

Mental spatial transformations of objects and perspective

JEFFREY M. ZACKS¹, JON MIRES², BARBARA TVERSKY² and ELIOT HAZELTINE³

¹Washington University, Psychology Department, St. Louis, MO 63130-4899, USA (E-mail: jzacks@artsci.wustl.edu); ²Stanford University, USA; ³NASA Ames Research Center, USA

Received 17 December 2001; accepted 10 January 2002

Abstract. This study sought evidence for the independence of two classes of mental spatial transformation: object-based spatial transformations and egocentric perspective transformations. Two tasks were designed to selectively elicit these two transformations using the same materials, participants, and task parameters: one required same-different judgments about pairs of pictures, while the other required left-right judgments about single pictures. For pictures of human bodies, the two tasks showed strikingly different patterns of response time as a function of stimulus orientation. Moreover, across individuals, the two tasks had different relationships to psychometric tests of spatial ability. The chronometric and individual difference data converge with neuropsychological and neuroimaging data in suggesting that different mental spatial transformations are performed by dissociable neural systems.

Key words: cognition, cognitive neuroscience, individual differences, mental rotation, sex differences, spatial reasoning

1. Introduction

The ability to imagine and reason about changes of objects and their spatial layout is important both for everyday cognition (turning a combination lock, retracing one's steps at the grocery store) and for reasoning in technical domains (air traffic control, architecture). Mental spatial transformations can vary in their geometry, in the stimuli that give rise to them, and in their neural implementation. To retrieve what is on the other side of the grocery aisle, for example, one might imagine the aisle rotating or ourselves moving around the aisle. Two classes of mental transformation seem particularly important to human cognition: *object-based spatial transformations* and *egocentric perspective transformations*. Object-based transformations are imagined rotations or translations of objects relative to the reference frame of the environment. Egocentric perspective transformations are imagined rotations or translations of one's point-of-view relative to that reference

frame. Movements of external objects often have important consequences that require actions in response. Those actions in turn give rise to changes of one's egocentric perspective. The ability to predict the consequences of object motion and perspective change is thus important for everyday reasoning, and probably confers adaptive fitness.

Both object-based transformations and egocentric perspective transformations involve updating of the relationship between the environmental reference frame, the intrinsic reference frames of objects in the environment, and the observer's egocentric reference frame. However, different relationships are updated in the two transformations. In the case of object-based transformations, the relationships between the environmental and egocentric coordinate frames remains fixed, while each of their relationships with an object's intrinsic coordinate frame are updated. In the case of egocentric perspective transformations, the relationships between the environmental coordinate frame and those of the objects in the environment remain fixed, while each of their relationships with the observer's egocentric coordinate frame are updated.

One reasonable hypothesis is that these two classes of transformation require different computations, and so are implemented by two dissociable neural systems. However, the data on which they operate and the behavioral responses to which they provide input can often be quite similar, even identical. For example, imagine one sees a car with a missing headlight approaching head-on, and is asked whether the working headlight would be visible if the car were to pass the viewer from right to left. This problem could be solved by imagining the car rotating about an axis through its roof and floor (an object-based transformation) or by imagining one's self rotating about that same axis (an egocentric perspective transformation). Given these shared processing constraints it is possible that the two transformations are actually implemented by the same underlying functional system.

The data available to date seem consistent with the view that egocentric perspective transformations and object-based transformations are implemented by different processing systems. The neuropsychological literature provides evidence that these two classes of mental spatial transformation can be independently impaired (see Zacks et al. 1999 for a review). Lesions to right posterior cortex give rise to selective impairments in the ability to mentally rotate external figures (Ratcliff 1979), whereas lesions to left posterior cortex give rise to selective impairments at simple navigation tasks that seem to require the participants to imagine themselves turning, as in following a marked route on a map (Semmes et al. 1963). Recently, we have collected functional brain imaging data that suggest that in the normal brain, performing egocentric perspective transformations gives rise to a different

pattern of activation than performing object-based transformations with the same stimuli and responses (Zacks et al. 1999a; Zacks et al. 1999b).

These neurophysiological data are consistent with behavioral experiments in which participants are explicitly asked to imagine themselves rotating (viewer rotations) or an array of objects rotating (array rotations). Spatial judgments following such transformations differ systematically, suggesting that though the problems are formally equivalent, they are performed using computationally different procedures (e.g. Huttenlocher and Presson 1973; Presson 1982; Wraga et al. 2000). Similar dissociations have been found for paradigms in which participants imagine a character rotating in a scene, and are induced to take either the character's perspective, or an external "god's eye" perspective on the whole scene (Bryant and Tversky 1999).

If these two classes of spatial transformation are indeed performed by dissociable neural processing systems, other consequences follow. First, some problems will be more easily solved using one transformation than the other. Therefore it should be possible to construct situations in which one can selectively evoke one class of transformation or the other, and produce systematic effects on the chronometric properties of task performance. Second, some people may be more efficient at one transformation than the other, as the two processing systems should have freedom to vary independently.

1.1. *Chronometric profiles of object-based spatial transformations and egocentric perspective transformations*

Regarding the first consequence, there is indirect evidence that the two classes of transformation have different chronometric profiles. One half of the picture comes from the extensive literature on mental rotation (e.g. Shepard and Cooper 1982). In the paradigm developed by Shepard and his colleagues, observers judge whether a pair of asymmetric objects is identical or are mirror images (Shepard and Metzler 1971). From trial to trial, the orientation disparity between the two objects is varied randomly. The typical finding is that response time is a monotonic function of the orientation disparity between the two objects, resulting in a strong correlation between orientation and response time. This finding holds whether objects are rotated in the picture plane, in depth, or through oblique planes (Parsons 1987b). From this result it has been concluded that observers solve the problem by mentally rotating one of the objects as a whole into alignment with the other, and comparing the two images (but see Just and Carpenter 1985; Pylyshyn 1979). Mental rotation is a prototypical example of an object-based transformation, and this paradigm establishes a chronometric profile for its performance. Similar patterns have been observed for judgments of whether a character is normal or mirror-reversed (e.g. Cooper and Shepard 1973), for which direc-

frame. Movements of external objects often have important consequences that require actions in response. Those actions in turn give rise to changes of one's egocentric perspective. The ability to predict the consequences of object motion and perspective change is thus important for everyday reasoning, and probably confers adaptive fitness.

Both object-based transformations and egocentric perspective transformations involve updating of the relationship between the environmental reference frame, the intrinsic reference frames of objects in the environment, and the observer's egocentric reference frame. However, different relationships are updated in the two transformations. In the case of object-based transformations, the relationships between the environmental and egocentric coordinate frames remains fixed, while each of their relationships with an object's intrinsic coordinate frame are updated. In the case of egocentric perspective transformations, the relationships between the environmental coordinate frame and those of the objects in the environment remain fixed, while each of their relationships with the observer's egocentric coordinate frame are updated.

One reasonable hypothesis is that these two classes of transformation require different computations, and so are implemented by two dissociable neural systems. However, the data on which they operate and the behavioral responses to which they provide input can often be quite similar, even identical. For example, imagine one sees a car with a missing headlight approaching head-on, and is asked whether the working headlight would be visible if the car were to pass the viewer from right to left. This problem could be solved by imagining the car rotating about an axis through its roof and floor (an object-based transformation) or by imagining one's self rotating about that same axis (an egocentric perspective transformation). Given these shared processing constraints it is possible that the two transformations are actually implemented by the same underlying functional system.

The data available to date seem consistent with the view that egocentric perspective transformations and object-based transformations are implemented by different processing systems. The neuropsychological literature provides evidence that these two classes of mental spatial transformation can be independently impaired (see Zacks et al. 1999 for a review). Lesions to right posterior cortex give rise to selective impairments in the ability to mentally rotate external figures (Ratcliff 1979), whereas lesions to left posterior cortex give rise to selective impairments at simple navigation tasks that seem to require the participants to imagine themselves turning, as in following a marked route on a map (Semmes et al. 1963). Recently, we have collected functional brain imaging data that suggest that in the normal brain, performing egocentric perspective transformations gives rise to a different

pattern of activation than performing object-based transformations with the same stimuli and responses (Zacks et al. 1999a; Zacks et al. 1999b).

These neurophysiological data are consistent with behavioral experiments in which participants are explicitly asked to imagine themselves rotating (viewer rotations) or an array of objects rotating (array rotations). Spatial judgments following such transformations differ systematically, suggesting that though the problems are formally equivalent, they are performed using computationally different procedures (e.g. Huttenlocher and Presson 1973; Presson 1982; Wraga et al. 2000). Similar dissociations have been found for paradigms in which participants imagine a character rotating in a scene, and are induced to take either the character's perspective, or an external "god's eye" perspective on the whole scene (Bryant and Tversky 1999).

If these two classes of spatial transformation are indeed performed by dissociable neural processing systems, other consequences follow. First, some problems will be more easily solved using one transformation than the other. Therefore it should be possible to construct situations in which one can selectively evoke one class of transformation or the other, and produce systematic effects on the chronometric properties of task performance. Second, some people may be more efficient at one transformation than the other, as the two processing systems should have freedom to vary independently.

1.1. Chronometric profiles of object-based spatial transformations and egocentric perspective transformations

Regarding the first consequence, there is indirect evidence that the two classes of transformation have different chronometric profiles. One half of the picture comes from the extensive literature on mental rotation (e.g. Shepard and Cooper 1982). In the paradigm developed by Shepard and his colleagues, observers judge whether a pair of asymmetric objects is identical or are mirror images (Shepard and Metzler 1971). From trial to trial, the orientation disparity between the two objects is varied randomly. The typical finding is that response time is a monotonic function of the orientation disparity between the two objects, resulting in a strong correlation between orientation and response time. This finding holds whether objects are rotated in the picture plane, in depth, or through oblique planes (Parsons 1987b). From this result it has been concluded that observers solve the problem by mentally rotating one of the objects as a whole into alignment with the other, and comparing the two images (but see Just and Carpenter 1985; Pylyshyn 1979). Mental rotation is a prototypical example of an object-based transformation, and this paradigm establishes a chronometric profile for its performance. Similar patterns have been observed for judgments of whether a character is normal or mirror-reversed (e.g. Cooper and Shepard 1973), for which direc-

1.2. Individual differences in mental spatial transformations

The second consequence of the dissociability of the two classes of spatial transformation is that they should be free to vary independently across individuals. In particular, there should be individuals who perform tasks requiring object-based transformations quickly and efficiently but perform tasks requiring egocentric perspective transformations slowly and inefficiently, and there should be individuals who show the opposite pattern. Of course, demonstrating that this is the case requires more than simply showing that two performance measures fail to correlate perfectly across individuals. A more rigorous criterion is discriminant validity. Multiple putative measures of object-based transformations should correlate better amongst themselves than with putative measures of egocentric perspective transformations, and vice versa.

1.3. Two tests of the independence object-based transformations and egocentric perspective transformations

The experiment reported here was designed to test these two consequences of the proposed dissociation. We hypothesized that judging whether two pictures of bodies were identical or mirror-images would give rise to an object-based transformation, whereas judging arm of a single body was outstretched would give rise to an egocentric perspective transformation. We predicted, first, that this would lead to two different chronometric patterns for the two tasks. Second, we predicted that, across individuals, overall performance on these two tasks would be systematically related to paper-and-pencil tests thought to rely differentially on one spatial transformation or the other.

2. Method

2.1. Participants

Forty-eight Stanford undergraduates (26 female) participated in exchange for course credit or \$12.

2.2. Psychometric tests of spatial ability

Three tests were administered. The Mental Rotations Test (Vandenberg and Kuse 1978) requires participants to identify rotated versions of three-dimensional objects composed of cubes. We hypothesized this would assay object-based transformation ability, because test-takers overwhelmingly report that they solve the problems by mentally rotating the cube figures.

tion an object would be facing if upright, and for naming novel pictures of familiar objects (Jolicoeur 1985, 1988).

The other half of the picture comes from a study of left-right judgments. Consider a situation in which observers make handedness judgments about a single asymmetric object (e.g. a human body with one arm outstretched) appearing at varying orientations. One might imagine that such judgments would be performed by mentally rotating the object to a canonical orientation (e.g. upright) and then comparing it to a stored template. If this were the case for left-right judgments, one would expect to see a chronometric profile similar to that for the Shepard paradigm: response time monotonically increasing as a function of disparity between the orientation of the object and its canonical orientation, independent of the plane through which the object is rotated. However, Parsons (1987a) conducted such a study, with pictures of human bodies as the stimuli, and observed a different pattern. Chronometric profiles differed for different planes of rotation. In particular, for rotations in the picture plane, which give rise to robust correlations between orientation and response time in the Shepard paradigm, there was essentially no relationship in Parsons' study. Critically, when participants were explicitly asked to imagine themselves in the position of the depicted figure and then press a button, the time to do so correlated with the time required to make a left-right judgment about the same figure. Based on these results, Parsons argued that the observers performed the left-right judgments by imagining themselves in the position of the figures (an egocentric perspective transformation).

Putting these two results together, one can speculate that for picture plane rotations object-based transformations lead to a reliable monotonic relationship between orientation and response time, whereas imagined perspective transformations do not. However, several concerns preclude such a conclusion. Mental rotation studies are typically conducted with stimuli such as abstract figures (Just and Carpenter 1985; Pylyshyn 1979; Shepard and Metzler 1971), while Parsons's (1987a) study used pictures of human bodies, and other studies of left-right judgments have used body parts (Cooper and Shepard 1975; Parsons 1994; Parsons et al. 1995; Sekiyama 1982). The different experiments also differ in the participants involved and in incidental features of the designs. To settle the issue, a direct comparison of the two paradigms using the same materials, task parameters, and participants is required. If object-based transformations and egocentric perspective transformations are performed by dissociable systems and the two chronometric patterns described here reflect differences between those two systems, those two patterns should be observed in such a direct comparison.

The other two tests were selected to tap into participants' ability to perform egocentric perspective transformations. The Map-reading Test requires participants to follow a route on a map and indicate whether each of 36 marked turns is a right or left turn (Money and Alexander 1966). Importantly, they cannot move the map while performing the task. Thus, each turn requires a mental updating of the participant's imagined egocentric perspective. Various clinical populations are impaired at performing this task, but neurologically normal adults typically score very well. To increase the sensitivity of the test for normal adults we sped the task, forcing participants to complete it in 20 sec. The Perspective-taking Test (Hegarty and Kozhevnikov 1999) presents participants with a picture of an array of objects. With the array in view, they are asked to imagine themselves standing at one object, facing a second, and indicate the angle to a third object. Each item consists of a circle with a line drawn from the center to the top of the circle. The center is marked with the object they are to imagine themselves standing at, the top is marked with the name of the object they are to imagine themselves facing, and the participant is asked to indicate the angle to the third object by drawing another line from the center of the circle. Participants were given four minutes to complete the 10 items on the test. This test explicitly asks participants to imagine an egocentric perspective transformation.

2.3. Chronometric spatial reasoning tasks

Participants performed two computer-administered tasks, run on a Macintosh computer with the PsyScope experimental package (Cohen et al. 1993). Both tasks involved spatial judgments about line drawings of male bodies with one arm extended (see Figure 1). The *same-different task* was similar to the widely-used paradigm developed by Shepard and his colleagues (Shepard and Metzler 1971). Participants saw pairs of line drawings of human bodies that differed in orientation and judged whether they were identical or mirror images. The pictures were arranged one above the other, and the upper picture (the "standard") always appeared at an upright orientation. The lower picture (the "test") appeared at a picture plane orientation that varied randomly in 30-degree increments from -180 degrees to 180 degrees. In the *left-right task*, only one ("test") picture was presented and participants judged whether it had a left or right arm extended. Its orientation was varied randomly in 30-degree increments from -180 degrees to 180 degrees (clockwise).

Before beginning a block of trials, participants were instructed to position their index fingers over two buttons on a button box. In the same-different task they were told to press the left button for "same" and the right button for "different." The buttons were marked accordingly. In the left-right task




Figure 1. Examples of pictures used in the chronometric tasks. The first (leftmost) depicts an upright (0 degree) figure with his left arm extended straight from the body. The second depicts a figure rotated -150 degrees, with his right arm extended straight from the body. The third depicts a figure rotated 60 degrees, with his left arm crossed across the body. The fourth depicts a figure rotated -30 degrees, with his right arm crossed across the body.

participants were told to press the left button for "left" and the right button for "right."

For both tasks, each trial began with a text cue presented at the center of the screen instructing the participant to "hit any button to go on." When a button was pressed, a cross appeared at the center of the screen for 1500 ms, followed by a stimulus. The stimulus remained on the screen until a button was pressed, ending the trial. The response time and button pressed were recorded. No feedback was given.

In both tasks, two poses of the body were used. In one pose, the outstretched arm was extended laterally away from the body. In the other pose, the outstretched arm was crossed over the chest (Figure 1). This was done because pilot testing had indicated that varying the visual appearance of the stimuli discouraged rote-learning strategies for performing the task. In the same-different task, the same pose was always used for both pictures.

In neither case was any instruction provided regarding how the problem should be solved. Participants were asked to respond as quickly as possible while remaining accurate. All trials were analyzed; i.e., no practice trials were given.

2.4. Procedure

After providing informed consent, participants were given written instructions for performing the same-different and left-right tasks, and any questions were answered. They then performed 96 trials of each task. For the same-different task, this allowed presentation of all possible combinations of orientation (12), match (2), handedness of upper picture (2) and pose (2). For

the left-right task, this allowed two presentations of each possible combination of orientation, (12), handedness of the picture (2), and pose (2). (Due to a bug in the experimental software, there was a slight departure from randomization in the trial order, such that in the last trials for each task one handedness occurred disproportionately often. None of the participants noticed this, and there were no indications that response times or error rates in these trials differed from the preceding trials.) Task order was counterbalanced across participants.

After each participant performed the chronometric tasks, the three spatial ability tests were administered: first the Mental Rotations Test (Vandenberg and Kuse 1978), then the modified Map-reading Test (Money and Alexander 1966), then the Perspective-taking Test (Hegarty and Kozhevnikov 1999).

Participants were run individually or in pairs. The psychometric tests were administered individually.

3. Results

3.1. Spatial ability test scores

Summary statistics for the three spatial ability tests are given in Table 1. In the Mental Rotations Test, participants identified rotated versions of sample figures; each correct identification is scored as a positive point, and each false alarm as a negative point. Possible scores range from -40 to 40, and were normally distributed over the positive values. A participant's score on the Map-Reading test is the total number of correctly-identified turns shown on the map. Scores covered most of the possible range (0-32), and were normally distributed. Scores on the Perspective-taking test are the average unsigned angular deviation from the correct answer, subtracted from 180 degrees. (Subtracting from 180 degrees produces scores that are higher for better performance, as for the other tests.) Thus, scores can range from 0 to 180; in this sample, scores were quite high and negatively skewed, indicating a restriction of the range of the test.

There were moderate correlations amongst the three tests (see Table 2).

3.2. Chronometric tasks

Accuracy was high overall, 96%, and similar between the same-different task (95%) and the left-right task (97%). Accuracy was stable across trials of each task, increasing an average of 1.4% from the first 24 trials of each task to the last 24. For all analyses of response time, only correct trials were used.

The primary variable of interest for the chronometric tasks considered by themselves is the relationship between orientation and response time. The

Table 1. Summary statistics for the three spatial ability tests

Test	Minimum	Maximum	Median	Mean	Standard deviation
Mental Rotations	0	38	17.0	18.48	8.76
Map-reading	3	32	16.5	17.80	6.90
Perspective-taking	119	173	164.0	159.10	13.00

Table 2. Correlations amongst the spatial ability tests. All were reliably greater than zero, min. $t(46) = 2.8, p < 0.007$

	Mental Rotations	Map-reading
Map-reading	0.39	
Perspective-taking	0.42	0.38

pattern observed for the two tasks is quite clear (see Figure 2). For the same-different task, response time increased monotonically with orientation. For the left-right task, there was no relationship between orientation and response time. This difference was tested by calculating mean response times for each participant for each combination of response time and orientation and entering these into a 2 (task) by 12 (orientation) repeated measures analysis of variance (ANOVA). There was a reliable task-by-orientation interaction, $F(11,517) = 16.6, p < 0.001$. The main effects of task and orientation were also reliable, $F(1,94) = 19.8, p < 0.001$ and $F(11,517) = 17.1, p < 0.001$, respectively.

The particular relationship between orientation and response time in each task was quantified by calculating, for each participant in each task, the correlation between the absolute value of the orientation of the test picture and response time. For the same-different task, these correlations were positive and robust, mean $r = 0.321, t(47) = 13.4, p < 0.001$. For the left-right task, there was no correlation, mean $r = -0.030, t(47) = -1.07, p = 0.29$. The difference between the two was statistically reliable, $t(47) = 8.86, p < 0.001$. (All reported t -tests are two-tailed.)

Alternatively, one can consider the slope of the regression line predicting response time from orientation. Reliably positive slopes indicate a relationship between orientation and response time. The results are equivalent. For the same-different task, the mean slope was 4.53 ms/degree, $t(47) = 10.8, p < 0.001$. For the left-right task, the mean slope was virtually zero: -0.339




Figure 2. Relationship between orientation and response time in the chronometric tasks.

ms/degree, $t(47) = -0.80$, $p = 0.42$. The difference between the two slopes was reliable, $t(47) = 7.32$, $p < 0.001$. The mean of the intercept for the regression line was 1438 ms for the same-different task and 1329 ms for the left-right task; this difference was not statistically reliable, $t(47) = 0.85$, $p = 0.40$.

3.3. Relationship between chronometric tasks and spatial tests

We calculated the correlation between each of the chronometric tasks and the three spatial tests. These correlations are given in Table 3. In general, test performance was correlated with speed in the chronometric tasks. This was statistically reliable for all the correlations except that between score on the Mental Rotations Test and the left-right task, $t(46) = -1.53$, $p = 0.14$. For the other five correlations, the smallest $t(46)$ was -2.26 , $p < 0.029$.

Following the method described by Meng, Rosenthal, and Rubin (1992), we tested the hypothesis that overall response time on the same-different task would be better predicted by the Mental Rotations Test than the Map-reading Test or the Perspective-taking test. Inspection of Table 3 shows the correlation coefficients were as predicted, and a planned contrast found this pattern to be statistically reliable ($p = 0.025$). We also tested the hypothesis that overall response time on the left-right task would be better predicted by the Map-reading Test and the Perspective-taking test than the Mental Rotations test. Again, inspection of Table 3 shows the correlation coefficients were as predicted; however, in this case the appropriate planned contrast was not statistically reliable ($p = 0.15$).

Table 3. Correlations between spatial ability tests and chronometric tasks (Higher test scores and lower chronometric task speeds indicate better performance, so the observed negative correlations indicate that faster participants had higher test scores)

Chronometric task speed	Same-Different Left-right	Spatial ability test score		
		Mental Rotations	Map-reading	Perspective-taking
		-0.60	-0.32	-0.35
		-0.22	-0.46	-0.36

Table 4. Performance on tests of spatial ability and chronometric tasks by sex (Standard errors in parentheses)

Spatial ability test score	Mental Rotations* Map-reading* Perspective-taking	Female mean		Male mean	
		Mean	SE	Mean	SE
Chronometric task speed	Same-different Left-Right	1836 ms (113 ms)	1387 ms (104 ms)	1850 ms (173 ms)	1189 ms (93.0 ms)
Chronometric task error rate	Same-different Left-right	4.41% (0.71%) 4.25% (0.97%)		4.99% (1.22%) 2.10% (0.59%)	

*The two sexes differed, $p < 0.05$.

3.4. Sex differences

Table 4 gives mean scores on the spatial ability tests and mean response time and error rate in the chronometric tasks, broken down by sex. There were reliable sex differences for the Mental Rotations Test ($d = 0.81$) and the Map-reading Test ($d = 0.67$), both favoring males. There were no reliable sex differences for the Perspective-taking Test or either of the chronometric tasks.

While there was no sex difference in overall response time for either chronometric task, there was a trend toward a difference in the mental rotation slope for the same-different task. Males rotated faster than females (mean 3.68 ms/degree for males, mean 5.26 ms/degree for females) and this difference approached statistical reliability, $t(46) = 4.6$, $p = 0.06$. However, there was clearly no difference in overall response time ($t(46) = -0.07$, $p = 0.95$), indicating that the function relating orientation to response time had a lower (faster) intercept for females in addition to a steeper (slower) slope.