

Research Report

PERCEIVING REAL-WORLD VIEWPOINT CHANGES

Daniel J. Simons¹ and Ranxiao Frances Wang²

¹Harvard University and ²Cornell University

Abstract—Retinal images vary as observers move through the environment, but observers seem to have little difficulty recognizing objects and scenes across changes in view. Although real-world view changes can be produced both by object rotations (orientation changes) and by observer movements (viewpoint changes), research on recognition across views has relied exclusively on display rotations. However, research on spatial reasoning suggests a possible dissociation between orientation and viewpoint. Here we demonstrate that scene recognition in the real world depends on more than the retinal projection of the visible array; viewpoint changes have little effect on detection of layout changes, but equivalent orientation changes disrupt performance significantly. Findings from our three experiments suggest that scene recognition across view changes relies on a mechanism that updates a viewer-centered representation during observer movements, a mechanism not available for orientation changes. These results link findings from spatial tasks to work on object and scene recognition and highlight the importance of considering the mechanisms underlying recognition in real environments.

Models of object and scene recognition largely fall into two groups, those positing view-dependent recognition and those arguing for view-independent recognition. Models predicting view dependence typically argue that each object is represented as one or more distinct views and recognition occurs by aligning current sensations to a stored view (Tarr & Pinker, 1989; Ullman, 1989) or by interpolating between stored views (Edelman & Bühlhoff, 1992). Models predicting view independence typically argue that objects are stored as structural descriptions and recognition occurs by accessing the structural description from any view that allows identification of the parts and their relations (e.g., Biederman, 1987). Recognition performance generally supports structural description models when objects are composed of distinctive, identifiable parts (Biederman, 1987; Biederman & Gerhardstein, 1993; Cooper, Biederman, & Hummel, 1992; Hummel & Biederman, 1992). Recognition performance is generally view dependent with stimuli that are not easily distinguishable by their parts, such as wire-frame objects (Tarr & Pinker, 1989) and blob-like objects (Edelman & Bühlhoff, 1992), and with highly overlearned stimuli such as letters (Corballis, Zbrodoff, Shetzer, & Butler, 1978). Recognition of spatial layouts of groups of objects also seems to be view dependent (Diwadkar & McNamara, 1997; see also Simons, 1996); observers are progressively slower to recognize the spatial layout of an array of objects with increasing changes in the orientation of the array (Diwadkar & McNamara, 1997).

Proponents of both positions typically study object recognition by presenting images that simulate rotation in depth or in the picture

plane in front of a stationary observer. Observers are asked to determine if an object is the same before and after a rotation (i.e., a same/different or old/new recognition task). If the latency increases or the accuracy decreases as the difference in orientation between the studied and tested view increases, recognition is taken to be view dependent. In contrast, if performance is relatively unaffected by the change in view, recognition is taken to be view independent. Although this approach has produced a number of insights into the structure of object representations, it has neglected a critical distinction: In the real world, the view of an object (its retinal projection) can change either because of object motion (e.g., the display rotations typically studied) or because of the movements of the observer. In fact, most real-world view changes are caused by observer movements and not by object rotations; people often move their heads and bodies, causing changes to their view of an object, but objects generally do not rotate in space in front of people. Although a change in the retinal projection due to object rotation may be equivalent to one produced by an observer moving to a different viewpoint, the underlying mechanisms for object recognition may be different. That is, object recognition in the real world may depend on more than just the retinal projection of a static image.

Researchers studying the development of spatial representations and spatial reasoning have long distinguished between observer movement and display rotation (e.g., Huttenlocher & Presson, 1973, 1979; Rieser, 1989; Rieser, Garing, & Young, 1994). For example, Huttenlocher and Presson (1973, 1979; Presson, 1982) asked subjects to imagine themselves moving to a novel observation point or to imagine a display rotating. Following these imaginary changes, subjects were asked to identify which object would now be to their left (or other specified directions). Interestingly, the task was easier after imagining self-motion than after imagining array rotation (with a different task, performance was better following imagined rotation).

More direct evidence for a distinction between performance following display rotation and performance following observer movement comes from studies demonstrating the importance of the actual position and orientation of an observer in spatial reasoning tasks. Although both young children and adults have difficulty imagining themselves facing a different direction and pointing to where hidden surrounding objects would be, the task becomes trivially easy when they have physically moved to the new observation point (Easton & Sholl, 1995; Huttenlocher & Presson, 1973; Rieser, 1989). In fact, observers can point more accurately following physical movement even when the target environment is entirely imagined from the beginning (Rieser et al., 1994).

Given that acting on the environment requires a viewer-centered representation, these findings are perfectly reasonable. In order to complete an action, observers must somehow update their spatial representations to accommodate changes in their body position and orientation. That is, the positions of target objects relative to the observer should be continuously adjusted to reflect the correct relationship. By updating their spatial representations as they move, observers can accurately interact with their surroundings from

Address correspondence to Daniel J. Simons, Department of Psychology, Harvard University, 820 William James Hall, 33 Kirkland St., Cambridge, MA 02138; e-mail: dsimons@wjh.harvard.edu. Sample stimuli will be available at the following World Wide Web address: <http://www.wjh.harvard.edu/~dsimons>.

Perceiving Real-World Viewpoint Changes

novel viewing positions. The same principles that apply in spatial reasoning problems may also be relevant to understanding the mechanisms underlying the recognition of objects or arrays of objects across views. In other words, a person's ability to recognize objects following changes to their orientation (e.g., simulated rotations) may not reflect his or her true ability to recognize objects following changes in viewing position in the real world. Recognition may be unaffected by changes to the observer's viewing position even when comparable view changes caused by display rotation produce view-dependent performance.

To examine this hypothesis, we used layouts of familiar objects on a large circular table. In order to demonstrate that recognition is accurate across shifts in the observer's viewing position, it is necessary to use displays that produce view-dependent recognition following display rotation. Otherwise, differences between viewpoint and orientation changes would be undetectable. Recent studies using spatial layouts of objects as stimuli have found view-dependent recognition performance across display rotations (Diwadkar & McNamara, 1997; see also Christou & Bühlhoff, 1997, for consistent evidence with virtual reality scenes). These studies also suggest an important parallel between recognition performance with spatial layouts of objects and with individual objects. Although studies of recognition of spatial layout may not directly constrain theories of individual object recognition, the intrinsic similarity of the two tasks and the striking parallels between the patterns of results suggest similar mechanisms for layout and object representation.

The task used in our studies draws on the methodology of change detection. In a typical recognition task, observers study a small set of objects or layouts and then at test try to determine whether a new instance matches a studied one. By asking subjects to perform a change detection task ("which object moved") rather than an old/new judgment task, we effectively increased the number of studied items; subjects viewed a new layout on each trial rather than a small set of layouts at the beginning of the task. As a result, subjects stored only a single view of each tested layout. This change detection task is fundamentally the same as an old/new recognition task in which the new layouts are slightly changed versions of the studied targets. The only difference is that observers must identify the change.

If object and scene recognition rely primarily on the retinal projection of a layout of objects, performance in this task should be identical regardless of whether a view change is caused by display rotations or observer movements. Therefore, performance differences between these conditions would suggest the need for an additional mechanism to account for real-world observer movements. Here we present three experiments that explored differences between orientation changes and viewpoint changes and their effects on the recognition of spatial layouts of objects. The first experiment was designed to pit changes in the observer's position (viewpoint changes) directly against display rotations (orientation changes).

EXPERIMENT 1

Method

Participants

Twenty-four undergraduates at Cornell University voluntarily participated in the experiment in exchange for course credit. All were informed of their rights as experimental participants.

Apparatus and procedure

Our apparatus was similar to that used by Diwadkar and McNamara (1997): Five familiar objects (brush, mug, goggles, stapler, and scissors) were randomly assigned to any of nine possible positions on a circular table 1.22 m in diameter (see Fig. 1). On each trial, the observer viewed a layout of the five objects for 3 s. The table was then occluded by a curtain for 7 s, and during the occlusion interval, one of the five objects was moved to a new position on the table. Following this retention interval, the curtain was raised, and the observer was asked to identify which of the five objects had moved. Participants were divided into two groups. Half the observers moved to a different viewing position during the retention interval of each trial, and half always viewed the array from the same viewing position on each trial. Observers who changed viewing positions simply walked to another chair located exactly 47° to the left or the right of the original viewing position (calculated from the center of the table).¹ As a control for the possibility that the act of walking from one position to another might affect the results, observers who viewed the array from the same viewing position before and after the retention interval walked halfway to the other viewing position and then returned to the initial position on each trial.




Fig. 1. Schematic illustration of the apparatus. Objects displayed on the table were a brush, a mug, goggles, a stapler, and scissors (represented by the symbols in the figure). Half of the subjects started each trial at Viewing Position 1, and half started at Viewing Position 2 (see Method for a complete description of the viewing conditions).

1. Subjects typically stood up, turned to face the other viewing position, walked four to five steps in a linear path, turned to face the display, and sat in a chair facing the center of the array. Thus, their motion included both a translation and a rotation first away from and then toward the display.

	Display Rotates (20 trials)	Display Doesn't Rotate (20 trials)
Viewpoint Condition (12 observers)	Same View	Different View (47° Viewpoint Change)
Orientation Condition (12 observers)	Different View (47° Orientation Change)	Same View

Fig. 2. Experimental conditions and the resulting retinal projections.

For all participants, the table rotated 47° on half of the trials. Thus, observers who changed viewing positions received the same view whenever the table rotated because they experienced a viewpoint shift and a compensating orientation change. When the table was stable, these observers experienced a viewpoint shift of 47°. In contrast, observers who maintained the same viewing position on each trial

experienced an orientation change of 47° when the table rotated, but an identical view when it was stationary (see Fig. 2).

Prior to each trial, observers in both conditions were told whether or not the table would rotate, thereby informing them if their view of the table would be the same or rotated. Observers were aware of the size and direction of the view change, which were constant on every different-view trial. Prior to the test trials, observers were given practice with the task to familiarize them with the objects, the viewing positions, and the magnitude of the view change. Each observer then viewed 40 trials, half of which involved either an orientation or a viewpoint change of 47° (depending on the condition to which they were randomly assigned) and half of which provided the same view before and after the delay. The positions of objects on the table and the direction of the observer's movement were fully counterbalanced across the four conditions.

Results and Discussion

This experiment was designed to pit the ability to detect layout changes when the orientation of the layout changed against that ability when the viewpoint of the observer changed. Consistent with previous findings of viewpoint dependence in the recognition of spatial layouts of objects (Diwadkar & McNamara, 1997; Simons, 1996), changing the orientation of the display in front of a stationary observer significantly disrupted the detection of changes; observers in the orientation-change condition identified the moved object significantly less accurately when

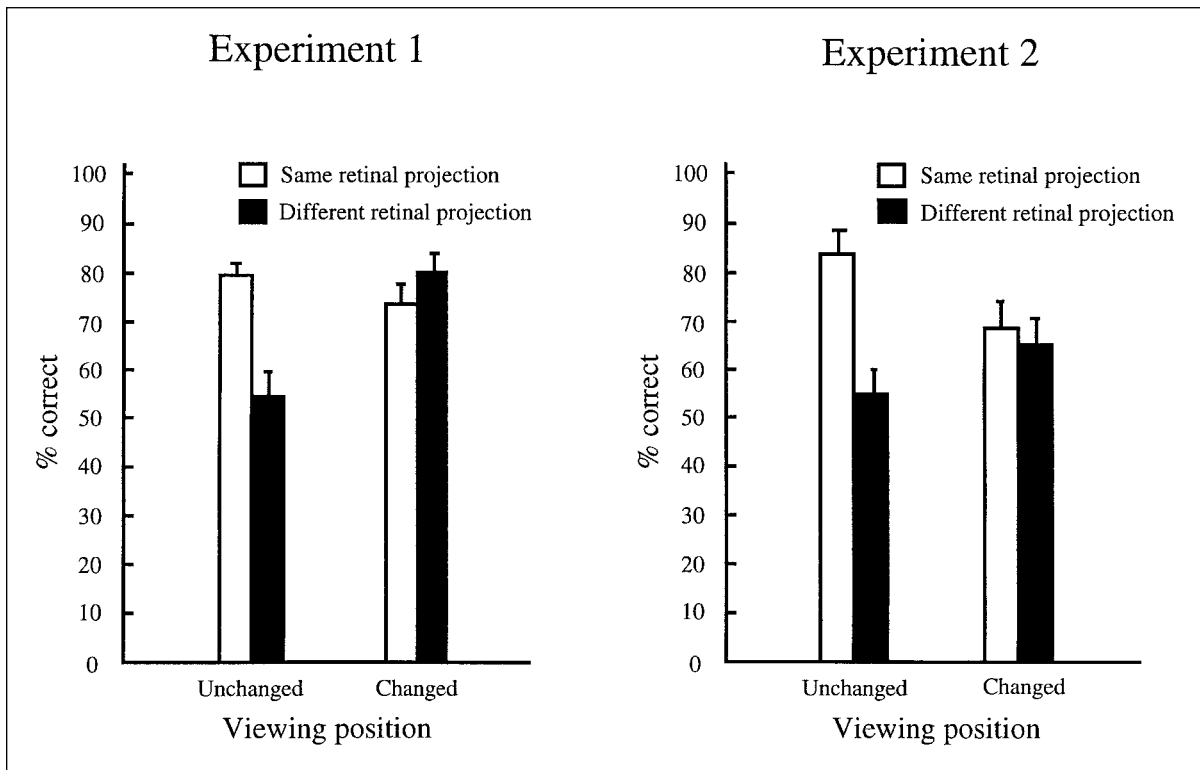


Fig. 3. Results of Experiments 1 and 2. Columns represent the percentage of correct responses. Error bars indicate standard errors. Observers who did not change viewing position during each trial received a different retinal projection when the table rotated and the same retinal projection when the table did not rotate. Observers who changed viewing position during each trial received the same retinal projection when the table rotated to match their movement and a different retinal projection when the table did not rotate.

Perceiving Real-World Viewpoint Changes

the orientation of the layout changed (i.e., the table rotated) than when they received an identical view before and after the retention interval, $t(11) = 5.9, p = .0001$ (see Fig. 3).

The retinal projections of orientation changes and viewpoint changes were equivalent in this study. Therefore, if recognition across views is based solely on the retinal projection, performance should have been equivalent for orientation and viewpoint changes. Yet, in contrast with the results for the orientation condition, viewpoint changes had no effect on accuracy in identifying the moved object. Observers in the viewpoint-change condition were equally accurate when they received a changed view (i.e., they moved to a different viewing position) and when they received an identical view (i.e., the table rotated to compensate for the shift in viewing position), $t(11) = 1.3, p = .22$. These findings suggest that different mechanisms underlie scene recognition across orientation changes and viewpoint changes. Even with displays that produce orientation dependence, observers are able to detect changes to the layout across shifts in their viewing position. Note that in both conditions, observers knew in advance of each trial whether the view would change and by how much; they were familiar with the individual viewing positions and the magnitude of the rotation.² In fact, the experimenters simply alternated trials with and without view changes, informing the subject in advance which sort of trial was next.

One potential factor complicating interpretation of the results of the first experiment was the presence of background landmarks in the room with the table. The presence of these landmarks meant that the two viewing positions in the viewpoint-change condition had additional indications of the magnitude of the change. Perhaps these extra landmarks somehow facilitated the alignment of the current view to a view-dependent representation. Alternatively, observers may have coded the relations of the objects on the table relative to other landmarks in the room, thereby establishing an environment-centered representation that would be unaffected by observer movements. In order to examine these possibilities more fully, we conducted a second experiment that eliminated background landmarks.

EXPERIMENT 2

Method

The procedure and materials were identical to those of Experiment 1 with one important exception. In this experiment, background landmarks were eliminated by turning off the room lights and coating the objects with phosphorescent paint. Twenty-four undergraduates at Cornell University participated in exchange for course credit, with half assigned to the viewpoint-change condition and half to the orientation-change condition.

Results

Eliminating the background landmarks had no effect on the pattern of results (see Fig. 3). Observers in the orientation-change condi-

2. It is possible that observers in the orientation-change condition had less information for the magnitude of the change than did subjects in the viewpoint-change condition. Ongoing studies are examining the possibility that allowing observers to control the rotation of the display will improve performance.

tion were significantly less accurate with a rotated view than with the same view, $t(11) = 7.0, p < .0001$. Observers in the viewpoint-change condition were equally accurate when they received a changed view (i.e., they walked to a different viewing position) and when they received the same view (i.e., the table rotated to compensate for their movement), $t(11) = 1.0, p = .339$. Even when orientation changes and viewpoint changes produced the same changes to the retinal projection, accuracy was disrupted only when objects changed orientation in front of a stationary observer. Performance was unaffected when view changes were caused by shifts in the observer's viewing position.

Discussion

As noted earlier, these results are not predicted by traditional models of object recognition, which do not distinguish between viewpoint changes and orientation changes. Although mechanisms like the mental alignment of the current image with a stored template or interpolation from multiple stored views may explain recognition performance when objects move in front of a stationary observer, additional mechanisms are needed to account for observers' ability to recognize objects in typical, real-world situations in which they view objects from different locations.

Two quite different mechanisms could account for why recognition is relatively independent of changes in viewing position. One possibility, suggested by research on spatial reasoning discussed earlier, is that observers form a representation of the objects in a scene relative to a larger spatial framework (Huttenlocher & Presson, 1973, 1979; Presson, 1982). In an environment-centered representation, observer position would not affect recognition. Thus, the lack of a disruption to recognition in our experiments may have resulted from the formation of a representation that placed the entire spatial array into the larger framework of the experimental room. Such a representation would have been unaffected by changes in observer position because the relation between the array and the surrounding environment was unchanged. Display rotations, in contrast, would have disrupted the relationship between the room and the display, thereby reducing recognition accuracy. This explanation may account for the results of the first experiment, but it is weakened by those of the second experiment. To code spatial layout into a larger spatial framework in Experiment 2, subjects would have had to use an imagined room; none of the larger spatial framework was visible during the trials.

Alternatively, observers may represent only the relation between objects in the array and their own position. This viewer-centered representation could then be updated by a mechanism that automatically takes visual, vestibular, or proprioceptive information into account to adjust the representation for changes in observer position. Consistent with prior research showing that accurate perception of self-motion can be achieved with vestibular and proprioceptive information (e.g., Berthoz, Israel, Francois, Grasso, & Tsuzuku, 1995; Loomis et al., 1993; Sholl, 1989), Experiment 2 suggests that if such an updating mechanism exists, it does not depend exclusively on visual information to accommodate the change in observer position. Even without such information, viewpoint changes caused no disruption to identification of changes in spatial layouts. Experiment 3 was conducted to examine the effect of eliminating visual, vestibular, and proprioceptive information during changes in viewing position on the ability to update representations across such changes. If objects in the display

are coded with respect to a larger, environment-centered spatial framework, then observers should not be disrupted as long as they can map their own position onto that larger framework. If, however, an updating mechanism accounts for the accurate recognition across viewpoint changes, then eliminating feedback during the change in observer position should disrupt performance.

EXPERIMENT 3

Method

The procedure was similar to the procedures of the first two experiments with several important exceptions. As in Experiment 1, the lights were on in the testing room so that observers could see the larger framework of the room. In this experiment, all 12 participants were in the viewpoint-change condition. On each trial, observers started at one viewpoint, sitting in a wheeled chair. After they viewed the array, the curtain was lowered, and observers closed and covered their eyes. An experimenter then wheeled them to the other viewing position while spinning the chair rapidly. The result of this manipulation was that observers could not sense their direction of motion. However, after they were stopped, they opened their eyes and could readily determine their position in the room. Furthermore, throughout the study, they always moved from the original position to the same ending position during each trial. That is, they knew the locations of the starting and ending points within the room and had experienced each viewing position. When they reached the new position, the curtain was raised, and the observers opened their eyes and responded.

Results and Discussion

Unlike in the first two experiments, observers in this experiment were less accurate when they received a different view (caused by the observer moving) than when they received the same view (i.e., when the table rotated to compensate for the change in their viewing position): Observers showed a consistent decline in accuracy in the different-view condition ($M = 64.6\%$ correct) relative to the same-view condition ($M = 72.5\%$). The mean difference between these conditions ($M = 7.92\%$, $SE = 2.08\%$) was significant, paired $t(11) = 3.8$, $p = .003$. Of the 12 participants, 9 were more accurate when receiving the same view, and none were more accurate with a different view (3 were equally accurate in the two conditions), $Z = 2.694$, $p = .007$ (by a Wilcoxon test). Although observers were exposed to the experimental room, were given the opportunity to place the display into the larger scene context of the room, were familiar with each viewing position, and knew the direction and extent of their motion through the room, accuracy was diminished by interfering with processing as observers changed viewing positions.

GENERAL DISCUSSION

The three experiments in this report suggest that view changes caused by rotating an array are not equivalent to those caused by observer movement. Although display rotations produce viewpoint-dependent recognition, observer movements apparently do not. When observers are given sufficient information to update their position with respect to the array, recognition appears to be independent of the observ-

ers' viewpoint. By investigating the sorts of view changes that observers typically experience, our experiments highlight a potentially important difference in layout recognition in these two conditions. Whereas display rotations require effortful processing and sometimes produce viewpoint dependence in recognition, movements of the observer are less disruptive. What differences in mechanisms underlie this distinction?

Although our experiments cannot conclusively pinpoint the processes underlying the relative viewpoint independence in the viewpoint-change condition, they do suggest some mechanisms that may be involved and others that probably are not. Although the results of Experiment 1 fit well with the hypothesis of a represented link between the display and the larger reference frame of the room, the results of Experiments 2 and 3 are problematic for the notion of an environment-centered representation. According to such models, disruptions to the formation of a mapping between the display and the larger reference frame should lead to a decline in accuracy. Yet in Experiment 2, in which displays were viewed in the dark, observers were not disrupted by a view change even though they were not given the opportunity to form a representation of the relationship between the display and the larger reference frame of the room. Furthermore, the hypothesized link between the display and the larger reference frame should be independent of the processes involved in getting from one viewing position to another. If observers form a link between the layout and the room, they should show equally accurate recognition from any position in that room; updating of spatial position over time should play no role in recognition. However, the results of Experiment 3 show that even when observers are given the opportunity to form a link between the display and the larger reference frame, the view change affects recognition if observers are not able to update their position over time.

Our results suggest that recognition of spatial layout across view changes caused by observer movement depends, at least partially, on updating the representation as the observer moves. A critical topic for future research is what information observers use to perform this updating. Experiment 2 demonstrated that visual information alone is not necessary to support updating. Experiment 3 eliminated visual, vestibular, and proprioceptive information, producing a significant disruption to recognition across view changes. Taken together, these results suggest that vestibular information, proprioceptive information, or both play a role in the updating process. Observers appear to form a viewer-centered representation upon first viewing a spatial layout. During a display rotation, they do not have information about the view change, so they cannot automatically update their representation. As a result, they show view-dependent detection. In contrast, during a viewing position change, observers have other information specifying the change. Presumably they can use these other sources of information to adjust or update their representation for their own shift in position. When they reach the new viewing position, they still have a viewer-centered representation of the layout, but it has been modified and now corresponds to their current viewing position. Further studies are needed to tease apart the separate contributions of visual, vestibular, and proprioceptive information to this updating mechanism.

In sum, these studies point to the importance of considering the typical behaviors of an organism in the study of representations. By limiting studies of recognition across views to changes in display orientation, previous studies neglected a potentially important component of object recognition—the updating of representations to compensate for changes in viewing position. Although both display rotations and observer movements can produce equivalent changes to

Perceiving Real-World Viewpoint Changes

the projection of a display on the retinas, the behaviors causing the changes may lead to strikingly different mental representations. Our data suggest that observers form a representation of a spatial layout that is dependent on observer position, but provided that sufficient information is available, they can flexibly adjust or update their representation to achieve viewpoint-independent recognition.

Acknowledgments—The authors contributed equally to this research, and authorship order was determined arbitrarily. Thanks to Michael Spivey-Knowlton, Vaibhav Diwadkar, and Linda Hermer for comments on earlier drafts of this manuscript and to Janellen Huttenlocher for helpful comments, criticisms, and suggestions. Thanks also to Richard Eibach, Seth Bowden, and Matt Zarnowiecki for help conducting Experiments 2 and 3. D.S. was supported by a Jacob Javits fellowship, and some of this research appeared in his doctoral dissertation at Cornell University.

REFERENCES

- Berthoz, A., Israel, I., Francois, P.G., Grasso, R., & Tsuzuku, T. (1995). Spatial memory of body linear displacement: What is being stored? *Science*, *269*, 95–98.
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, *94*, 115–147.
- Biederman, I., & Gerhardstein, P.C. (1993). Recognizing depth-rotated objects: Evidence and conditions for three-dimensional viewpoint invariance. *Journal of Experimental Psychology: Human Perception and Performance*, *19*, 1162–1182.
- Christou, C., & Bühlhoff, H.H. (1997). *View-direction specificity in scene recognition after active and passive learning* (Technical Report No. 53). Tübingen, Germany: Max-Planck-Institut für Biologische Kybernetik.
- Cooper, E.E., Biederman, I., & Hummel, J.E. (1992). Metric invariance in object recognition: A review and further evidence. *Canadian Journal of Psychology*, *46*, 191–214.
- Corballis, M.C., Zbrodoff, N.J., Shetzer, L.I., & Butler, P.B. (1978). Decisions about identity and orientation of rotated letters and digits. *Memory and Cognition*, *6*, 98–107.
- Diwadkar, V.A., & McNamara, T.P. (1997). Viewpoint dependence in scene recognition. *Psychological Science*, *8*, 302–307.
- Easton, R.D., & Sholl, M.J. (1995). Object-array structure, frames of reference, and retrieval of spatial knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*, 483–500.
- Edelman, S., & Bühlhoff, H.H. (1992). Orientation dependence in the recognition of familiar and novel views of three-dimensional objects. *Vision Research*, *32*, 2385–2400.
- Hummel, J.E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review*, *99*, 480–517.
- Huttenlocher, J., & Presson, C.C. (1973). Mental rotation and the perspective problem. *Cognitive Psychology*, *4*, 277–299.
- Huttenlocher, J., & Presson, C.C. (1979). The coding and transformation of spatial information. *Cognitive Psychology*, *11*, 375–394.
- Loomis, J.M., Klatzky, R.L., Golledge, R.G., Cicinelli, J.G., Pellegrino, J.W., & Fry, P.A. (1993). Nonvisual navigation by blind and sighted: Assessment of path integration ability. *Journal of Experimental Psychology: General*, *122*, 73–91.
- Presson, C.C. (1982). Strategies in spatial reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *8*, 243–251.
- Rieser, J.J. (1989). Access to knowledge of spatial structure at novel points of observation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*, 1157–1165.
- Rieser, J.J., Garing, A.E., & Young, M.F. (1994). Imagery, action, and young children's spatial orientation: It's not being there that counts, it's what one has in mind. *Child Development*, *65*, 1262–1278.
- Sholl, M.J. (1989). The relation between horizontality and rod-and-frame and vestibular navigational performance. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*, 110–125.
- Simons, D.J. (1996). Accurate visual detection of layout changes requires a stable observer position. *Investigative Ophthalmology & Visual Science*, *37*, s519.
- Tarr, M.J., & Pinker, S. (1989). Mental rotation and orientation-dependence in shape recognition. *Cognitive Psychology*, *21*, 233–282.
- Ullman, S. (1989). Aligning pictorial descriptions: An approach to object recognition. *Cognition*, *32*, 193–254.

(RECEIVED 8/22/97; ACCEPTED 1/2/98)