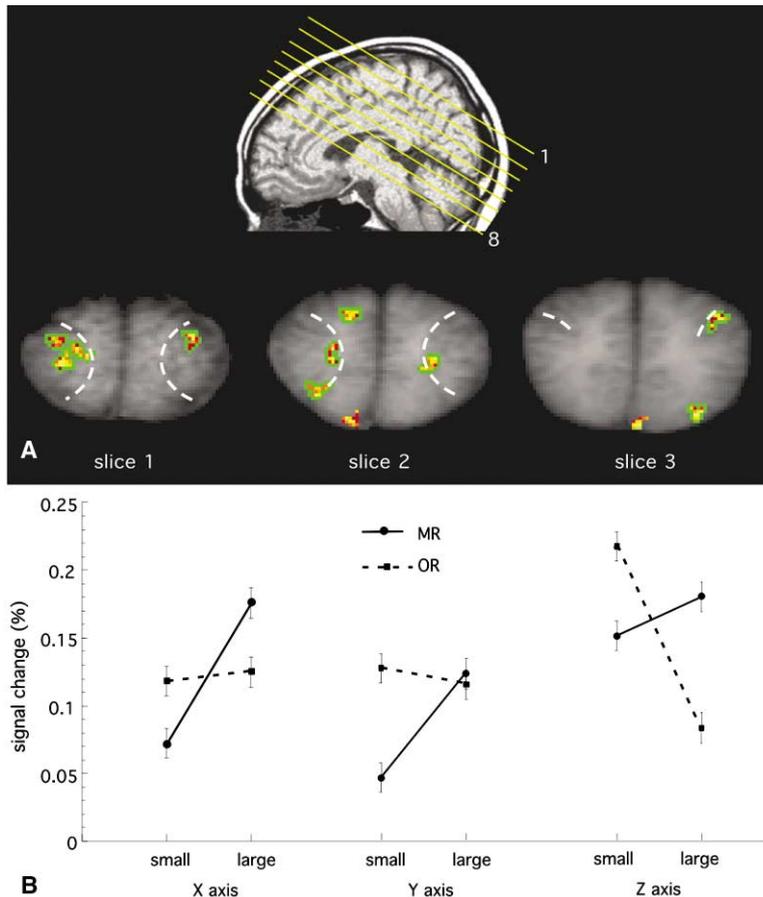
Figure 3. Activity in SPL during Mental Rotation and Object Recognition

(A) SPL/IPL region of interest (ROI). The right hemisphere is shown on the left. The ROI (outlined in green) was defined as the region in the area surrounding the intraparietal sulcus that showed a viewpoint effect in the MR task ($p < 0.05$, one-tailed), using only a subset of the data for each subject. Broken lines show the approximate location of the IPS. The location of the slices shown in Figures 3, 4, and 5 is illustrated on a sagittal image.

(B) Interaction between task (MR/OR) and the magnitude of rotation (small/large) for each axis in the SPL/IPL ROI.

point, the response of the parietal lobe to increasing viewpoint has been shown to be monotonic for MR in other studies (Richter et al., 1997; Carpenter et al., 1999). Thus, while we cannot extrapolate the activity obtained during the OR task to predict what would be found in a parametric design, our effects clearly indicate that the parietal lobe plays a different role during MR and OR tasks, even in cases where highly comparable behavioral effects of viewpoint are observed.

Activity in the SPL also depends on the axis of rotation. It is clear from Figure 3 that this region was recruited more heavily in OR judgments for small rotations around the z axis as compared to any other OR condition. Small rotations around the z axis may present a situation that is particularly unambiguous (with no feature occlusion and relatively little perturbation of top-bottom relations), leading subjects to spontaneously adopt a mental rotation strategy. This strategy may be less efficient when part occlusion is present (Rock et al., 1989). In contrast, subjects may have had no other option than to use a mental rotation strategy in the MR trials, when handedness judgments were required.

Brain Areas Sensitive to Viewpoint Effects in Mental Rotation versus Object Recognition
The relationship between task and magnitude of rotation can be explored further by asking whether other regions of the brain show a significant interaction between these

two factors. We investigated this issue in a voxel-based analysis with either all axes of rotation combined or each axis separately.

The interaction map (weights of 1, −1, −1, and 1 for the MR-large, MR-small, OR-large, and OR-small conditions, respectively) identified several regions exhibiting a significant interaction (Figure 4). Consistent with the ROI analysis discussed above, the right and left parietal lobes (SPL/IPL) show increased activity with larger rotations in the MR task but not in the OR task. This pattern was obtained for all three axes, although in slightly different parietal regions for each one. In addition to these parietal areas, other regions showed a similar pattern of activity, including a region of BA19 in the right lateral occipital sulcus. However, none of these additional areas exhibited this pattern significantly for all three axes (most did only for one axis).

There were also regions with increased activity with larger rotations in OR, but not MR, including part of left BA19 and the right inferior temporal gyrus. Although nonoverlapping, regions of the right middle temporal gyrus showed this pattern for rotations around both the y and z axes.

Thus, although these analyses can be interpreted as indicating that parietal regions are more closely associated with the viewpoint effect in mental rotation whereas temporal regions play a similar role for object recognition (Wilson and Farah, 1999), this interpretation should
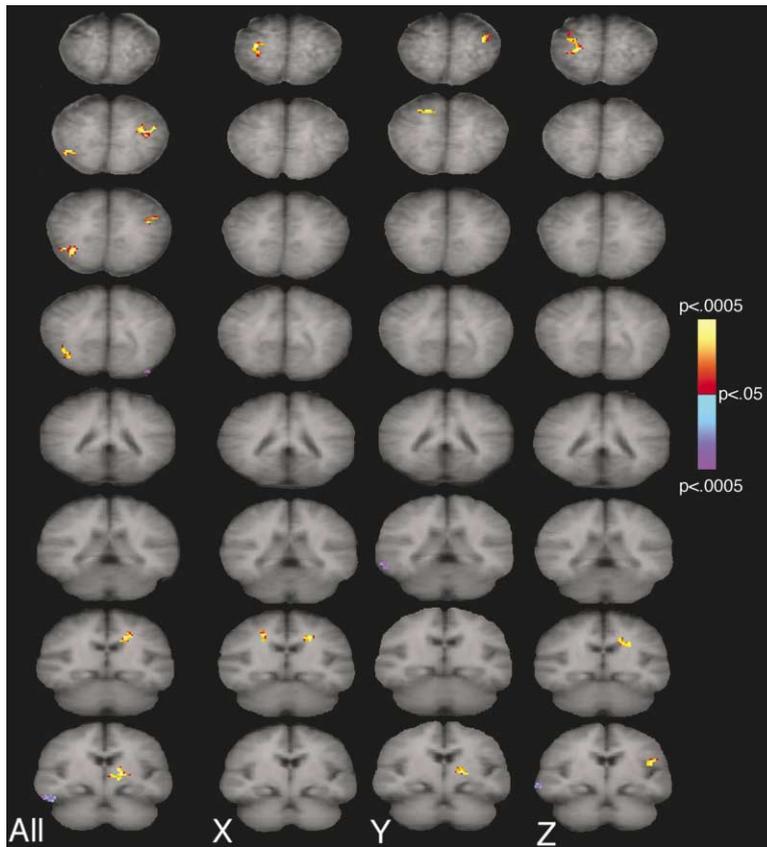
Figure 4. Brain Regions Showing a Task by Magnitude of Rotation Interaction (p < 0.05) with Talairach Coordinates

Only contiguous areas with a peak reaching p < 0.001 and a volume of activation of at least 0.22 cm³ at a p < 0.05 are shown. In yellow to red are areas showing an increase in activity with larger rotations for MR and a decrease for OR (from top to bottom). All axes: IPL (−36, −48, 46; 39, −72, 26; −41, −34, 44), BA19 (35, −68, 16; 39, −60, 11), cingulum (−21, 9, 30), thalamus (−15, −7, 5). x axis: SPL (27, −59, 49), cingulum (±21, 10, 30). y axis: SPL (−33, 47, 58), precentral gyrus (19, −29, 61), thalamus (−16, 7, 5). z axis: SPL (24, −57, 50), cingulum (−23, 5, 26), insula (−40, 6, 15). In blue to purple are areas showing an increase in activity with larger rotations for OR and a decrease for MR. All axes: BA19 (−38, −81, −6), inferior temporal gyrus (52, −31, −15); y axis: middle temporal gyrus (56, −43, 0); z axis: middle temporal gyrus (62, −20, −6).

be qualified. Extensive differences are found in the regions exhibiting viewpoint-related activity depending on the axis of rotation. (Sections of this study describing the role of the SPL in MR versus OR and viewpoint by axis interaction present more direct tests of such differences.) This suggests that the differences in the measured neural substrates associated with specific viewpoint effects are mediated by the geometrical and spatial constraints of the task (e.g., different processes may be involved depending on whether parts are occluded or not by rotation and whether the spatial relations between parts change; Biederman and Gerhardstein, 1993; Hayward and Tarr, 1997). Thus, similar geometrical and spatial constraints could also account for some of the differences obtained between the OR and MR judgments.

These findings should not be interpreted as evidence that there is *no* overlap between the neural substrates associated with the normalization processes invoked during MR and OR. An analysis isolating areas with a viewpoint-dependent BOLD response in both MR and OR yielded several common regions of activity. However, at a 0.05 α level, the only regions showing a common pattern for both tasks actually exhibited *decreased* activity with larger viewpoints (see Table 1). At an even less stringent α level of 0.10, three areas appear with a common pattern of increase with larger viewpoints in both OR and MR: part of the right SPL, the left IPL, and a region in the right precentral gyrus. However, an ANOVA on the activity in these regions that included these three ROIs as a factor (i.e., task × viewpoint ×

ROI) revealed that the task × viewpoint interaction was dependent on the brain area [the three-way interaction was significant, $F(2,28) = 7.89$, p = 0.002]. Specifically, only the left IPL ROI showed an increase in activity with viewpoint in both OR and MR (% signal change in MR small: 0.04; MR large: 0.08; OR small: 0.01; OR large: 0.05). In contrast, the right SPL ROI showed an increase in activity with viewpoint only for MR (% signal change in MR-small: 0.04; MR-large: 0.14; OR-small: 0.12; OR-large: 0.10), whereas the precentral gyrus showed an increase in activity with viewpoint only for OR (% signal change in MR-small, 0.06; MR-large, 0.03; OR-small, 0.01; OR-large, 0.07).

Whereas our analyses reveal many differences in the neural substrates of MR and OR, we find only one area in the left IPL showing a similar increase with larger rotations for both MR and OR. This area was slightly more dorsal to the left IPL areas that showed the same interaction as the SPL (activity increased with larger rotations for MR and the reverse for OR). Such results support the hypothesis of limited overlap in the neural substrates underlying viewpoint-dependent processes in MR and OR.

Differences between Mental Rotation and Object Recognition Independent of Viewpoint Effect

The next question we address is whether there are differences in the neural substrates underlying MR and OR tasks that are not dependent of the effect of viewpoint. The most basic model would assume that the two tasks share all components up to the additional handedness

Table 1. Talairach Coordinates for the Center of Areas Showing a Main Effect of Magnitude of Rotation for Both MR and OR Judgments, All Axes Combined

	Right/Left	Front/Back	Dorsal/Ventral
Small-large rotations			
($p < 0.05$, volume > 0.22 cm^3)			
Precuneus	6	−48	45
Supramarginal gyrus	44	−52	29
Post cingulate gyrus	3	−44	24
Middle frontal gyrus	42	14	46
Large-Small rotations			
($p < 0.10$, volume > 0.22 cm^3)			
Superior parietal lobe	33	−51	55
Inferior parietal lobe	−43	−38	53
Precentral gyrus	59	3	13

judgment required in MR. In OR judgments, the two stimuli must be aligned and then compared ignoring handedness; in contrast, in MR judgments, the two stimuli must be aligned and then compared, taking handedness into account. This account predicts that some brain areas should be more active during MR than OR judgments, but equally for small and large rotations (this does not preclude that the activity in such areas is equally sensitive to viewpoint in both tasks).

The results of analyses comparing MR and OR are shown in Figure 5. The simple model of MR as an OR task plus handedness judgment appears to be refuted; in particular, there are several brain areas more active for OR than MR judgments. With all axes combined, the left IPL was more active for MR than OR, but the fusiform gyrus (BA19) was more active for OR than MR judg-

ments. When considering the rotation axes separately, several ventral occipito-temporal areas were more active for OR than MR judgments; for the x axis, this included the right precuneus, right BA19 in the middle occipital gyrus, and BA20 in the right inferior temporal and right fusiform gyri and, for the z axis, BA19 in the middle occipital gyrus and the inferior temporal gyrus. For z axis rotations, the right middle frontal gyrus and bilateral parietal areas in the IPL/angular gyrus showed more activity for MR than OR, as did occipital areas in the precuneus and BA17.

In general, consistent with results in the previous section, MR appears more strongly associated with dorsal areas whereas OR is more strongly associated with ventral areas (there is no dorsal area more active for OR than MR). Again, the differences in brain activity observed
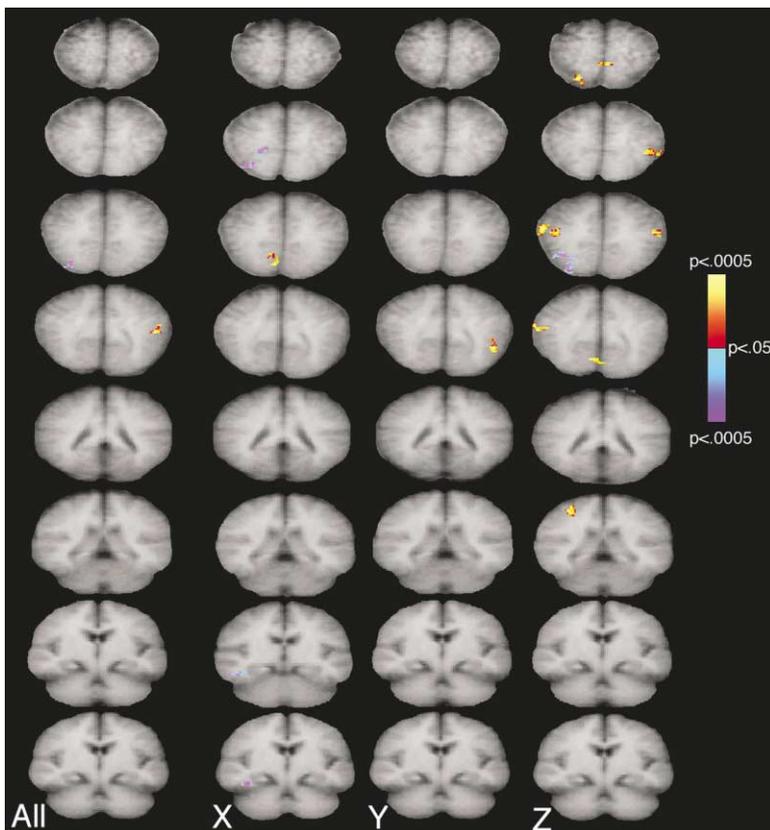
Figure 5. Brain Regions Showing a Difference between MR and OR Independently of the Viewpoint Effect ($p < 0.05$) with Talairach Coordinates

Only contiguous areas with a peak reaching $p < 0.001$ and a volume of activation of at least 0.22 cm^3 at a $p < 0.05$ are shown. All axes: IPL (−50, −39, 28); x axis: BA17 (7, −72, 13); y axis: superior temporal gyrus (−45, 55, 15); z axis: precuneus (−3, −70, 39), BA19 (22, −88, 25), angular gyrus (−45, −64, 32), IPL (±50, −42, 37; 57, −32, 33), BA17 (2, −68, 4), middle frontal gyrus (30, 11, 43). In blue to purple are areas showing more activity for OR than MR. All axes: BA19-fusiform gyrus (31, −78, −5); x axis: precuneus (21, −61, 34), BA19 (32, −77, 22), inferior temporal gyrus (46, −36, −7), fusiform gyrus (37, −27, −12); z axis: inferior temporal gyrus (36, −67, 17), BA19 (31, −81, 7).

Table 2. Talairach Coordinates of Areas which for Each Task Show a Larger Viewpoint Effect for One Axis than for the Other Two Axes

	Left/Right	Front/Back	Dorsal/Ventral
Object recognition			
x axis			
Fusiform gyrus (BA19)	38	−66	6
Fusiform gyrus (BA37)	43	−43	−5
Middle frontal gyrus	21	1	48
Parahippocampal gyrus	−22	−27	−5
Thalamus	−14	−17	8
z axis			
Inferior parietal lobe	52	−35	31
Mental rotation			
x axis			
Superior temporal gyrus	−54	−48	20
Inferior frontal gyrus	−53	7	28

between MR and OR occur even in cases where behavioral performance, as measured by accuracy and RTs, are matched between the two tasks (this is the case for the x axis trials, as well as for small rotation z axis trials, which in a separate analysis of brain activity, resulted in the same pattern as seen for all z rotations; data not shown).

We investigated whether any of the regions that showed preferential activity for one task over the other also showed a viewpoint effect by using the regions in Figure 5 as ROIs. In each ROI, we compared the activity during each small and large rotations for each subject (using only the axis or axes of rotations used to define the ROI). The only ROI to show a significant main effect of viewpoint was the fusiform gyrus area (BA19), for all axes combined. That is, this area not only showed significantly more activity for OR than MR judgments, but it also showed an overall viewpoint effect (small rotations, 0.19% signal change; large rotations, 0.24% signal change; $F(1,14) = 5.11$, p = 0.04). There was no interaction of this viewpoint effect with task (p = 0.36).

Viewpoint by Axis Interaction

The analyses presented above all point toward important differences in the neural substrates recruited by normalization mechanisms around different axes of rotation. However, only the ROI analysis in the SPL explicitly tested for these effects. Finding a considerable influence of axis of rotation on the neural substrates of MR and OR is important in the context of the differences observed between these two tasks in our study. If our experiment revealed mostly MR/OR differences in brain activity but little influence of the axis of rotation, these differences could be interpreted as by-products of the different strategies by which observers must necessarily approach these two tasks (Milner and Goodale, 1995). However, an axis of rotation effect would suggest that normalization mechanisms operating on different principles are recruited according to task. Moreover, obtaining such differences in MR would suggest that mental rotation is not a direct analog of physical rotation (as argued by Shepard and Cooper, 1982), but rather depends on the effect of 3D transformations on the 2D projected images.

In Table 2, we report the areas showing a significantly larger viewpoint effect for one axis compared to the other two axes combined, in each task separately. Note that this analysis is more easily interpretable for MR than OR, because in the MR task, the effect of viewpoint was comparable in both accuracy and RTs for all three axes of rotation, whereas there were significant differences in observed behavior between axes for OR. In both tasks, some brain areas showed axis-selective viewpoint effects, even in the face of equivalence in behavioral performance during MR judgments. In fact, the analysis shows a more important effect of axis for MR than OR, perhaps suggesting that OR judgments use a representation that is closer to an analog of physical rotation than MR judgments.

When they are present, we attribute these differences to the fact that rotations around different axes differentially change the appearance of an object in terms of feature visibility, as well as top-bottom and left-right spatial relations between features. Our measurement of the BOLD effect during these transformations supports the hypothesis that separable normalization procedures are used to compensate for these different kinds of changes. However, our results do not allow us to make any specific statements about the computational properties of the processes supported by these brain areas, but instead point to the need for psychophysical studies that do not treat all object rotations as identical (e.g., Biederman and Gerhardstein, 1993; Hayward and Tarr, 1997).

Discussion

Neural Substrates

The brain areas isolated in our study can be divided into four subsets based on their different response patterns and their theoretical implications. First, certain parietal brain areas showed an increase in activity with larger viewpoint differences in MR, but a decrease in the same conditions for OR. This includes the SPL and IPL, the areas most often associated with MR in prior studies. The neural network associated with MR thus appears to respond quite differently during OR judgments, suggesting that despite their superficial similarity, the two tasks recruit very different processes.

A second subset of areas showed an increase in activity with larger viewpoint disparities in OR, but a decrease under the same conditions for MR. This includes several

areas of the classical ventral temporal pathway (right inferior and middle temporal gyrus, BA20, and BA21; see above) associated with object recognition (Ungerleider and Mishkin, 1982; Perrett et al., 1987). This is consistent with findings of viewpoint sensitivity in object-selective cells in the inferior temporal lobe (e.g., Logothetis et al., 1995) and the observation that visual agnosic patients with ventral temporal lobe lesions can show increased sensitivity to viewpoint transformations in OR tasks (e.g., Suzuki et al., 1997). However, this result is inconsistent with suggestions that viewpoint-dependent object recognition merely reflects the spurious activity of processes associated with MR in the parietal lobe during object recognition (Biederman and Gerhardstein, 1993, 1995). As an aside, it should be noted that although areas associated with MR more than OR were found bilaterally, by far most of the areas more active (or showing a larger increase of activity with viewpoint changes) for OR than MR were found in the right hemisphere.

A third subset of areas showed a similar increase in activity with increasing angular disparities in both MR and OR. Part of the left IPL (Table 1 and see above) and a region of the right fusiform gyrus located in what would be the object-selective area dubbed "LO" in studies of object processing (e.g., Grill-Spector and Malach, 2001) showed similar viewpoint-dependent responses during MR and OR. However, in contrast to the IPL area, the fusiform area also showed a significant task effect, with more activity during OR than MR. Different interpretations can be made regarding these common areas. It is possible that some of the processes involved in viewpoint-dependent object processing are relatively automatic and are recruited regardless of the task. This is likely to be true for the ventral areas more heavily involved in OR; although responses in the object-selective areas of the ventral temporal lobe can be influenced by attention, object-selective responses in these areas can be obtained under a wide range of conditions, even when subjects attend to dimensions other than object identity (e.g., Gauthier et al., 2000). Thus, the viewpoint-dependent activity in LO during MR may reflect the automatic recruitment of processes that typically mediate object recognition (perhaps the "default" task of vision) and may not be central in mental rotation (Carpenter et al., 1999). Conversely, it is possible that IPL activity reflects processes associated with MR automatically engaged during OR, but the automaticity of such processes is more questionable, especially given past evidence that we do not encode handedness unless required to do so (Corballis, 1982; Hinton and Parsons, 1981). Finally, it is possible that the constant alternation between the OR and MR has led us to overestimate the common responses in the two tasks.

Finally, a fourth subset of areas showed more activity for MR than OR *or* more activity for OR than MR, but did not show a viewpoint-dependent response in either case. Interestingly, these areas selectively recruited by MR or OR also reflected a general pattern of dorsal areas associated with MR and ventral areas associated with OR. Even though the two tasks that we compared are among those that have been most extensively studied and for which behavioral performance is surprisingly similar, they clearly invoke a complex set of visual and cognitive processes, certainly not limited to viewpoint-

dependent ones, of which we have an incomplete understanding (see Just and Carpenter, 1976). Neuroimaging offers a new means of investigating the componential complexity of these tasks. At the very least, the finding of brain areas more active for OR than MR (in addition to those more active for MR than OR) allows us to firmly reject one characterization of these two tasks in which MR consists of no more than viewpoint-dependent object recognition followed by a handedness judgment.

Axis Matters
Considering some of our more specific results, we find that activity in part of the SPL exhibits a viewpoint-dependent BOLD effect during mental rotation that was the opposite of the viewpoint-dependent effect in object recognition during z axis rotations. This suggests that the viewpoint-dependent parietal activity observed here and in earlier fMRI studies may reflect processing that is particularly important for mental rotation at large viewpoint disparities and object recognition at small viewpoint disparities for rotations in the picture plane. How could this be so? Consider that many models of object recognition assume that multiple views are encoded for each object (Tarr and Pinker, 1989; Poggio and Edelman, 1990; Tarr, 1995). One of the reasons for multiple views is distinctiveness; that is, observers represent views of an object that are different from one another using some metric of geometric or image-based similarity (Biederman and Gerhardstein, 1993; Tarr and Bülthoff, 1998; Tarr and Gauthier, 1998; Tarr and Kriegman, 2001). In some viewpoint-dependent accounts of object recognition, these multiple views are matched to *nearby* input shapes through normalization procedures akin to mental rotation, but input shapes far from any known view are matched using geometric inferences and other non-transformation-based processes (Hayward and Tarr, 1997; Tarr and Bülthoff, 1998). Normalization processes can only match views when there is some correspondence between features; large rotation differences are far more likely to yield new views because feature correspondences cannot be established. Thus, object recognition at small disparities may be performed using a process similar to that used during mental rotation, but this same process cannot be used for object recognition judgments at large disparities—an interpretation consistent with our present fMRI results.

The activation we and others observe in SPL and IPL in association with MR reinforces the interpretation of mental rotation as an analog of physical rotation (Shepard and Cooper, 1982) because these areas have been associated with perception for action (Milner and Goodale, 1995; Ungerleider and Mishkin, 1982). However, the axis-dependent nature of our results suggest that mental rotation may not be a continuous transformation process, but rather may be sensitive to the changing image geometry of the object being rotated (Hochberg and Gellman, 1977; Folk and Luce, 1987).

Common Normalization Mechanisms for Viewpoint-Dependent Object Recognition and Mental Rotation?
When two tasks yield very similar performance, it is tempting to infer common underlying mechanisms. One

example is Tarr and Pinker's (1989) (Tarr, 1995) observation that mental rotation and viewpoint-dependent object recognition of identical objects yield highly similar patterns of performance across changes in viewpoint during initial learning, following extensive practice, and when new viewpoints are introduced.

However, we found striking dissociations between the neural substrates of viewpoint-dependent OR and MR, in many cases accompanied by comparable behavioral performance. Such results require a reassessment of claims regarding the relationship between viewpoint-dependent object recognition and mental rotation. In particular, as suggested by Perrett et al. (1998) and Yuille and Steiger (1982), the putative transformation mechanisms used in the two tasks may not be as similar as originally hypothesized (Tarr and Pinker, 1989; Jolicoeur, 1985). In contrast, our findings are consistent with a recent model offered by Perrett et al. (1998). These authors propose that object recognition is mediated by the accumulation of evidence in populations of viewpoint-specific neurons tuned to specific object features. The response of populations of viewpoint-selective cells in the ventral cortical stream would explain both why subjects learn multiple familiar views and why we observe a general pattern of viewpoint-dependent behavior in object recognition tasks (Bülthoff and Edelman, 1992; Logothetis et al., 1994; Tarr and Pinker, 1989; Tarr, 1995; Tarr and Gauthier, 1998). This explanation is supported by the significant viewpoint-dependent activity we observed in ventral areas for the OR task and is consistent with the finding of viewpoint-selective neurons in inferotemporal cortex (Perrett et al., 1987; Logothetis et al., 1995).

In conclusion, we used fMRI to examine the neural substrates underlying performance for two very similar tasks: mental rotation and viewpoint-dependent object recognition. Although some functions are probably shared between the two tasks, there are critical differences that indicate that MR and OR are, to a large extent, distinct from one another. In particular, mental rotation should not be seen as an additional "check" that is executed postrecognition, and viewpoint-dependent object recognition should not be seen as being mediated by mental rotation. Supporting these two points, we observed the association of parietal activity with mental rotation and of ventral temporal activity with object recognition—patterns consistent with the hypothesis that mental rotation is supported by areas often associated with motor planning (e.g., Georgopoulos et al., 1989), and viewpoint-dependent object recognition is supported by areas associated with object representation (e.g., Perrett et al., 1998). It is clear that fMRI allows an understanding of perceptual function that moves beyond behavioral comparisons.

Experimental Procedures

Subjects
Fifteen neurologically normal right-handed volunteers (8 males, 7 females) from the Yale University community gave informed and written consent and were recruited in exchange for payment.

Stimuli
Six novel three-dimensional objects were created using StrataVision 3D (Strata, St. George, UT). Examples are shown in Figure 1. The objects were based on six of the seven objects created by Tarr (1995). Each object was asymmetrical around any possible axis. All objects were illuminated from a single overhead point and rendered with a ray-tracing algorithm without cast shadows. Each object had a "foot" (three cubes perpendicular to the main axis of the object) to mark the object's top-bottom orientation. For each object, a canonical view was established at a view in which the foot was at the bottom and the other parts were oriented either left-right or front-back and rotated 10° around the horizontal (x), vertical (y), and line-of-sight (z) axes. The objects are depicted from their canonical views in Figure 1. New views of each object were generated by rotating a given object in its canonical view around either the x, y, or z axis in 15° steps (only one axis at a time, from 0° to 135°); this procedure produced nine views per axis of rotation. Mirror reflections of each object were are also generated and then subjected to the same procedure for generating views. Therefore, each object was depicted in 28 views in its standard version and 28 views in its mirror-reflected version. Although the visual angle subtended varied with each view, the maximum dimension of any image was 388 pixels horizontally and 310 pixels vertically, and objects were included within a square area of approximately 6° × 6° of visual angle.

Task Design
Subjects performed two tasks on sequentially presented views of the objects described above. In the MR task, trials were initiated by a central fixation cross for 200 ms. Following a blank interval of 175 ms, an object was shown for 500 ms. After a second blank interval of 500 ms, the same object was shown again, and it remained on the screen for 3000 ms, regardless of how long subjects took to respond. The second stimulus presentation always showed the object rotated relative to its first presentation; half of the time, it was also mirror reversed. For MR trials, the task was to determine whether the two images depicted the same mirror version or whether the two images showed mirror reflections. Subjects pressed one button to respond "same version" and a different button to respond "mirror version". To prevent transfer from one task to another, only objects 4, 5, and 6, as shown in Figure 1, were used in the MR task.

In the OR task, the structure of each trial was very similar to the MR task. A central fixation cross appeared for 200 ms, followed by a 175 ms blank interval, then the first object for 500 ms, a second blank interval of 500 ms, and then finally a second object for 3000 ms. In OR, however, the same object was shown in the second stimulus presentation on only half of the trials; in the other half, the second stimulus presentation depicted a different object from the first (not simply a mirror reflection). If the second stimulus presentation showed the same object, it was always rotated relative to the first stimulus presentation, and half of the time it was also mirror reversed. For OR trials, the task was to determine whether the two images depicted the same object or two different objects. They were explicitly instructed to ignore any differences in viewpoint or mirror reversal. Subjects pressed one button to respond "same object" and a different button to respond "different object". To prevent transfer from one task to another, only objects 1, 2, and 3, as shown in Figure 1, were used in the OR task.

Stimuli for MR and OR trials were drawn from the set of rendered images of objects with the following constraints. First, in the MR condition, a given object was always paired with a rotated version of itself or a rotated version of its mirror reflection. In contrast, in the OR condition, a given object was paired either with a rotated (and possibly flipped) version of itself or a rotated (and possibly flipped) version of a different object. Second, for both the MR and OR conditions, when the second object was rotated relative to the first object, the view difference was always restricted to a rotation around a single axis (x, y, or z). Moreover, object views were selected so that the object pairs, whether the same or different, were separated by rotation differences of 15°, 30°, 75°, or 90°.

For example, MR trials around the x axis could use any single object for which there were available ten images (canonical plus nine x axis rotations) of the standard version and the ten images of the mirror-reflected version. The particular images chosen upheld the constraint that the view differences be 15°, 30°, 75°, or 90°. In this example, "same version" trials consisted of two images that either showed two standard or two mirror-reflected views of the

same object—in either case, a 3D rotation of one object would align it perfectly with the other object. "Different version" trials consisted of one object in the standard version and one object in the mirror-reflected version; here, a rigid body rotation of one object could never bring it into perfect alignment with the other object.

An example OR trial around the x axis could use any single object for which there were available ten images (canonical plus nine x axis rotations) of the standard version and the ten images of the mirror-reflected version. For same object trials, the second image was also drawn from this set of 20 images and, again, the particular images chosen upheld the constraint that the view differences be 15°, 30°, 75°, or 90°. For different object trials, the second image was drawn from 20 images of an entirely different object, but the view difference constraint was still applied.

Practice trials outside of the scanner included 12 trials per axis and per task (total of 72 trials), blocked both by axis and by task as in the fMRI task (6 trials in a row for the same axis and task). Feedback was provided during practice only.

In the fMRI experiment, there were nine series of trials, each lasting 252 s. There were 432 trials in total, 48 trials in each series; 24 were MR and 24 were OR (of these, 12 were small rotations and 12 were large). In series 1, 4, and 7, all rotations were about the x axis. In series 2, 5, and 8, rotations were about the y axis, and series 3, 6, and 9 showed the objects rotating about the z axis (picture plane rotation). Within each series, subjects performed eight blocks of six trials, where each block was either MR or OR and either small rotations (15° or 30°) or large rotations (75° or 90°). To designate whether a block of trials was MR or OR, they were preceded by a screen containing the words "Same Version?" (to denote MR) or "Same Object?" (to denote OR), shown for 4.5 s. The factors of task and magnitude of rotation were manipulated within fMRI series; the factor of axis of rotation was manipulated between fMRI series.

Behavioral Analyses

The behavioral data for 2 of the 15 subjects was unavailable due to equipment malfunction. Trials were omitted from the behavioral analyses if subjects failed to respond within 3000 ms of the onset of the second object. This resulted in the exclusion of 11% of the total trials for the remaining 13 subjects. An ANOVA on the number of missing trials revealed that there were significantly more missing trials in the MR than OR task [7% versus 15%; $F(1,12) = 13.66$, $p < 0.01$] and more missing trials for y axis judgments (14%) than for x or z axis judgments [10% and 9%, respectively; $F(2,12) = 20.2$, $p < 0.01$]. No other main effect or interaction was significant. These results are consistent with the analyses on the collected responses showing that the MR task led to overall more errors and slower RTs than the OR task. However, the absence of main effects or interactions with magnitude of rotation suggests that the missing data should not have a strong influence on our estimate of the viewpoint functions. Two dependent variables were assessed. First, accuracy for each subject was determined by calculating the number of correct trials in each condition as a proportion of the number of total trials included in the analysis. Second, mean RTs for each subject were calculated from trials on which subjects responded correctly.

fMRI Methods

Imaging was performed on a GE 1.5 T Signa (Milwaukee, WI) scanner with an ANMR (Advanced NMR, Wilmington, MA) resonant gradient echoplanar imaging system. All images were acquired using a standard quadrature head coil and a T2*-sensitive gradient-recalled single-shot echoplanar pulse sequence. Nine contiguous coronal T1-weighted slices parallel to the tentorium and covering the parietal and occipital lobes and the posterior portions of the frontal and temporal lobe were prescribed based on sagittal localizers acquired at the beginning of each scanning session. Functional T2*-weighted images were acquired at the same locations in runs of 1008 images (112 per slice). fMRI acquisition parameters were as follows: repetition time (TR) of 2250 ms, echo time of 60 ms, flip angle of 60°, acquisition matrix of 128 × 64, and field of view of 40 × 20 cm. Voxel resolution was 3.12 × 3.12 mm in plane and 7 mm thick (no skip).

fMRI Data Analyses

One run was excluded for one subject because of excessive motion (exceeding 0.5 of a voxel in any direction). The SPM96 algorithm was used to correct for motion between successive images in each run of the remaining data. The anatomical and BOLD images for each subject were then transformed into an average Talairach-like coordinate system (Talairach and Tournoux, 1988), using eight anatomical anchor points (AC, most posterior point of fourth ventricle, and the superior, inferior, anterior, posterior, left, and right most points on the cortical surface).

The percent signal change data (corrected for a linear drift during a run; see Skudlarski et al., 1999) were analyzed using a priori contrasts between specific conditions in software developed at Yale and implemented in Matlab. Maps of t values and percent signal change, both corrected for a linear drift in the signal, were generated. These maps were then spatially smoothed using a Gaussian filter with a half-width half-maximum (HWHM) value of 0.5 voxel. In group composite maps, conditions were compared by multiplying their percent signal change relative to a fixation baseline with contrast weights for each subject. Under the null hypothesis of no effect, the expected value for this contrast is equal to zero. We used a bootstrap method to assess the statistical significance of effects. A population distribution for each voxel was generated by calculating mean values of the contrast in which randomly chosen subsets of half of the subjects were assigned reversed weights. For example, for the comparison of MR versus OR, we assigned the real MR trials to the OR condition for a randomly selected half of the subjects. This was done 3000 times so as to generate for each voxel a distribution of expected activity under the assumption of no difference between conditions. The observed contrast (with the data for all subjects assigned to the conditions as experienced) was assigned a p value indicating the proportion of times that the observed contrast was more extreme than the randomized contrast. An area was considered active if the peak of activity reached $p < 0.001$ (uncorrected) *and* if the volume of contiguous activation at $p < 0.05$ was larger than 0.22 cm^3 (except for analyses reported in Table 1, where the threshold was dropped to $p < 0.10$ to explore the possibility of common areas for MR and OR). The maps were then overlaid on the average normalized anatomical images (because of artifacts in the T1 images for one subject, only 14 of 15 subjects were included in the anatomical composites).

References

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

Biederman, I. (1987). Recognition-by-components: a theory of human image understanding. Psychol. Rev. *94*, 115–147.

Biederman, I., and Gerhardstein, P.C. (1993). Recognizing depth-rotated objects: evidence and conditions for three-dimensional viewpoint invariance. J. Exp. Psychol. Hum. Percept. Perform. *19*, 1162–1182.

Biederman, I., and Gerhardstein, P.C. (1995). Viewpoint-dependent mechanisms in visual object recognition: reply to Tarr and Bülthoff (1995). J. Exp. Psychol. Hum. Percept. Perform. *21*, 1506–1514.

Bülthoff, H.H., and Edelman, S. (1992). Psychophysical support for a two-dimensional view interpolation theory of object recognition. Proc. Natl. Acad. Sci. USA *89*, 60–64.

Carpenter, P.A., Just, M.A., Keller, T.A., Eddy, W., and Thulborn, K. (1999). Graded functional activation in the visuospatial system with the amount of task demand. J. Cogn. Neurosci. *11*, 9–24.

Cohen, M.S., Kosslyn, S.M., Breiter, H.C., DiGirolamo, G.J., Thompson, W.L., Anderson, A.K., Brookheimer, 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.

Corballis, M.C. (1982). Mental rotation: anatomy of a paradigm. In Spatial Abilities: Developmental and Physiological Foundations, M. Potegal, ed. (New York: Academic Press), pp. 173–198.

Corballis, M.C., and McLaren, R. (1984). Winding one's ps and qs: mental rotation and mirror-image discrimination. J. Exp. Psychol. Hum. Percept. Perform. *10*, 318–327.

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., et al. (1998). A common network of functional areas for attention and eye movements. Neuron *21*, 761–773.

Coull, J.T., and Nobre, A.C. (1998). Where and when to pay attention: the neural systems for directing attention to spatial locations and to time intervals as revealed by both PET and fMRI. J. Neurosci. *18*, 7426–7435.

Folk, M.D., and Luce, R.D. (1987). Effects of stimulus complexity on mental rotation rate of polygons. J. Exp. Psychol. Hum. Percept. Perform. *13*, 395–404.

Gauthier, I., Tarr, M.J., Moylan, J., Skudlarski, P., Gore, J.C., and Anderson, A.W. (2000). The fusiform "face area" is part of a network that processes faces at the individual level. J. Cogn. Neurosci. *12*, 495–504.

Georgopoulos, A.P., Lurito, J.T., Petrides, M., Schwartz, A.B., and Massey, J.T. (1989). Mental rotation of the neuronal population vector. Science *243*, 234–236.

Grill-Spector, K., and Malach, R. (2001). FMR-adaptation: a tool for studying the functional properties of human cortical neurons. Acta Psychol. (Amst.) *107*, 293–321.

Hayward, W.G., and Tarr, M.J. (1997). Testing conditions for viewpoint invariance in object recognition. J. Exp. Psychol. Hum. Percept. Perform. *23*, 1511–1521.

Hayward, W.G., and Williams, P. (2000). Viewpoint dependence and object discriminability. Psychol. Sci. *11*, 7–12.

Hinton, G.E., and Parsons, L.M. (1981). Frames of reference and mental imagery. In Attention and Performance, Volume 9, A.D. Baddeley and J. Long, eds. (Hillsdale, NJ: Erlbaum), pp. 261–278.

Hochberg, J., and Gellman, L. (1977). The effect of landmark features on mental rotation times. Mem. Cognit. *5*, 23–26.

Hummel, J.E., and Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. Psychol. Rev. *99*, 480–517.

Humphrey, G.K., and Khan, S.C. (1992). Recognizing novel views of three-dimensional objects. Can. J. Psychol. *46*, 170–190.

Jolicoeur, P. (1985). The time to name disoriented natural objects. Mem. Cognit. *13*, 289–303.

Jolicoeur, P. (1990). Identification of disoriented objects: a dual-systems theory. Mind and Language *5*, 387–410.

Just, M.A., and Carpenter, P.A. (1976). Eye fixations and cognitive processes. Cognit. Psychol. *8*, 441–480.

Lawson, R., Humphreys, G.W., and Watson, D.G. (1994). Object recognition under sequential viewing conditions: evidence for viewpoint-specific recognition procedures. Perception *23*, 595–614.

Logothetis, N.K., Pauls, J., Bülthoff, H.H., and Poggio, T. (1994). View-dependent object recognition in monkeys. Curr. Biol. *4*, 401–414.

Logothetis, N.K., Pauls, J., and Poggio, T. (1995). Shape representation in the inferior temporal cortex of monkeys. Curr. Biol. *5*, 552–563.

Luna, B., Thulborn, K.R., Strojwas, M.H., McCurtain, B.J., Berman, R.A., Genovese, C.R., and Sweeney, J.A. (1998). Dorsal cortical regions subserving visually guided saccades in humans: an fMRI study. Cereb. Cortex *8*, 40–47.

Milner, A.D., and Goodale, M.A. (1995). The Visual Brain in Action (Oxford: Oxford University Press).

Perrett, D.I., Mistlin, A.J., and Chitty, A.J. (1987). Visual neurones responsive to faces. Trends Neurosci. *10*, 358–364.

Perrett, D.I., Oram, M.W., and Ashbridge, E. (1998). Evidence accumulation in cell populations responsive to faces: an account of generalisation of recognition without mental transformations. Cognition, *67*, 111–145.

Poggio, T., and Edelman, S. (1990). A network that learns to recognize three-dimensional objects. Nature *343*, 263–266.

Richter, W., Ugurbil, K., Georgopoulos, A., and Kim, S.-G. (1997). Time-resolved fMRI of mental rotation. Neuroreport *8*, 3697–3702.

Rock, I., Wheeler, D., and Tudor, L. (1989). Can we imagine how objects look from other viewpoints? Cognit. Psychol. *21*, 185–210.

Shepard, R.N., and Cooper, L.A. (1982). Mental Images and Their Transformations (Cambridge, MA: The MIT Press).

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

Skudlarski, P., Constable, R.T., and Gore, J.C. (1999). ROC analysis of statistical methods used in functional MRI. Neuroimage *9*, 311–329.

Suzuki, S., Peterson, M.A., Moscovitch, M., and Behrmann, M. (1997). Viewpoint specificity in the identification of simple volumetric objects (geons) is evident in control subjects and very exaggerated in visual object agnosia. Presented at the Fourth Annual Meeting of the Cognitive Neuroscience Society, Boston, Massachusetts.

Tagaris, G.A., Kim, S.-G., Strup, J.P., Andersen, P., Ugurbil, K., and Georgopoulos, A.P. (1996). Quantitative relations between parietal activation and performance in mental rotation. Neuroreport *7*, 773–776.

Tagaris, G.A., Kim, S.G., Strupp, J.P., Andersen, P., Ugurbil, K., and Georgopoulos, A.P. (1997). Mental rotation studied by functional magnetic resonance imaging at high field (4 Tesla): performance and cortical activation. J. Cogn. Neurosci. *9*, 419–432.

Talairach, J., and Tournoux, P. (1988) Co-planar Stereotaxic Atlas of the Human Brain (New York: Thieme).

Tarr, M.J. (1995). Rotating objects to recognize them: a case study of the role of viewpoint dependency in the recognition of three-dimensional objects. Psychon. Bull. Rev. *2*, 55–82.

Tarr, M.J., and Bülthoff, H.H. (1995). Is human object recognition better described by geon-structural-descriptions or by multiple-views? J. Exp. Psychol. Hum. Percept. Perform. *21*, 1494–1505.

Tarr, M.J., and Bülthoff, H.H. (1998). Image-based object recognition in man, monkey, and machine. Cognition *67*, 1–20.

Tarr, M.J., and Gauthier, I. (1998). Do viewpoint-dependent mechanisms generalize across members of a class? Cognition *67*, 71–108.

Tarr, M.J., and Kriegman, D.J. (2001). What defines a view? Vision Res. *41*, 1981–2004.

Tarr, M.J., and Pinker, S. (1989). Mental rotation and orientation-dependence in shape recognition. Cognit. Psychol. *21*, 233–282.

Ungerleider, L.G., and Mishkin, M. (1982). Two cortical visual systems. In Analysis of Visual Behavior, D.J. Ingle, M.A. Goodale, and R.J.W. Mansfield, eds. (Cambridge, MA: The MIT Press), pp. 549–586.

Wilson, K.D., and Farah, M.J. (1999). The neural basis of orientation normalization in object recognition. Soc. Neurosci. Abstr. 141.8.

Yuille, J.C., and Steiger, J.H. (1982). Nonholistic processing in mental rotation: some suggestive evidence. Percept. Psychophys. *3*, 201–209.