

13 *Memory for Pictures, Maps, Environments, and Graphs*

Barbara Tversky
Stanford University

Historically, there have been two kinds of theories used to explain memory for the visuospatial world: theories imported from those about language or theories relying on the picture metaphor. This chapter develops an alternative theory that derives from perceptual processing and accounts for systematic errors in memory for maps, environments, and graphs.

An early process in perception is distinguishing figures from ground. Contour is a critical cue to "figuralness," as well as to memory for objects. Parts and internal features are also operative. Once isolated, figures are related to other figures at the same level of analysis and to a frame of reference at a higher level of analysis. Both perceptual and conceptual factors affect choice of reference figures and frames. This analysis accounts for distortions of shape as well as of distance and location, such as those due to hierarchical organization, landmarks, perspective, alignment, and rotation. Analogous biases appear in nonspatial cognition and analogous coding appears in spatial language.

INTRODUCTION

Pictorial Representations as Like Language

In the 1960s, as the nonsense syllable lost its hold as the stimulus par excellence in research on memory, some ventured so far as to study memory for pictures. It happened gradually, first words, then different kinds of words,

leading to a surprising finding. Memory for concrete words was better than memory for abstract words, even controlling for just about anything anyone could think of (Paivio, 1971). Then came a bolder step: actual pictures. The extant explanations for memory for pictures, however, were based on words. One early example is the "verbal loop hypothesis" (Glanzer & Clark, 1963), according to which people remember pictures by describing them to themselves and remembering the descriptions. If so, then pictures receiving longer descriptions should be more difficult to remember, a finding Glanzer and Clark obtained for sequences of geometric figures. Similarly, Brown and Lenneberg (1954) found that more codable colors were remembered better, where "codable" was also related to length of description. With the influence of artificial intelligence modeling, this view became both more general and more sophisticated. Rather than being represented as simple verbal descriptions, pictures were thought of as being represented as propositions. Not just pictures, but discourse as well, and for some, all thought could be represented as propositions, providing the mind with a single underlying language (e.g., Anderson, 1978; Clark & Chase, 1972; Pylyshyn, 1973).

Pictorial Representations as Like Pictures

At the same time, and spurred by some of the same findings, a different view of memory for pictures was evolving. Not only were concrete words remembered better than abstract ones, but also pictures of objects were remembered better than their names or descriptions, and words that were imaged were remembered better than words that were not imaged (Bower, 1972; Paivio, 1969, 1971, 1986). It is odd to explain better memory by poorer memory; if pictures are remembered better than descriptions, then memory for pictures must be something more than memory for descriptions. Furthermore, the world is full of visual stimuli that are difficult to describe, or to distinguish in words, yet are remembered. The prime example is faces. Of course, there are ways to defend the claim that mental representations are propositional even from these attacks, but those arguments have become more contrived and less appealing.

Other research showed qualitative as well as quantitative differences between memory for pictures and memory for words. Memory for pictures of faces reflects the visual similarity of the faces, whereas memory for names of faces reflects the acoustic similarity of the names (Tversky, 1969). Parallel findings were demonstrated for images, which were assumed to have pictorial qualities (e.g., Kosslyn, 1976; Shepard & Chipman, 1970; Shepard & Podgorny, 1978). Against persistent opposition and in an effort to legitimize the study of imagery, researchers accumulated demonstration after demonstration of pictorial characteristics of images. Images were

shown to reflect physical properties, such as size, shape, and color. Processing images, such as scanning or rotating them, was shown to be analogous to physical processes (for reviews, see Finke & Shepard, 1986; Kosslyn, 1980). In short, researchers repeatedly demonstrated that mental images were similar to pictures.

Undoubtedly, mental images bear resemblances to pictures. But this evidence has been taken by many to imply that pictorial representations are like pictures, and that memory for the visual world is like faded pictures. This view has simplicity and elegance, but it is contradicted by considerable evidence, some of which is reviewed later in this chapter. Put simply, a theory based on the likeness of memories of the pictorial world to pictures of it predicts that memory should fade the way that pictures fade, unsystematically. If picture memories just get dimmer, then errors in memory should be random. In fact, as is shown later, memory errors are systematic and biased.

In this chapter, I develop another approach to understanding memory for the visuospatial world. That approach derives from an analysis of how scenes are perceived and comprehended, with the underlying assumption that the processes involved in perception and comprehension affect memory. The approach also derives from an analysis of how information is retrieved from memory, reconstructive processes that bear similarities to the constructive ones used in perception and comprehension. Supporting evidence comes from studies of practical stimuli, particularly maps and environments, but also graphs. These are not only interesting in their own right, but especially informative to the general case of picture memory because of their schematic character. They are real and ecologically valid stimuli and also allow systematic variation of important physical attributes. Much of the relevant data comes from studies of systematic errors and biases, with the special insight errors and biases provide into perceptual and comprehension processes.

A PERCEPTUAL ANALYSIS

One of the earliest processes in perceiving a scene is distinguishing figures from backgrounds (e.g., Hochberg, 1978). In contrast to backgrounds, figures have closed contours, and it is the tendency toward closed contours that provides viewers with clues to figuralness. Figures are more likely to be symmetric than grounds, so symmetry is also a clue to figuralness. After figures have been determined, they can be located. Rather than absolutely, figures are located relatively, in two ways: relative to other figures at the same level of analysis, and relative to a frame of reference at a higher level of analysis.

This analysis suggests a way to subdivide memory phenomena into memory for the figures themselves and memory for the relations among figures.

Put differently, it is convenient to distinguish memory for objects or elements of a scene and memory for the spatial relations among objects in a scene. In fact, this also provides a sensible separation of the memory phenomena. Memory for elements is briefly reviewed first, followed by a brief review of memory for relations among elements.

MEMORY FOR ELEMENTS

There are many ways to look at memory for elements. One way is to examine the factors that facilitate memory, for example, exposure time, familiarity, elaboration, reinstatement of context, dissimilarity, and organization, factors that affect memory for anything. Although there is evidence for all of these, it is not reviewed here. Instead, evidence that sheds light on the quality of visuospatial memory representations, evidence culled from different paradigms, is reviewed, in continuing the perceptual analysis.

Once a figure is isolated, there is usually an attempt to identify it. Figures can be difficult to identify when portions of them are deleted at random. When figures can be identified despite such degradation, they are remembered far better than when they cannot be identified and thus appear to be a collection of unconnected blobs (Wiseman & Neisser, 1974). Just as contour is important in isolating figures, it is also important in identifying them, because different classes of figures tend to have different contours or shapes (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976). In maps, figures are often land masses, towns, and lakes; in graphs, they may be curves or lines. For the ordinary objects people experience, identification is usually at the basic level, that is, the level of table or dog or apple rather than the level of furniture or animal or fruit (e.g., Murphy & Smith, 1982; Rosch, 1978). Basic level identification can often be made solely on the basis of contour or shape, as different basic level objects differ in shape.

Despite the importance of shape, representations of objects contain more information than shape, and that additional information may be influential in identification. Shape can be decomposed into parts that are distinguished as local discontinuities (Hoffman & Richards, 1984) that constitute a relatively small vocabulary (Biederman, 1987) and that relate perception to function (Tversky & Hemenway, 1984). Because parts have an integrity of their own that transcends a particular viewpoint, they may operate in recognition independent of shape. Identifying people and some objects often requires a step beyond identification at the basic level, to identification at the level of an individual. Because individual people or category members share the same contours for the most part, identifying at a more specific level generally requires more than contour information. In the important case of faces, for example, individual identification uses information about the configuration of features, many of them internal to the contour of the

face (Carey & Diamond, 1977). Identification of objects may also continue to a more refined, subordinate level, such as the level of coffee table or terrier or red apple. Again, contour may not entirely differentiate objects at the subordinate level; other features, such as texture or color or variants on parts and shape, may be needed (Tversky & Hemenway, 1984).

Contour: Symmetry

Many of the common natural objects and artifacts in the world are bilaterally symmetric, frequently along the vertical axis. Perhaps because of the prevalence of symmetry in the world of objects, there is a bias toward symmetry in both perception and memory. Figures that are visibly asymmetric, but only slightly so, are perceived as more similar to a more symmetric figure than to a less symmetric figure, where both alternatives are equally different physically (Freyd & Tversky, 1984). When people are asked to assign interpretations to random filled polygons, they interpret symmetric ones as symmetric objects and asymmetric ones as bilaterally symmetric objects that have been turned (McBeath, Schiano, & Tversky, 1995). Figure 13.1 has examples of each. For example, the upper left figure was called "cartoon dog" (facing right) or "laughing mouse" (facing left), and the upper right figure was called "man in a sombrero." Finally, slightly asymmetric curves in graphs and rivers in map sketches were remembered as more symmetric than they actually were (Tversky & Schiano, 1989).




FIG. 13.1. Stimuli used in studies of symmetry bias. (a) From Freyd and Tversky, 1984. (b) From McBeath, Schiano, and Tversky, 1995. Used with permission.

Canonical View

Despite the bias toward symmetry in perception and memory of contours, the preferred view of objects and faces is not always the symmetric view. In fact, the bias toward perceiving asymmetric objects as turned symmetric ones may encourage canonical views that are not symmetric. Objects, after all, are generally three-dimensional, and a symmetric view lacks one of them. For faces, a $\frac{3}{4}$ view is often better recognized than either a symmetric, frontal view or a profile (e.g., Hagen & Perkins, 1983; Shapiro & Penrod, 1986). For a face, a $\frac{3}{4}$ view provides depth information about facial features that is absent in a frontal view and configurational information about facial features absent in a profile. For common objects, the view that is preferred and recognized fastest is the view that conveys relatively more of the information important to identifying the object (Palmer, Rosch, & Chase, 1981). For many cases, such as a car, a chair, and a piano, this was a $\frac{3}{4}$ view, probably for reasons similar to faces. For other objects, such as a house, an alarm clock, and a telephone, a more frontal view was preferred. For these objects, configural information seems more important and depth information less. For still other objects, such as a horse and a teapot, a side view was preferred. In these cases, it seems that a profile conveys the essential configuration of features for objects with known symmetry. Of course, it is also possible that canonical views are also more familiar views, though Palmer, Rosch, and Chase found the same relationships between judged informativeness of view, rated goodness of view, and identification time of view for novel objects. Together, these findings add to the argument that although contour is important in object identification and representation, other features may also be operative.

Part of the appeal of a canonical view, then, is that it contains some information about contour, some information about configuration and internal features, and often some information about depth. Might memory representations of objects correspond to canonical views of them? The large number of experiments indicating that the time it takes to recognize rotated objects increases with the angle of rotation from upright suggests exactly that (e.g., Jolicoeur, 1985). However, a theory proposing multiple representations of objects with priority to familiar views accounts for those findings, and others as well (e.g., Tarr & Pinker, 1990; Ullman, 1989).

Familiarity

In general, people are faster at recognizing familiar objects than unfamiliar ones (e.g., Rosch, 1975). Similarly, people are better at recognizing faces from their own race than faces from other races (e.g., Shapiro & Penrod, 1986). Presumably, greater experience distinguishing same-race faces leads to learn-

ing the features that vary in that race and selectively attending to those features when observing faces. The features that account for variability in one race may be different from those that account for variability in another.

Caricature Effect

Greater attention to more distinguishing features may account for the super-fidelity of caricatures, if, in fact, caricatures turn out to be superfaithful. At issue is whether memory representations of individuals are closer to caricatures of the individuals than to more realistic drawings or photographs. Because memory representations cannot be directly observed, they must be inferred from measures such as recognition speed or recognition accuracy or judgments of likeness. But factors other than similarity to memory representation can account for performance, especially speed of recognition. An early study cited as support for the caricature effect used photos, line drawings, and cartoons not of faces, but of a hand, switches, and valves (Ryan & Schwartz, 1956). The dependent measure, however, was speed of recognition. The cartoons were the simplest of all the stimuli, and the others were so cluttered with detail that it was difficult to discriminate what they were. Later work comparing photographs and caricatures of faces reported disadvantages of caricatures in recall and recognition accuracy, in recognition speed, and in judgments of likeness (Hagen & Perkins, 1983; Tversky & Baratz, 1985). Yet, comparing photographs to drawings may not be a fair comparison.

Brennan (1985) developed a computer program that produced caricatures by exaggerating individual features away from a group average. This technique was used by Rhodes, Brennan, and Carey (1987) to study recognition of drawings of students and faculty and was applied to photographs by Benson and Perrett (1991) to study recognition of faces of public figures. Caricatures had a small advantage in speed of recognition. Notably, caricatures did not have an advantage over base veridical drawings or photos in either accuracy or likeness judgments. Thus, to the extent that there is evidence for superfidelity of caricatures, it is in speed of identification. It certainly makes sense that making distinctive features slightly more prominent increases the speed of identification, but it does not follow from that that memory representations of faces are exaggerated relative to a norm. The fact that likeness judgments were on the whole highest for the base veridical drawings and photos is evidence to the contrary.

Exposure Is Not Sufficient to Render Features Distinctive

Passive observation without the need to distinguish among individuals does not seem to be sufficient to learn features. For example, Americans can neither correctly recall nor correctly recognize the features in a penny,

despite years of daily experience with them. Because they can distinguish pennies from other coins on the basis of color (and size), they do not need to encode the exact words and pictures on each face of a penny, and for the most part, judging from poor performance, they do not (Nickerson & Adams, 1979).

Schemas, Typicality

In a classic study, Bartlett (1932) showed a mask-like drawing of a face (see Fig. 13.2) to a subject for reproduction from memory. That subject's drawing was shown to another, who drew it from memory, and so on. Gradually, the drawings evolved toward a more typical (Western) sketch of a face. Do memory representations evolve toward the more typical or familiar? A generation of psychologists asked a related question: Do memory representations evolve toward good gestalts? Years of research, pro and con, led Riley (1962) and Postman (1954) to conclude that memories of drawings tend toward the familiar, often not necessarily toward good gestalts. Although true for reproduction or recall, the opposite seems to hold for recognition. In recognition, atypical faces are recognized better than typical (e.g., Light, Kayra-Stuart, & Hollander, 1979; Shapiro & Penrod, 1986). This seems to be a general memory phenomenon: frequent or familiar words are better recalled but rare ones are better recognized (Kintsch, 1970), and schema-consistent elements of a scene or scenario are better recalled and schema-consistent intrusions are common, but deviant elements are better recognized and rarely intruded (e.g., Bower, Black, & Turner, 1979; Brewer & Nakamura, 1984; Brewer & Tenpenny, 1994). Such opposing effects in recognition and recall must be a consequence of retrieval processes rather than of representation.

Labels

In the same year that Bartlett published his work on schemas, Carmichael, Hogan, and Walter (1932) published a study demonstrating the effects of verbal labels on reproductive memory for objects. Thus, a figure like O-O, when labeled as "dumbbells" was drawn closer to barbells but when labeled as "eyeglasses" was drawn closer to glasses. This finding lent support to the early theories accounting for memory of the visuospatial world as descriptions. However, distortion toward a label turned out to occur only in reproductive memory, not in recognition memory (Prentice, 1954). The effect of labels, then, is similar to the effect of schemas; it appears to guide reproductive recall, but not recognition, indicating that its effect is in retrieval rather than in representation (cf. Riley, 1962).




FIG. 13.2. Successive reproductions of a mask-like face. Adapted from Bartlett, 1932. Reproduced with permission of Cambridge University Press.

Is There a Memory Representation?

Contradictory effects of typicality-schemas-familiarity-labels in recognition and in recall indicate that in the end, it may be impossible to disentangle the "memory representation" from the retrieval task (cf. Anderson, 1978). In fact, many current views of memory, especially distributed views, dispute the idea that there is a single, unitary memory representation independent of retrieval (e.g., Rumelhart & McClelland, 1986). Schematization in recall and schema-consistent intrusions may be a consequence of schema-guided retrieval as well as schema-guided encoding. Schemas may be useful both in construction at encoding and in reconstruction at retrieval, but they do not seem to be operative in recognition. In recognition, on the other hand, atypical features may have an advantage because of their very distinctiveness. Unusual features that are easy to recognize are at the same time more difficult to integrate into a schema. The information remembered about an event, visuospatial or otherwise, is not independent of how that information is retrieved.

Conclusions So Far

Contour figures prominently in the isolation and identification of figures. Other aspects of figures, most importantly configuration of parts and internal features, can also be influential. Symmetry, another clue to detecting figures, appears to slightly distort memory for objects. Thus, expectations about objects, that they are closed figures, that they are bilaterally symmetric along a vertical axis, affect object perception and memory. Memory by recall or reproduction is often altered toward the typical or the schematic or the label. Usually, these effects do not appear in recognition. Disparate findings for different memory measures demonstrate that memory is inextricable from the processes used to retrieve it.

EFFECTS OF OTHER ELEMENTS ON MEMORY FOR SPATIAL RELATIONS

Locations cannot be remembered absolutely. There is no device in the eye or brain that gives coordinates of objects. Instead, memory for locations is relative in two ways: to other elements on the same level of analysis, and to a frame of reference on a higher level of analysis. Selection of reference objects and reference frames is determined by both perceptual and conceptual factors. These processes, of relating elements to other objects and to a frame of reference, can distort memory for location and orientation

of elements. Most of the examples demonstrating these phenomena come from memory for maps and environments. These are relatively simple stimuli where spatial relations can be controlled or isolated (see Tversky, 1991, for a review).

Alignment

When students were asked to select which of the two maps of the world displayed in Fig. 13.3 is correct, a significant majority selected the lower map, which is incorrect (Tversky, 1981). In the incorrect map preferred by respondents, the Western hemisphere has been moved northward relative to Europe, Africa, and Asia so that the United States is more parallel to Europe and South America is more parallel to Africa. This shift underlies the reason for the error. Presumably, the locations of the major continents are remembered relative to each other and they are consequently remembered




FIG. 13.3. Maps of the world used by Tversky, 1981. Reproduced with permission from Tversky, 1991.




FIG. 13.4. Maps of the Americas used by Tversky, 1981. Reproduced with permission from Tversky, 1991.

bered as more aligned than they actually are. People make similar errors in judging relative directions of cities in the various continents. For example, a majority of subjects think that Rome is south of Philadelphia, when in fact Rome is north of Philadelphia. Although European cities at latitudes comparable to American cities seem warmer, weather cannot account for alignment, as the same effect was obtained for the east-west shift in the maps in Fig. 13.4. A significant majority of subjects selected the left map as the more veridical looking one. In that map, South America has been moved westward relative to North America so that the two Americas are more aligned than they truly are. Alignment errors occur in estimates of directions between pairs of cities in countries and pairs of roads in towns (Byrne, 1979, 1982; Tversky, 1981). They also appear in artificial maps and in visual blobs (Tversky, 1981).

Landmarks

In the previous cases, the elements were approximately equal in size and prominence. In other cases, however, some elements are more prominent than others. In environments, prominent elements are often called landmarks. Landmarks may differ from ordinary elements in a number of ways. For example, they may be larger or more centrally located or better known (e.g., Couclelis, Golledge, Gale, & Tobler, 1987). Because of any of these features, landmarks seem to organize the elements in their vicinity. When asked where

they live, people often reply with near the nearest landmark likely to be known to the questioner (Shanon, 1983). Using a landmark as a reference object, however, distorts distance estimates. Sadalla, Burroughs, and Staplin (1980) found out what campus buildings were regarded as good landmarks. They then asked students to make distance estimates of campus buildings, using either ordinary buildings or landmarks as anchors. Students judged ordinary buildings to be closer to landmarks than landmarks to ordinary buildings. These errors are inconsistent with any metric model of space.

Perspective

According to popular cartoons, in the New Yorker's view of the world, Manhattan, the Hamptons, the Jersey Coast, and Connecticut loom large, far from one another. Squeezed together into the receding distance are the Mississippi, the Rockies, and California to the West and England and the European continent to the East. Similar cartoons have mocked the egocentric perspectives of other denizens of the world; in each, the immediate environment is magnified relative to the distant environment. In fact, in perception, distances between near points *appear* larger than distances between far points. Might this distortion of distance due to perspective also occur in memory? In order to check, Holyoak and Mah (1982) asked subjects to imagine themselves on either the Atlantic or the Pacific Coast and to give judgments of distances between pairs of cities more or less equidistant along an east-west swath across the United States. In fact, subjects with an Atlantic perspective judged the distance from New York to Pittsburgh greater than those with a Pacific perspective, and vice versa. Those with a Pacific perspective judged the distance between San Francisco and Salt Lake City to be greater than those with an Atlantic perspective. Additionally, the distorting effects of perspective was flexible. All of the subjects were actually students in Ann Arbor, Michigan, who had been asked to take one perspective or another.

EFFECTS OF A REFERENCE FRAME ON MEMORY FOR SPATIAL RELATIONS

Grouping

Space is not inherently hierarchical, but people impose a hierarchical structure on it. This may be a simple matter of grouping some elements together in one region, and others in another region, for example, by proximity or by natural geographic boundaries or by function. The very fact of grouping some elements together and consequently separating them from other elements can create distortions of distance and orientation.

In a well-known example, Stevens and Coupe (1978) asked students to indicate the directions between pairs of cities such as Reno and San Diego. Although Reno is in fact west of San Diego, most subjects indicated that it was east of San Diego. Stevens and Coupe proposed that, rather than remember the relative locations of all pairs of cities, people remember the relative locations of states and also which cities are in which states. Then, they use the relative locations of the states to make judgments about the relative locations of the cities within them. Because in general California is west of Nevada, people judge all cities in California to be west of all cities in Nevada, though this is not always the case. Effects of hierarchical organization have also been demonstrated for constructed worlds (e.g., McNamara, 1986, 1992; Stevens & Coupe, 1978).

Distances as well as directions may be distorted by grouping. Hirtle and Jonides (1985) asked one group of Michigan students to group related buildings in Ann Arbor. On the whole, students put campus buildings together and separated them from town buildings, which they also put together. Another group of students was asked to estimate distances between pairs of buildings. Distance estimates between groupings were overestimated relative to distance estimates within groupings, even though the students who estimated distances were not the same as those who produced the groupings. Coren and Girgus (1980) found similar effects on memory for dots. Pairs of dots grouped into the same figure were judged to be closer than pairs of dots grouped into different figures although the absolute distance between dot pairs was the same.

Reaction times to make distance judgments also reflect grouping. People are faster to judge that one city is east or north of another when the cities are in different geographic entities than when they are in the same geographic entity, even when actual distances are smaller (Maki, Maki, & Marsh, 1977; Wilton, 1979).

Reference Frames

Elements may be related to a reference frame as well as to other elements and to groups of elements. The walls of a room, the sides of a piece of paper, and the north-south east-west longitude and latitude are examples of reference frames for locating elements. In the case of dots, they are remembered as closer (but not too close) to their relevant frames of reference than they actually are (Huttenlocher, Hedges, & Duncan, 1991; Nelson & Chaiklin, 1980; Taylor, 1961).

Reference frames also affect the remembered orientation of figures. An important part of identifying a figure is determining its orientation (e.g., Braine, 1978; Rock, 1973). Figures without a clear orientation are unstable

in perception (Atneave, 1971). Even meaningless figures are perceived to have a natural orientation, with a preference for vertical elongation, balance, and important features at the top (Braine, 1978). When the natural orientation of a figure conflicts with its actual orientation within a frame of reference, its orientation is remembered as closer to the frame of reference than it actually is. An example is the orientation of South America, which appears tilted. Not so much in memory, however. When asked to place cutouts of South America in a north-south, east-west frame of reference, most people rotated South America so that it appeared closer to upright (Tversky, 1981). Rotation errors also occurred in judgments of the directions between pairs of cities and in artificial stimuli; it also appeared for other geographic regions (e.g., Chase, 1983; Chase & Chi, 1981; Moar & Carleton, 1982).

The effects of cognitive factors in selection of a reference frame are apparent in memory for graphs. A line interpreted as a function in a pair of axes is remembered as closer to the implicit 45-degree line. That same line interpreted as a path in a map or as simply a line in a framework is remembered differently (Schiano & Tversky, 1992; Tversky & Schiano, 1989). For graphs, but not for paths or simple lines, the implicit 45-degree line serves as a meaningful referent in the interpretation of straight-line functions. Apparently for this reason it comes to serve as a frame for remembering the orientation of the line as well.

REFERENCE OBJECTS, REFERENCE FRAMES

Figures, then, are located and oriented with respect to each other and with respect to a reference frame. These phenomena bear similarities to the gestalt principles of grouping and common fate respectively. Reference elements and reference frames then serve as anchors for the target elements so that the target element is remembered to be closer to or more aligned with the reference object or frame than it really is.

Both perceptual and conceptual factors determine selection of reference objects and reference frames. For reference objects, the perceptual features, proximity, salience, and similarity all seem to contribute to selection of a reference object, for example, the Americas for each other, Europe and the United States for each other, South America and Africa for each other. Salience, centrality, accessibility, and familiarity are some of the factors making for a good reference object. Horizontal and vertical lines make excellent reference frames, for a variety of reasons. Both are defined in the world, the horizontal by the horizon or ground, the vertical by gravity and its effects in the world. Perhaps because of their importance and salience in the world, the vertical and horizontal have privileged status in perception (e.g., Howard, 1982). Other, related, perceptually compelling reference