How visual landmarks are selected during small-scale navigation

Vlada Aginsky (Knowlton)

April 25, 2001

Department of Cognitive and Linguistic Sciences

Brown University

1 General Introduction

When thinking about a frequently traveled route people generally recall a few discrete objects or places connected by stretches of indistinct paths. As a child I used one particular navigation strategy to get to my best friend's house. After reaching Coolidge Corner I traveled up Harvard St. and looked for a rounded corner wall of yellow bricks. That rounded wall at the corner of Coolidge Street marked the turn that would lead me to my friend's house. Over the years I must have turned that corner several hundred times. However, if someone asked me to recall what was across the street from the yellow brick wall I wouldn't have the slightest clue. For some reason, at that particular juncture, I used that wall as a "landmark" and invariably ignored numerous other objects that could have played the same role. Why did I choose that wall?

A traveler can use several different forms of information to navigate. Most of these forms of information are exploited by one of two main navigation strategies. One strategy entails using information about self-motion to update one's distance and direction from a particular starting or ending position. This strategy is often referred to as "path integration" (e.g., Mittelstaedt & Mittelstaedt, 1982) or "dead reckoning". When using this strategy the traveler relies on internally-generated information such as proprioceptive or vestibular signals to monitor his body's translations and rotations as he moves along a path. He can also use optic flow, or the retinal motion pattern created as the traveler moves through the environment, to make similar orientation judgements (for review see Warren, 1998). Path integration does not, however, entail the use of discrete objects, or "visual landmarks".

Visual landmarks are utilized as part of a strategy that relies on associations between the positions of various objects in the environment. Therefore, a location can be recognized according to its spatial relationship with another recognized object, without the need to reference information about self-motion. This strategy is often referred to as "piloting" (for review see Papi, 1992).

One of the challenges in human navigation research has been to explain why particular objects are used as visual landmarks. That is, while it is simple to determine which objects are selected to be used as landmarks in a given environment it is difficult to determine the underlying principles that guide the selection. The problem lies partially in the fact that so many different objects are used as landmarks in various situations. While a traveler may rely upon a view of a mountain to orient himself in one place he may use a brightly colored sign to do the same in another place. The traveler's reasons for using such different objects as visual landmarks are difficult to ascertain because in a real-world environment such objects are presented arbitrarily. Furthermore, people generally rely on landmarks that are too large to manipulate for experimental purposes. For instance, without replacing one object with another we cannot tell whether or not a different object in the same place would serve just as well. Without the ability to physically displace a potential landmark it is also difficult to tell whether or not it would still be used as a landmark if it stood in a slightly different position.

In studies using small animals researchers have found much greater success in revealing the underlying principles that guide their subjects' choices of visual landmarks. This is because the objects encompassed within the environments of such small travelers have been rather easy to manipulate. By displacing or removing different objects from the foraging environments of small animals researchers have been able to observe which objects influence the animals' navigating abilities (e.g. Cheng, 1986; Collett, Cartwright, & Smith, 1986; Spetch, Cheng, & MacDonald, 1996). In this manner they have been able to determine which aspects of potential visual landmarks animals weight most heavily when selecting the landmarks they will rely on.

In the current study the question of how humans select visual landmarks will be approached in the same way that it has been approached by researchers studying animal navigation. In the environment where human subjects will perform a place finding task potential visual landmarks will be manipulated to see how much influence they exert on the subjects' performance. The problem of object manipulation will be solved by using a virtual environment in which any large- or small-sized objects can be easily displaced or removed.

Two main questions will be addressed in this study. The first question asks whether or not people are able to switch from a piloting strategy to path integration when visual landmarks are known to be unreliable. Subjects will be encouraged to abandon the piloting strategy by being informed that objects in the virtual environment will shift positions throughout the course of the experiment. Thus, they will be aware of the fact that these objects do not make useful landmarks. This question will address the degree to which people are compelled to use visual landmarks even when it would be more sensible to use other means of navigation under the given circumstances.

The second question asks which objects will people choose to rely on as visual landmarks if the piloting information provided by the objects *cannot* be ignored. For example, will people prefer to rely on relatively large or small objects? Will they use a single object or a configuration of several objects to find the target location? The answer to the second question will help us to understand how visual landmarks are selected.

1.1 VISUAL LANDMARKS

What are visual landmarks and how do humans and animals use them when traveling from point A to point B? Perhaps it would suffice to invoke an intuitively valid definition and be done with the matter. For instance, it would not be unreasonable to say that a landmark is a visually salient and distinctive object which, by being associated in memory with a specific location, another landmark, or an action, helps a traveler identify his/her position in the environment. Unfortunately it's not that simple.

Even if we accept the above definition we are still left with a very impoverished understanding of how visual landmarks are selected and used. There are clearly many different spatial environments in which visual landmarks can play different roles. In a grid-like arrangement of city streets a sign on a corner may simply tell a traveler that he is "two blocks" away from his destination. In an irregular desert terrain a tall cactus may function as a point within a configuration of rocks and cacti that is a half day's walk from a camp that lies in the direction of the early morning sun. Or, perhaps, the edge of a building, when viewed so that it lines up with an edge of a taller building in the distance indicates the direction which one must follow at that juncture along a known route.

Animal studies have often converged on the notion that there are two general landmark categories, each serving distinct navigation functions. Global landmarks are objects that act like beacons by helping the animal choose a direction of travel. In the real world global landmarks are generally large objects that can be seen from great distances, such as mountains, tall trees, or tall, man-made structures. In order for a global landmark to be useful in defining a specific location it must be approached from a specific direction so that the traveler can match the location with the contours of the global landmark "view" (e.g., Cartwright & Collett, 1982).

Local landmarks can be employed once the animal enters the vicinity of a destination and needs to home in on an exact spot (e.g., Collett, 1996). Local landmarks are smaller, positioned somewhere near the desired location, and thus can only be seen once the traveler is in fairly close proximity to the goal. Unlike global landmarks, local landmarks are small enough to be approached from any direction and still be used successfully to find the goal.

In this study participants will be presented with both local and global landmarks that will be optimally positioned for a successful location search. That is, as the subjects approach the goal the global landmarks will be seen directly behind it. Therefore these objects will always be approached from the same direction and their contours will offer a useful context within which the goal can be fixed. The local landmarks will also be readily available when the subjects make their directional decisions. These smaller objects will be placed at various distances from the goal, but none will be too far so that the subjects cannot incorporate them into their search. With both types of landmarks available it will be interesting to see if participants will prefer one type over another or if they will use a combination of the two.

1.2 HOW ARE LANDMARKS USED?

1.2.1 Some functions of local landmarks

The process of finding a point in space using visual landmarks, or "piloting", has been shown to occur in several different ways. For instance, one can choose to rely on a single visual landmark or one can use a configuration of landmarks to indicate the position of the goal. One study investigating how configurations of identical landmarks are used by pigeons and humans showed that different strategies are used by the two groups (Spetch, Cheng, & MacDonald, 1996; Spetch, Cheng, MacDonald, Linkenhoker, Kelly, & Doerkson, 1997). When configurations of landmarks were expanded or contracted humans looked for a spot that preserved the same spatial relation to all the landmarks (for e.g., was in the "center"), whereas pigeons looked for a spot that was in the same place as before relative to one landmark. In another experiment gerbils behaved similarly to pigeons, preferring to search a particular distance away from one landmark, rather than calculate a spot that would maintain a previously observed relationship to several landmarks (Collett, Cartwright, & Smith, 1986).

Such studies, therefore, suggested that humans define a location in relation to the whole configuration of landmarks while pigeons and gerbils preferred to rely on an individual landmark. However, we should not be too quick in dismissing the remaining objects when only one appears to be used to specify a location. Evidently, other surrounding landmarks can play an important role even when a single landmark strategy appears to be employed. For example, in the Spetch et al. experiments it was shown that pigeons did take note of all the landmarks in the array, as the experiment provided several target locations which the pigeons encoded with respect to different landmarks. Therefore, the authors reasoned, the birds must have been aware of the landmarks' spatial relations to each other. This observation implied that whereas pigeons did not use the landmark configuration to locate the target spots, they did use the configuration to identify individual landmarks.

A somewhat similar strategy to using multiple-object configurations is using configural information about the overall environment, such as the shape of a room. Some studies have shown instances when animals resort to using this kind of configural information to navigate (e.g., Ramos, 2000; Cheng, 1986; Margules & Gallistel, 1988). In an experiment where rats searched for food inside a rectangular arena these animals made some remarkable systematic errors. Although the walls and corners contained many distinctive visual cues the rats often searched in a location near a corner that was 180 degrees around the center of the arena from the correct corner. The searched location was metrically correct in that it was the correct distance and direction from both the long and short wall. However, inside a rectangle there are two such possible places and sometimes the rats picked the wrong one. This meant that the animals were using the shape of the arena and disregarding the other visual cues (Cheng, 1986).

Some experiments have revealed the technique of using 2-D image patterns or "views" to define locations. Experiments with bees, for instance, have shown that the insects will search for a goal in a place from which landmarks create a retinotopic pattern that matches a pattern previously viewed from the goal (Cartwright & Collett, 1983; Collett, 1996). Therefore when the landmark array is distorted, in order to match the view that was experienced during training, the bees will search at much greater or shorter distances from the landmarks than necessary. The hoverfly has also been observed to use retinal image matching to find its way home (Collett & Land, 1975).

Some recent studies involving human subjects have proposed spatial learning models in which local views are used to learn and plan travel routes (e.g., Mallot & Gillner, 1998; Scholkopf & Mallot, 1995). Such view-based models suggest that navigation behaviors may simply be reduced to learned sequences of views associated with certain movement decisions. Accordingly, different views of the same place would be incorporated into whatever sequence of views and movements allowed the traveler to arrive at each particular vantage point. This type of model is related to view-based object recognition models which state that newly encountered objects are represented as collections of 2-D views experienced by the observer (e.g., Edelman & Bulthoff, 1992; Tarr & Pinker, 1989). Whether we think about a view of an object or of a scene, we are bound to consider the spatial relations of features (in the case of the object) or objects (in the case of a scene) as being an important aspect of the representation.

1.2.2 Global vs. Local Landmarks

Is it true, then, that people are sensitive to the shape or configuration of the local visual cues which guide them to their goal? Does configural scene information carry a greater weight in people's navigational decisions than individual objects? Furthermore, if this is so, then does the distance of the visible objects have any effect on whether or not they are attended to and incorporated into the "landmark" bundle? That is, are global landmarks an equally important part of the visual information that allows a traveler to find the correct spot or are they merely beacons used to find local landmarks from great distances?

Distant or "global" landmarks may, in fact, function as more than just beacons used to determine the direction of a far-off place. Several studies conducted inside the laboratory have shown that rodents often use the configuration of visual cues external to the local landmark array to find local food sources (e.g., Collett, Cartwright, & Smith, 1986; Teroni, Portenier, & Etienne, 1987). This can be witnessed in cases when small transformations are made within the local object arrays (e.g., Collett, Cartwright, & Smith, 1986) and also when the entire local arrays are rotated with respect to the global environment (e.g., Teroni, Portenier, & Etienne, 1987).

Collett & Kelber (1988) found that a surrounding panorama influences the search behavior of honeybees even when visually distinctive local cues are available. They trained the insects to search for sugar water in sites that could be encoded relative to an array of local landmarks of varying colors and shapes. However, when positions of the local landmarks were interchanged the bees picked landmarks to search around according to their positions relative to the surrounding scene and not according to their distinctive features. The implication of this study was that in certain conflict situations the context of the distant panorama could override other potential visual cues such as local landmark features. The authors reasoned that this was an advantageous strategy for the bees to use, since distant landmarks provide a more stable image on their retina than the local landmarks as they fly near the goal.

A recent experiment conducted with human subjects tested the role of local and global landmarks in a virtual environment (Steck & Mallot, 1997). Participants used mouse buttons to navigate along a road system projected on a 180 degree surrounding screen. Combinations of global and local landmarks could be seen at junctions where turn decisions needed to be made. After training the landmarks were manipulated so that several types of cue conflict conditions were presented. The results showed that subjects did not favor one type of landmark over another. They chose to rely on either global or local landmarks at different decision points and the choices varied between individual subjects. Therefore, it was shown that global, as well as local, landmarks were used to make local-scale movement decisions, although the criteria used for choosing one type of landmark over another were not elucidated.

Overall, the evidence is still equivocal about whether or not global landmarks are considered good sources of visual information for a traveler who's trying to zero in on a specific location. This study will attempt to add some clarification to this issue.

1.3 PILOTING VS. PATH INTEGRATION

An alternative method of navigation, one that does not involve the use of visual landmarks, is called dead reckoning. Dead reckoning, originally a nautical term, is the activity of updating one's position in space by using internally generated signals as well as some visual cues such as optic flow or a celestial compass (for review, see Papi, 1992). Many animal navigation studies have investigated this mostly non-visual strategy, also referred to as "path integration" (e.g., Mittelstaedt & Mittelstaedt, 1982). In mammals path integration is believed to be accomplished through the use of proprioceptive signals and motor commands sent to the muscles (efference copies), as well as vestibular signals to calculate bodily translations and rotations (for review see Etienne, Berlie, Georgakipoulos, & Maurer, 1998). Arthropods such as spiders and crabs can also make use of such internally-generated information to integrate their movements. Ants and bees use celestial patterns, while birds make use of both a celestial compass and the earth's geomagnetic field to complete their migrations (see Papi, 1992).

Whatever mechanism is employed, the traveler's task is to be able to update it's direction vector relative to some starting or "home" position at any given point along the path. However, any such estimations invariably contain some error. The errors can accumulate until the traveler is far off course unless some kind of external environment-based correction method is applied. The correction method of choice is, not surprisingly, learning the layout of visual landmarks and using associations of known landmarks with other locations to adjust the homing vector. The interrelation between path integration and visual landmarks had been observed in many navigation studies (e.g., Etienne, Maurer, & Seguinot, 1996).

In their articles Etienne (1992) and Etienne et al. (1996) discussed the complementary relationship between piloting and path integration in small mammals. Without using environmental features these animals mainly rely on vestibular signals to monitor rotational changes and proprioceptive information and efference copies to assess linear displacements. However, according to several triangle completion experiments (returning to the starting point after making an L shaped outward journey) with hamsters, dogs, (and humans) it was observed that, without visual feedback, all groups of subjects made systematic errors when trying to update their homing vectors and return to point one (Etienne et al., 1996). Hamsters who learned to find a nest box at the periphery of an arena could do so efficiently in the dark by using dead reckoning. When the lights were turned on, however, they would switch to a visual strategy, making use of the stable visual landmarks that were available outside the arena (Etienne, 1992). The dominant strategy for each situation was easily ascertained by observing the animals' behavior in conflict situations, when the arena was rotated relative to the visual cues.

A seminal series of studies by Loomis (e.g., Loomis, Klatsky, Golledge, Cicinelli, Pellegrino, and Fry, 1993; Fujita, Klatsky, Loomis, & Golledge, 1993) has shown that humans are also capable of performing path integration when deprived of visual cues. However, similarly to many animal demonstrations human performance in, for instance, triangle completion tasks maintains a fair degree of accuracy up to a certain point, possibly to a distance of about 15 meters (e.g., Fukusima, Loomis, Da Silva, 1997). Loomis and colleagues have reported that in situations of uncertainty people resort to adjusting estimated distances or rotations towards mean values. Therefore they tend to underestimate longer path legs and larger turn angles and overestimate the shorter and smaller ones (Fujita et al., 1993).

As for the influence of motion-related visual information, a study by Riecke, van Veen, & Bulthoff (2000) showed that humans are able to use optic flow in homing tasks, although, again with some degree of error. Participants performed triangle completion by moving (via mouse buttons) through a virtual environment projected onto a half-cylindrical viewing screen. In contrast to their less than accurate performance in the optic flow condition subjects showed perfect homing performance when stable visual landmarks were available. This, according to the authors, "demonstrated that piloting by salient landmarks and visual scene-matching plays a dominant role in visual navigation…" (Riecke et al., 2000, p. 26).

But how does this "dominant role" affect travelers' decisions in situations when visual information is unreliable? There have been many accounts of conflict situations in which the travelers had to make a decision to use either external visual cues or internally generated signals. Ants usually rely on landmark-position associations over path integration when the two sources of information are incongruous (e.g., Collett & Collett, 2000). Hamsters will depend on visual landmarks over idiothetic signals up to a point when the visual information diverges too greatly from previous experiences or proves to be too variable across trials (e.g., Etienne, 1992; Teroni et al., 1987). Even birds, when faced with inconsistent magnetic and celestial signals, choose to use the visual cues, that is, the celestial patterns to navigate (see Papi, 1992). However, in all of these situations the navigators did not know which sources of information were veridical - they had to make their best guess.

What would they do if they were made aware of the unreliability of one and the stability of another? One might assume that they would ignore the unreliable information and attend to the stable information. Part of the current study will attempt to answer this question. Specifically, the experiments described in this paper will place human participants in a situation where the visual landmarks are known to shift positions from trial to trial. Thus they are without a doubt an unreliable source of information. Will the participants be able to ignore such landmarks and instead rely solely on dead reckoning which they know will lead them to the correct target position?

1.4 THE TRANSFORMATIONAL METHOD

In order to study how humans use visual landmarks this series of experiments will use a technique that has been widely adopted in animal navigation studies - the transformational approach (Cheng & Spetch, 1998). The crux of this technique is the displacement of objects that, during training, may have been employed as visual landmarks. Displacing landmarks and observing the subject's subsequent behavior allows us to learn about the visual information that is crucial for the navigation task, without having to manipulate the subject's own visual ability. That is, the "landmark" role of certain objects can be confirmed or repudiated depending on whether or not the subject's orientation is affected by the movement of the objects.

The transformational approach was introduced by Tinbergen (1972) and has since been employed in a multitude of studies (e.g., Cheng, 1986; Collett, Cartwright, & Smith, 1986; Spetch, Cheng, & MacDonald, 1996). Until the emergence of virtual reality (VR) technology, however, this technique was practically infeasible in human studies. This is because it is difficult, if not impossible, to move the types of objects - buildings, trees, mountains etc. - that are generally believed to serve as visual landmarks for humans. In a virtual environment (VE), though, we are freed from physical constraints and can quickly and easily displace or remove both large and small-scale structures. Thus we can see what people do when their world is played with like a doll's house.

1.5 USING A VIRTUAL ENVIRONMENT TO STUDY HUMAN NAVIGATION

Studying visually-guided navigation in a virtual environment reveals many interesting things about people's ability to get around in a VE, but can we assume that such information tells us anything about the way that people navigate in the real world? After all, we are still heavily constrained by limitations in current VR technology and can hardly claim that we can offer a VE experience that comes close to mimicking a real-world environment. Fortunately there is evidence indicating that visual and spatial learning in a VE translates readily into real-world visual and spatial knowledge. Some of this evidence will be summarized in the following sections.

1.5.1 Presence

A VE must offer its navigator a sense of "presence" or "immersion" in order to be an effective simulation of the real environment (RE). That is, the person moving within the VE should feel that his physical actions will change his relationship to the environment in a predictable manner. For example, if the person moves towards a distant object that object should increase in size at a rate that corresponds to the person's approaching speed just as it would in a RE. But how many other features of the VE need to faithfully mimic their RE counterparts in order for the simulation to be considered adequate? To what extent does the visual and auditory, as well as other types of sensory feedback have to resemble that in a RE? And how many different types of sensory data have to be available in order to experience immersion?

This thesis and the experiments which it describes will approach these questions from the ecological point of view. That is, it will embrace the view derived from Gibson's (e.g., 1979) approach to understanding visual and spatial perception : that "the measure of (a VE's) fidelity is the degree to which the simulation captures the richness of natural couplings between perception and action" (Flach & Holden, 1998, p. 93). In their 1998 article, Flach and Holden proposed that the quality of visual, acoustic, or tactile information in a VE does not need to resemble corresponding RW information too closely as long as an agent can interact with the environment in the same way that it interacts with the RW. They suggested that "the experience depends more on what can be "done""(Flach & Holden, 1998, p. 94).

Zahorik and Jenison (1998) argued further in favor of the Gibsonian approach to measuring presence by reinforcing it with a theory of a mid-twentieth century philosopher, Martin Heidegger. Heidegger also believed that our representations of the surroundings are defined by the actions that we perform in them. He also proposed that we represent objects as components of actions we perform when we interact with them. In other words, when one is acting in the world the physical properties of objects are inconsequential because the actor is concentrating on the action. Or, as stated in Zahorik and Jenison's example of hammering, "the concernful action of hammering precludes the user from possessing a stable representation of the equipment utilized in the action, thereby rendering the equipment transparent to the user" (p. 84). Zahorik and Jenison concluded that: "Successfully supported action in the environment is a necessary and sufficient condition for presence" (p. 87).

In the spirit of the above approach I will also adopt the assumption that a virtual environment can offer a valid simulation of the world as long as the participant feels that his or her actions yield the same body-environment transformations as would be experienced in the real world. Luckily enough, a host of studies has surfaced in recent years showing that VEs can be used successfully to study real-world spatial learning.

1.5.2 Spatial knowledge translates from a virtual environment to the real world

Koh, Wiegand, Garnett, and Durlach (1999) compared acquisition of spatial knowledge in the real world to several types of virtual conditions: immersive, nonimmersive, and 3D model environments. The model environment offered a computer model of the space being explored (in this case a section of a building with the roof off) from an outside perspective. In the immersive condition subjects explored the environment using an HMD and in the nonimmersive condition subjects viewed a desktop display. After exploring the experimental space subjects were asked to estimate the bearing and range of various landmarks from different test locations in the space. The investigators did not find any significant accuracy differences among the different training methods. They found much greater variability between the different sub-tasks (various landmark/test location pairs) and between different subjects. Thus in their study VEs offered an adequate alternative to training in a RE.

Waller, Hunt, and Knapp (1998) also found that under certain conditions training in a VE resulted in the ability to navigate while blindfolded in a RE maze as well as subjects who were trained in a RE. During training the VE subjects used an HMD and a joystick to move about. However, the VE subjects were only able to perform as well as RW subjects after relatively long periods of exposure to the environment (5 minutes in the VE as opposed to 1 minute in the RE). Those who did not have extensive exposure to the VE (only 2 minutes) took more time to subsequently walk while blindfolded through the maze. Furthermore Waller et al. found that subjects who trained in a desk-top environment performed only slightly worse than those who trained for the short durations in the immersive VE.

Therefore, Waller's study implied that spatial knowledge can transfer from VEs to the RW as long as enough time is offered to explore the VE. It was suggested by the authors that longer exposure time in VEs is necessary simply because many people are not accustomed to moving about in a VE, and not because of some fundamental limitation of learning in a VE as opposed to the RW. In fact, it was noted that the VE group performed surprisingly well considering that during training they did not have proprioceptive information to complement visual information as did the RW group (Darken, Allard, & Achille, 1998).

The issue of familiarization with VEs was also investigated by Ruddle, Payne and Jones (1998). In one portion of this study subjects navigated two large-scale virtual buildings using a desk-top display and a mouse and keyboard to control movement. Measures of spatial knowledge, such as route-finding ability, sense of distance, and direction estimations were taken throughout a series of sessions in each building. As could be expected participants improved over time as they navigated within a particular building. Interestingly, though, the participants' performance was better on all measures and across all sessions when they were tested in the second building. Thus, simply the experience of moving about in one virtual environment caused an overall improvement in spatial knowledge acquisition in the second environment.

1.5.3 Differences between VR interfaces

Although studies have shown that spatial knowledge transfers (more or less) between real and virtual environments (e.g., Waller et al., 1998; Koh et al., 1999), a substantial number of them have found that some consistent knowledge and performance differences do result from various types of training environments. It is worth noting some of these differences, especially those arising when different interfaces are used, as my study will also reveal divergent performance levels due to differing modes of interaction with the VEs.

With the exception of those who are blind, travelers may consider visual cues to be the most salient and informative elements in the environment. After all, visual information offers unambiguous targets that can be located farther from the observer than the furthest perceptible sources of taste, touch, smell, or even sound. However, visual information alone may not allow an agent to form as rich a mental representation of the world as would visual cues accompanied by proprioceptive feedback (e.g., Kearns, Warren, Tarr, & Duchon, submitted). Studies looking at the ability to recognize configurations of single or multiple objects have shown that looking at such displays while actively moving around them contributes to more accurate representations than being shown different views while remaining stationary (e.g., Simons & Wang, 1998). Evidently static images do not allow the observer to use extra-retinal information (such as vestibular, proprioceptive, efference copies, etc.) that would otherwise contribute to a richer mental representation of the object layout.

Likewise spatial learning in VEs may progress more quickly when the environment can be explored both through vision and natural motion. In order to test the importance of proprioceptive and vestibular signals in visually-guided navigation Chance, Gaunet, Beall, and Loomis (1998) asked participants to learn the locations of objects while moving through a virtual maze. One of three possible locomotion modes (all requiring an HMD) was used for the task: a Walk mode, in which subjects walked freely in the VE, a Real Turn mode, in which subjects could physically turn in place to change directions but were translated only visually by a constant forward motion, and a Visual Turn mode, in which a joystick was used to control turns. After familiarizing themselves with the maze subjects had to indicate the directions to several objects encountered in the maze from a position where the objects could not be seen.

Chance et al. found that the Walk mode yielded the most accurate direction estimations and there was a trend for better accuracy in the Real Turn mode vs. the Visual Turn mode. The authors regarded these results as evidence of a better understanding of one's position relative to other points in space than when non-visual cues are available along with the visual information.

In a study that had participants learn the layout of a virtual building using an HMD (along with a handheld button box) or a desk-top display (along with a mouse and keyboard) Ruddle, Payne, and Jones (1999) observed some interesting differences in the way that participants in the two groups accomplished the task. For example, those using the HMD navigated the virtual building faster than the desk-top group but not because they traveled shorter distances. The desk-top group was slower because they often stopped to look around before proceeding to new sections of the building. By contrast the HMD subjects kept moving while they looked around, creating a smoother flow in their exploratory behavior. The allowance of more natural combinations of head and body movements caused the HMD group to look around more often, but also more gradually. That is, they demonstrated a greater range of head turns, ensuring that they wouldn't travel past important landmarks as the desk-top participants often did.

Although the aforementioned study by Koh et al. (1999) did not find any significant performance differences between subjects trained in different VR systems, the authors acknowledged that the design of their experiment may have served certain natural spatial biases that overshadowed any mode-of-training effects. For instance, in one subtask that required making a bearing estimation in a rectangular room subjects in all groups tended to "square the rectangle".

Overall the studies described above suggested that VEs can be useful and reliable testing grounds when studying human spatial learning. By observing people's behavior in relatively simple virtual environments it is possible to gain a better understanding of how people navigate in the complex real world. The current study used the Virtual Environment Navigation Laboratory at Brown University to study how people select visual landmarks.

1.5.4 The VENlab (Virtual Environment Navigation Laboratory)

The experiments in this study were conducted in the VENlab at Brown University. The VENlab incorporates a Silicon Graphics (SGI) Onyx2 with a Reality Engine graphics board. In combination with Sense8's World Tool Kit software we can build and present high fidelity virtual worlds. These worlds are then displayed on the head mounted display (HMD) at a frame rate of 60Hz, with a resolution of 640 x 480 pixels, with little delay in the graphics display. The VENlab interfaces include an HMD, Intersense Tracker System, and Joystick Controller. The HMD is a Kaiser Electro-optics Proview 80, which is non-see-through, fully immersive, with binocular lenses.

Video input to the HMD controller provides stereoscopic (two independent video channels) viewing, with a 60 degree horizontal field of view, (40 degrees vertical) with binocular viewing. The Tracking System is Intersense's IS-900, a hybrid acousto-inertial 6 degree of freedom position and orientation tracking system. It tracks changes in orientation and position by integrating the outputs of its gyroscopes and accelerometers, and corrects drifts by using a room-referenced ultrasonic time-of-flight range measuring system. The system detects the precise position and orientation of a person in the room and updates the graphics relative to this position. To support this tracking system, the laboratory has 181 position emitters fixed on the ceiling and three microphones plus an inertial system fixed on top of the HMD. The microphones pick up the information from the emitting beacons to track the position. This system allows a person to move freely within the 11 x 11 meter laboratory space.

2 Experiment I

2.1 INTRODUCTION

The purpose of the first experiment was to test whether or not people can abandon a piloting strategy in favor of path integration if the available visual landmarks are known to be unreliable. The second goal of this experiment was to determine which type of landmarks, if any, people prefer to use - local or global - if they cannot completely ignore such visual information. The task in this experiment involved learning to walk to a particular location at the perimeter of an arena after entering the center of the arena through a tunnel. The several objects that stood inside and outside the arena shifted positions in various combinations from trial to trial. Since subjects were informed that such random shifts would occur, they knew that it was impractical to rely upon these objects as visual landmarks. The question was - would they be compelled to rely upon some or all of them anyway? And if so, which ones?

This experiment was divided into two parts: Experiment I - Joystick and Experiment I - Walking. The walking group had more non-visual information available to them so they were more likely to succeed in completely ignoring the objects positioned throughout the environment. That is, with the availability of proprioceptive and vestibular signals they presumably had enough information to perform adequately in the task. The group that moved through the environment with the use of a joystick should have had a tougher challenge in trying to ignore the objects. Apart from visual landmarks the only other directional cues available to these subjects were some optic flow information offered by poles surrounding the arena and hand motor signals related to manipulations of the joystick. If subjects succeeded in ignoring potential landmarks in the Walking condition, then perhaps the Joystick condition would tell us something about how visual landmarks are chosen in this type of task.

Local and global objects were shifted in several possible combination from trial to trial. The magnitude of directional errors associated with each shift combination represented the degree to which subjects depended on that combination of objects to find the target location.

2.2 METHOD

Subjects

Five adult subjects participated in the Joystick experiment and Ten adult subjects participated in the Walking experiment. They received monetary compensation for their participation.

Virtual Environment for Experiment I

A green colored arena 1100 cm. in diameter served as the ground plane on which the subjects walked. White poles separated by 4 degree spaces stood along the arena's perimeter. Each pole was 65 cm. high. Beyond the arena a brown colored ground extended as far as the location of the global landmarks.

A 300 cm. high, 100 cm. wide tunnel stood in the arena. One opening of the tunnel was located just outside the bounds of the arena and the other end was located at the arena's center. The tunnel had two straight legs of equal length, joined by a 160° bend (see Figures 1a and 1b). This bend kept the participants from seeing the arena until they entered the second leg of the tunnel. The end of the tunnel that was located outside the arena served as the starting point and the opening of the tunnel at the center of the arena is where the subjects emerged and began their walk to the target pole. The target pole was situated 46° to the left of the straight ahead viewing direction when standing at the tunnel's point of emergence (POE).

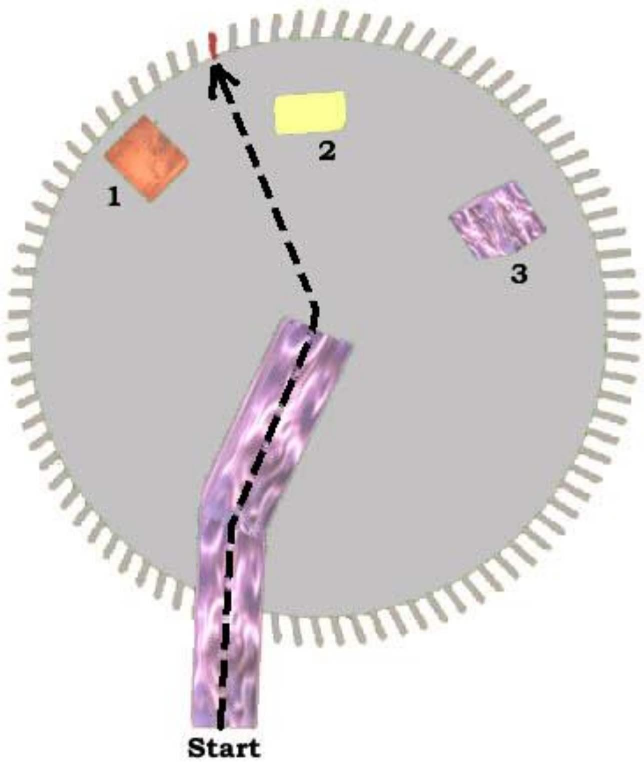

Figure 1a This diagram shows a bird's eye view of the environment in which subjects were tested. The dashed line shows the path to the target pole, beginning at the "Start" position. The numbered objects are the "local" objects. The numbers and the dashed line are included for the purposes of this diagram - they were not included in the actual environment.

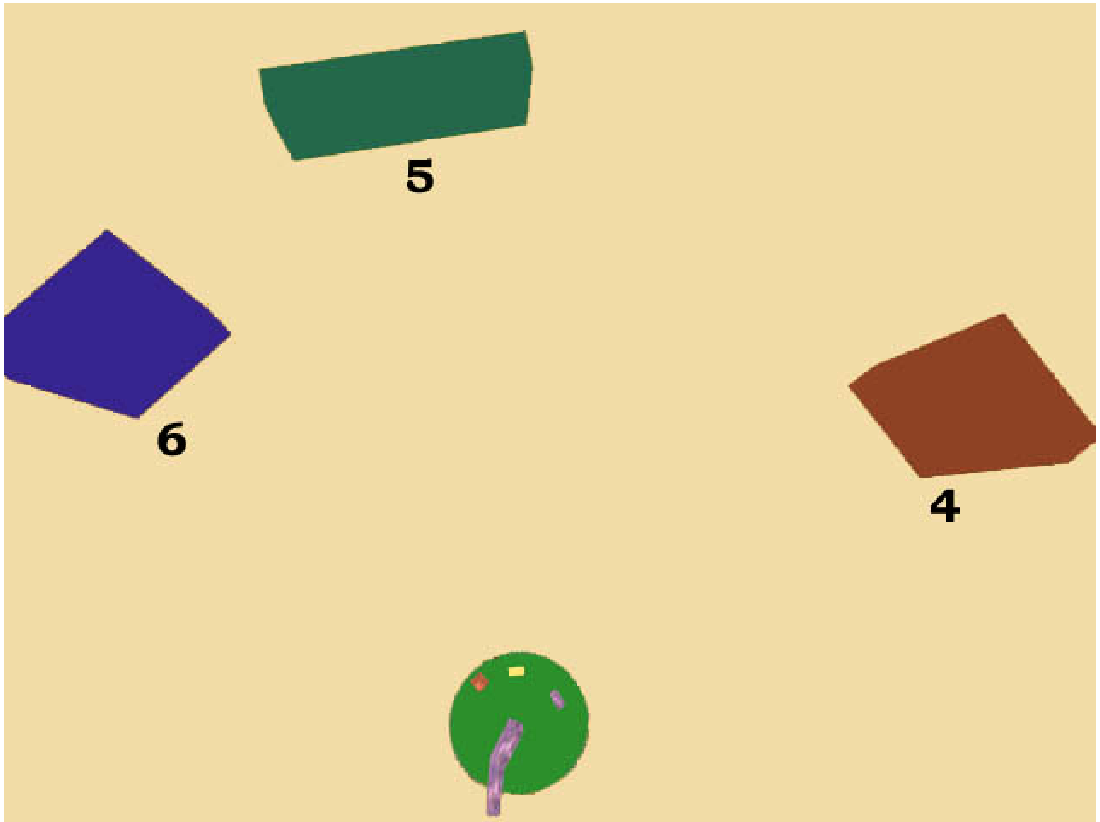

Figure 1b This wide-angle view of the virtual environment shows the "global" objects (numbered 4, 5, and 6) as well as the arena (bottom of diagram) in which the subjects were tested. The numbers in the diagram were not included in the actual environment.

Three rectangular blocks ("local objects") were situated inside the arena. The positions of all objects will be described in relation to the straight ahead viewing position at the POE. Block 1 was 450 cm. in front of and 70 degrees to the left of the POE. Block 2 (the center block) was 400 cm. in front of and 30° to the left of the POE. Block 3 was 350 cm. in front of and 30° to the right of the POE. Block 1 was colored red with a "terrain" texture. Its physical dimensions were: 100 cm.(X) x 100 cm.(Y) x 100 cm.(Z). Block 2 was colored yellow with a metallic texture. Its physical dimensions were: 120 cm.(X) x 100 cm.(Y) x 50cm.(Z). Block 3 was a pinkish-silver color with a "metallic" texture. Its physical dimensions were: 100 cm.(X) x 250 cm.(Y) x 65 cm.(Z).

Three larger rectangular blocks ("global landmarks") were situated at a distance outside the arena. Block 4 was 4000 cm. in front of and 70° to the left of the POE. Block 5 (the center block) was 5000 cm. in front of and 37° to the left of the POE. Block 6 was 4000 cm. in front of and 28° to the right of the POE. Block 4 was colored blue. Its physical dimensions were 1000 cm. (X) x 2000 cm. (Y) x 300 cm. (Z). Block 5 was colored green. Its physical dimensions were 2000 cm. (X) x 1000 cm. (Y) x 300 cm. (Z). Block 6 was colored red. Its physical dimensions were 1000 cm. (X) x 2000cm. (Y) x 300 cm. (Z).

Procedures

Training trials

The HMD was put on and adjusted to fit snugly on the subject's head. Calibrations were conducted to make sure that the visual display was centered and that the subject could see the display in 3D. The subject was led by the experimenter to one of two starting positions located at opposite corners of the room. When the subject was facing the opposite corner the tunnel was made to appear so that the subject was facing the entrance. The subject was then asked to walk through the tunnel and upon emerging into the arena to turn towards the "red" pole. Then the subject had to walk to the red target pole and "walk through it" at which time the experimental world disappeared, the trial was finished, and the subject had to adjust his or her position to prepare for the next trial.

In order to be oriented correctly for the beginning of the next trial the subject had to look for a reorienting pole (either blue or red) somewhere near where he or she was standing. Then the subject had to walk to the pole so that it was just in front of him or her. At that point the subject had to turn about 180 degrees so that he or she was facing the other reorienting pole in the distance (blue if the red reorienting pole was closer and vice versa). The distant reorienting pole was actually in the opposite corner of the room. When the subject was standing with his or her back to the near reorienting pole and facing the distant reorienting pole he or she was ready to begin the next trial. Thus there were two possible positions in the real room (at one of two reorienting poles) from which the subject began a trial. Of course, in the VE the starting position was always the same.

The subject completed 6 training trials during which he or she was asked to try to get a good sense of where the red target pole was located. The subject was informed that during the test trials the target pole would be colored white and would be indistinguishable from the other poles. In addition the subject was warned that the different colored blocks standing inside and outside the arena might shift positions randomly from trial to trial. Therefore the subject was given enough information to know that the blocks would be unreliable as visual landmarks and should be ignored. Lastly, the subject was told that throughout the test sequence a practice trial (with the red target pole and the original configuration of landmarks) would occur after every five test trials. This would ensure that the correct location of the target pole would be kept fresh in the subject's mind throughout the experiment.

After the training trials the subject had a good idea of where the target pole was located and so the test trials began.

The procedures were the same for the "Joystick" experiment. The only difference is that subjects were sitting down and using a joystick to "move" themselves through the virtual world. They wore the same HMD that was used in the walking experiments and received the same visual input.

Test trials

Test trials used the same procedure as the training trials except that the target pole was no longer colored red. It was white like the other poles surrounding the arena. Before each trial began the blocks standing inside and outside the arena shifted positions in different combinations, depending on the condition. Every six trials the target pole appeared red and the blocks appeared in the original positions, as in the training trials. This ensured that the subjects maintained a good memory of the target location throughout the experiment.

Subjects were told that they could take a break at any time during the experiment. They were asked to wait until the end of a trial before requesting a break. Finally, they were told that the experimenter would inform them when they were half-way through the experiment at which point they could also take a break if they wished.

The position of the pole that the subject walked to in each trial was recorded and used in the data analyses. The angular distance of the chosen pole from the target pole was deemed the error made during that trial.

Experimental Design

Two factors were crossed in this experiment: Shift Sizes and Landmark Shift Combinations. There were 6 different shift sizes: 18°, 12°, 6°, -6°, -12°, -18°. The angles represent how far the objects were shifted around the center of the arena (also the POE). Negative angles mean that objects were shifted to the left relative to the POE and positive angles mean that objects were shifted to the right relative to the POE.

Five different landmark combinations were used: *All, All Local, All Global, One Local,* and *One Global.* In the *All* condition all the objects were shifted together. In the *All Local* and *All Global* conditions either all the local objects or all the global objects were shifted together. In the *One Local* and *One Global* conditions either one local object (Object 2) or one global object (Object 5) was shifted.

With 6 shift conditions and 5 landmark conditions there were 30 possible conditions in all. There were 4 trials in each shift/landmark condition. Along with the training trials that occurred every six trials there was a total of 145 trials.

2.3 RESULTS

Experiment 1 - Joystick

The results for the Joystick task appear by landmark condition in Figure 2, and by shift condition in Figure 3. A two-factor, within subjects ANOVA showed significant main effects of both the shift, $F(5,20) = 4.86$, $p = .005$, and landmark combination, $F(4,16) = 7.24$, $p = .002$, conditions. The largest errors were made during the *All* landmark combination condition (mean error = 9.1°), followed by the *All Local* (mean error = 5.2°), *One Local* (mean error = 4.3°), *All Global* (mean error = 3.6°), and *One Global* (mean error = 2.1°) conditions. There was also a significant interaction between the shift and landmark combination conditions, $F(20,80) = 4.0$, $p = 0.000$. This interaction can best be understood when looking at whether or not the error means in each landmark combination condition are correlated with shift size (see Fig. 4). Specifically, the means in the *All* ($r^2 = 0.99$, $p = 0.0001$, slope = 3.92), *All Local* ($r^2 = 0.7$, $p = 0.05$, slope = 1.66), and *One Local* ($r^2 = 0.88$, $p = 0.006$, slope = 1.28) conditions were positively correlated with the shift sizes, whereas the means in the *All Global* ($r^2 = 0.08$, $p = 0.58$, slope = 0.34) and *One Global* ($r^2 = 0.0025$, $p = 0.93$, slope = -0.022) conditions were not.

Matched-pairs t-tests yielded significant differences between the *All* condition and all the other landmark combinations. The differences between the *All Local* and the *One Local*, t(4) = 2.0, p = 0.11, and *One Global*, t(4) = 1.99, p = 0.12, conditions approached significance. Also approaching significance was the difference between the *All Global* and *One Global* conditions, t(4) = 2.17, p = 0.09.

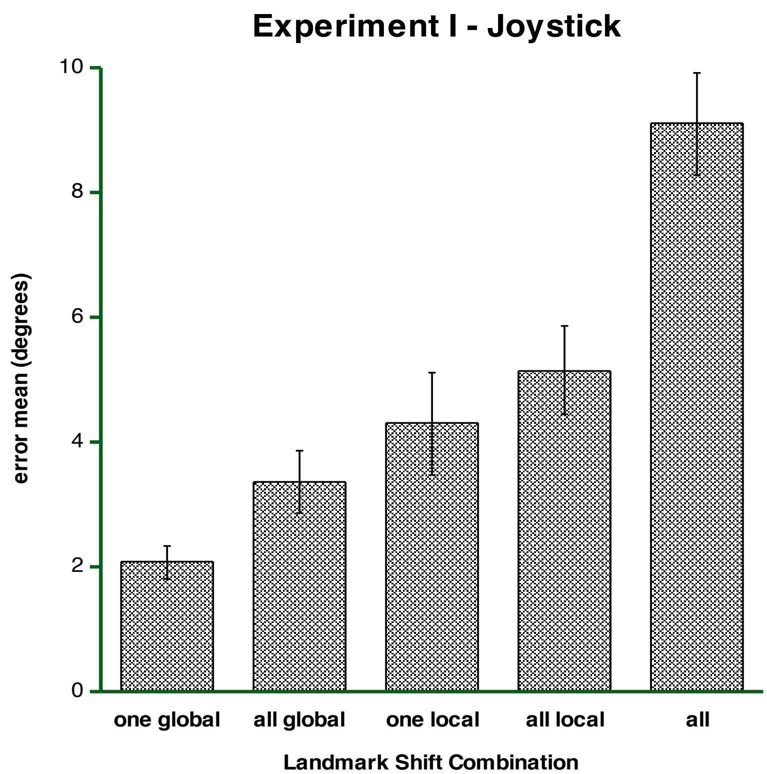

Figure 2 This graph shows the overall error means for each Landmark Shift Combination in Experiment I - Joystick.

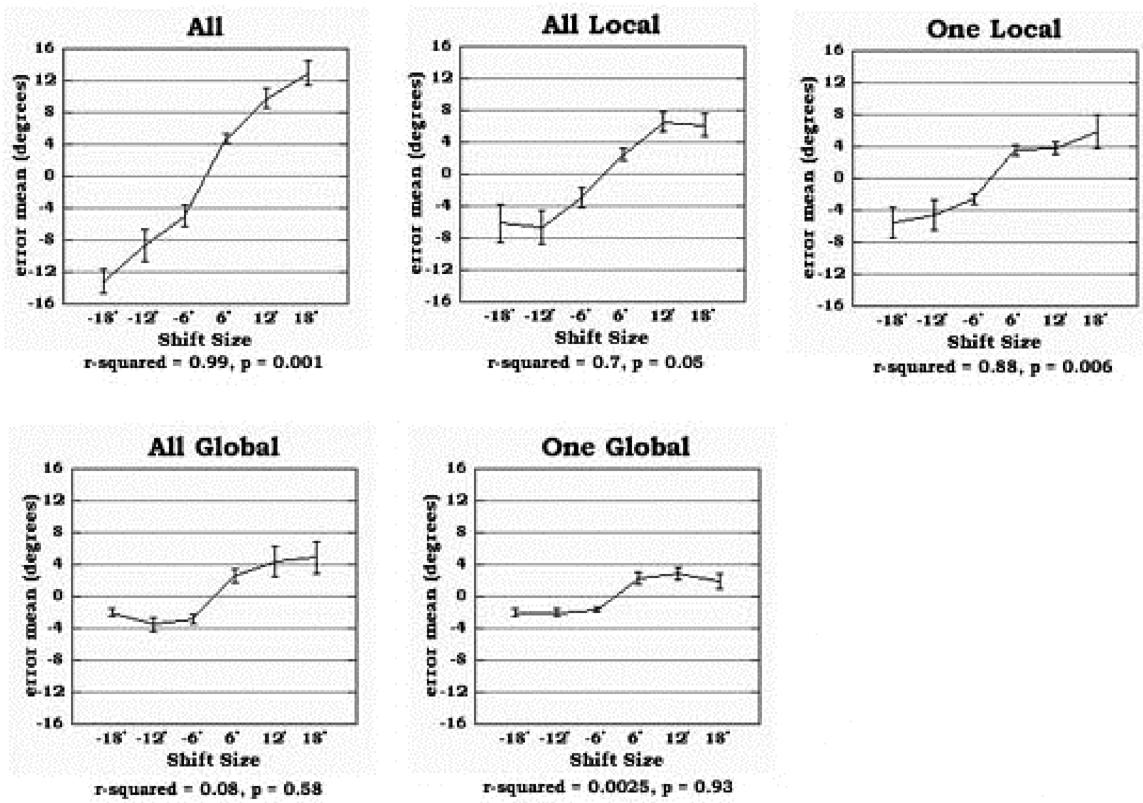

Figure 3 The above graphs show the error means in each Landmark Combination condition by each Shift Size in Experiment I - Joystick. The correlation coefficients are included for each Landmark Combination/Shift Size cross.

Experiment I - Walking

In the walking condition a two-factor, within subjects ANOVA resulted in significant main effects of both the Shift, $F(5,45) = 10.6$, $p < 0.001$, and the Landmark, $F(4,36) = 19.99$, $p < 0.001$ (see Figure 4) conditions. As in the Joystick experiment, the greatest errors occurred during the *All* landmark combination condition (mean error = 6.8°), followed by the *All Local* (mean error = 5.8°), *One Local* (mean error = 3.1°), *All Global* (mean error = 2.5°), and *One Local* (mean error = 2.1°) conditions.

To visualize the interaction between the Shift and Landmark combination conditions, $F(20,180) = 4.27$, $p = 0.000$, please see the correlation graphs for each Landmark condition in Figure 5. As in the Joystick experiment the error means in the *All* ($r^2 = 0.76$, $p = 0.02$, slope = 2.24), *All Local* ($r^2 = 0.85$, $p = 0.009$, slope = 2.1), and *One Local* ($r^2 = 0.82$, $p = 0.01$, slope = 0.61), conditions were positively correlated with the shift sizes, whereas the error means in the *All Global* (r^2 = 0.36, p = 0.21, slope = 0.29) and *One Global* (r^2 = 0.007, p = 0.88, slope = -0.04) were not.

Matched-pairs t-tests showed that the errors in the *One Global* condition were significantly lower than the errors in the *One Local*, t(9) = 4.0, p = 0.003, *All Local, and All* conditions. The difference between the *One Global* and *All Global* conditions was not significant, t(9) = 1.88, p = 0.09. The *All Local* errors were significantly lower than the *All* errors, t(9) = 2.62, p = 0.03. Both the *All Local* and *All* error means were significantly larger than the error means in the *One Local, All Global,* or *One Global* conditions.

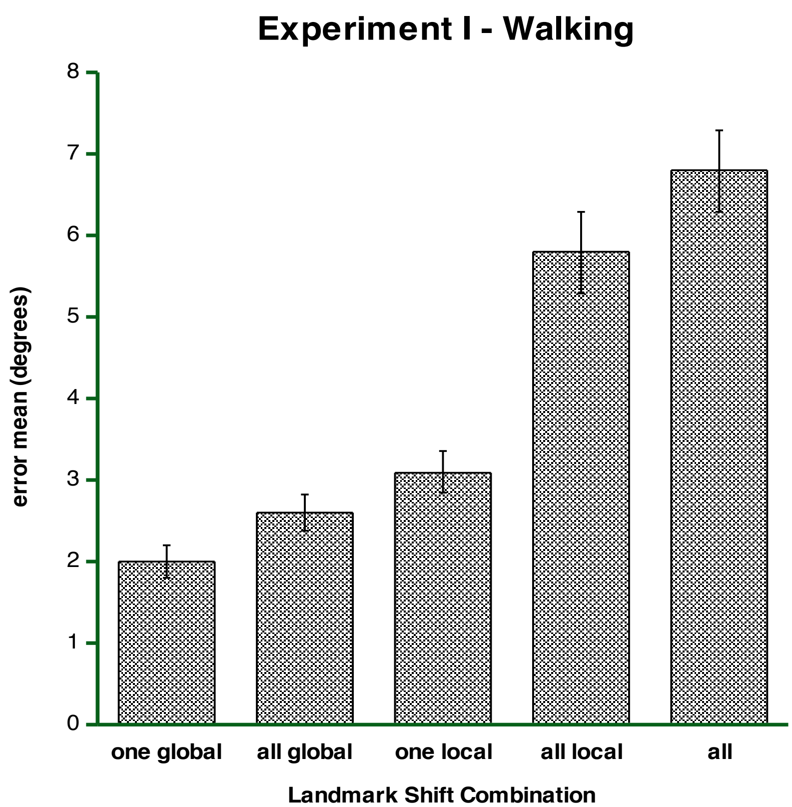

Figure 4 The above graph shows the overall error means for each Landmark Shift Combination in Experiment I - Walking.

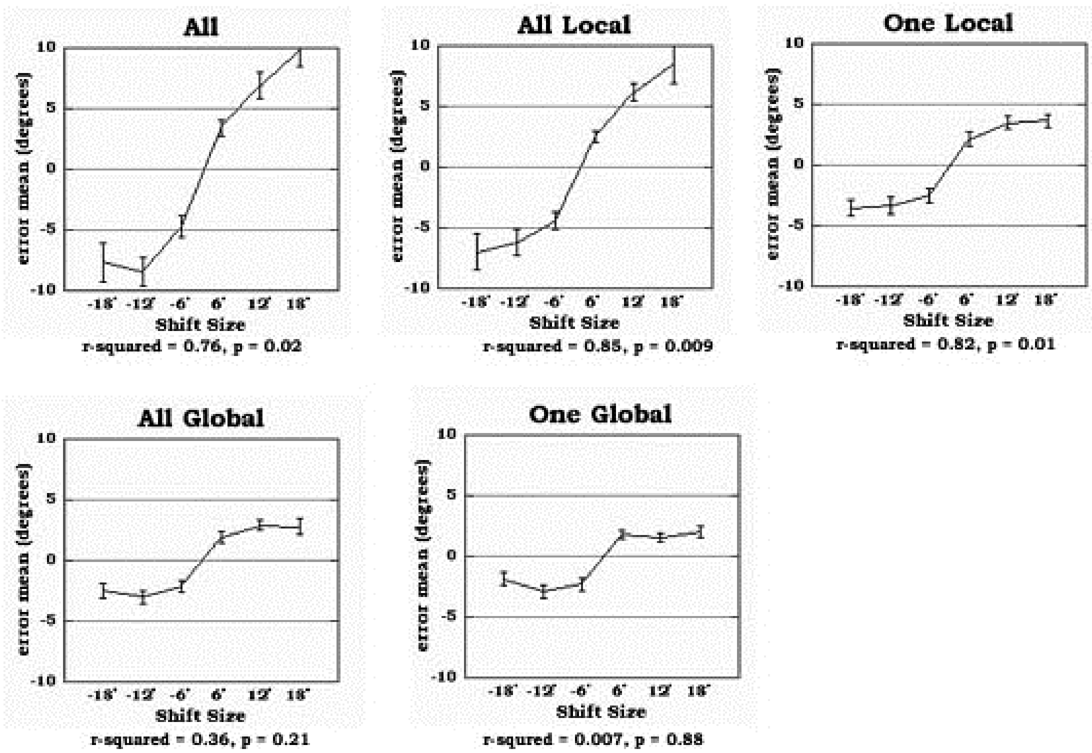

Figure 5 The above graphs show the error means in each Landmark Combination condition by each Shift Size in Experiment I - Walking. The correlation coefficients are included for each Landmark Combination/Shift Size cross.

2.4 DISCUSSION

Experiment I-Joystick and Experiment I-Walking shared the same error pattern. The greatest final position errors away from the target pole were made when *All* the objects were shifted together. The rest of the error means in descending order were in the *All Local, One Local, All Global,* and *One Global* conditions. It was not surprising that shifting the entire array of objects had the biggest effect on subjects' performances. If any of the objects were used as visual landmarks then shifting all of them guaranteed that the crucial visual cues were displaced. Furthermore, because the configuration of the array didn't change it was impossible to determine the extent of any one object's displacement relative to other elements in the array. Therefore if subjects were going to make any errors in their movement towards the target pole they would occur in at least the *All* condition.

One of the most interesting findings in this experiment was that *both* the Joystick and the Walking groups made directional errors that were associated with object shifts. Not surprisingly, the error angles were larger in the Joystick experiment than in the Walking experiment. In the *All* landmark condition of the Joystick experiment the errors were sometimes as great as the magnitudes of object shifts themselves. This showed that it was difficult for subjects to use anything other than the visual information offered by the arrangement of objects. In the Walking experiment the errors were smaller than the object shifts. These results were also expected because the Walking experiment offered subjects the use of additional information from internally generated signals. This information could be used to assuage the errors that were induced by displaced visual landmarks. However small, though, in several Walking conditions the errors did follow the patterns of object shifts, showing that path integration could not completely override landmark-driven navigation.

Specifically, the highest correlations between errors and shift magnitudes occurred in the *All, All Local,* and *One Local* object shift conditions. This implies that subjects could not resist using the local objects as landmarks when they tried to guess the location of the target pole. The results also suggest that just one local object was having some influence on location choices, more so than all of the global objects combined.

If local, but not global landmarks were being used in this task the first experiment could not reveal exactly how these local landmarks were being used. Since either all or one of the local objects were displaced it is not yet clear whether or not one local object or a particular configuration of local objects was used. For instance, the results would be consistent with a strategy of using the two landmarks closest to the target pole. Therefore when one of them moved the errors decreased but were still greater than errors associated with movements of other non-attended objects. The results could also support a strategy of using the entire configuration of local objects as a landmark. In that case, when the configuration broke down, the subjects simply guessed which elements of the original configuration they wanted to follow and on some occasions guessed wrong. Lastly, the results could also mean that each local object had some degree of influence on the subjects' decisions and that the individual effects of each object could be added to yield the final magnitude of error. Experiment

II was designed to further investigate the question of how local objects were being used as landmarks.

3 Experiment II - Local Objects Only

3.1 INTRODUCTION

Experiment I showed that local objects were preferred over global objects to serve as landmarks in the task of locating the target pole. This experiment was conducted to better understand how the local objects were used to solve this navigational problem. In the previous experiment either all local objects or one local object were shifted in various trials. Both types of shifts proved to have an effect on where subjects walked. However, it was unclear whether each local object or only the middle object, Object 2, was deemed a useful landmark. In this experiment the three local objects will be shifted more systematically so that a possible influence of each object or any combination of the three can be revealed. Besides shifting all three objects at the same time, we will also observe the effects of shifting each object individually, as well as the effects of shifting all possible combinations of two objects.

Global objects were not included in this environment. As they did not exert a substantial influence on subjects' performance in the previous experiment their presence was not thought to be of interest in this experiment. Their absence, however, may be more informative. If subjects react to local object shifts (*All Local* or *One Local*) in the same way in this experiment as they did in the last then the inconsequential role of the global objects will be confirmed. If, however, a different pattern of errors emerges it will mean that global objects did in some way affect the selection of visual cues in the previous experiment.

3.2 METHOD

Subjects

Ten adult subjects participated in this experiment for monetary compensation.

Environment

The environment in Experiment II was the same as in Experiment I with the exception that there were no global objects. Only the three smaller blocks (the local objects) were used in this environment.

Procedures

The procedures in Experiment II were the same as the procedures in Experiment I.

Design

Two factors were crossed in this experiment: Shift Sizes and Landmark Shift Combinations. Since the previous experiment showed that in some conditions error were, in fact, correlated with the shift angles it was considered unnecessary to include as many different shift sizes in this experiment. To decrease the number of trials only 4 different shift sizes were included: 12°, 6°, -6°, -12°.

Seven different landmark combinations were used: *All, Object 1, Object 2, Object 3, Objects 1&2, Objects 1&3,* and *Objects 2&3.* In the *All* condition all the objects were shifted together. The names of the other conditions state exactly which object or objects were displaced.

With 4 shift conditions and 7 landmark conditions there were 28 possible conditions in all. Each shift/landmark cross occurred in 4 trials. Along with the training trials that occurred every six trials there was a total of 135 trials.

3.3 RESULTS

A two-factor, within subjects ANOVA yielded significant main effects of both the Shift, $F(3, 27) = 3.19$, $p = .04$, and the Landmark, $F(6,54) = 4.92$, $p = .000$ (see Figure 6) conditions. The largest error means occurred during the *Objects 1&2* (mean error = 4.96°), followed by the *Object 2* (mean error = 4.56°), *All* (mean error = 4.48°), *Objects 2&3* (mean error = 4.32°), *Objects 1&3* (mean error = 3.35°), *Object 1* (mean error = 3.19°), and *Object 3* (mean error = 2.84) conditions.

The interaction between the Landmark and Shift conditions was not significant. The errors in the *Object 2* condition were positively correlated with the shift sizes (r^2 =

0.98, p = 0.012, slope = 0.97). The regression analyses for the *Objects 1&2* (r^2 = 0.78, p = 0.12, slope = 1.25), *All* (r^2 = 0.8, p = 0.10, slope = 1.56), *Objects 2&3* (r^2 = 0.94, p = 0.033, slope = 1.93), and *Objects 1&3* (r^2 = 0.76, p = 0.13, slope = 0.63) approached significance. There was no significant correlation between the shift sizes and the *Object 1* (r^2 = 0.22, p = 0.53, slope = 0.47) and *Object 3* (r^2 = 0.0002, p = 0.99, slope = 0.003) conditions (see Figure 7).

Matched-pairs t-tests showed that the errors in the *Object 2* condition were significantly larger than the errors in the *Object 1*, t(9) = 2.48, p = 0.035, and the *Object 3*, t(9) = 4.1, p = 0.003, and *Objects 1&3*, t(9) = 2.17, p = 0.058 conditions. There were no significant differences between the error means of the *Object 2, Objects 2&3, Objects 1&2, and All* conditions. Similarly there were no significant differences between the *Object 3, Object 1,* and *Objects 1&3* conditions. The *Objects 1&3* error mean was not significantly lower than the *Objects 2&3* error mean, t(9) = 1.4, p = 0.19, but it was significantly lower than the *Objects 1&2* error mean, t(9) = 2.67, p = 0.025. The *All* error mean was not significantly higher than the *Objects 1&3* error mean, t(9) = 1.4, p = 0.19.

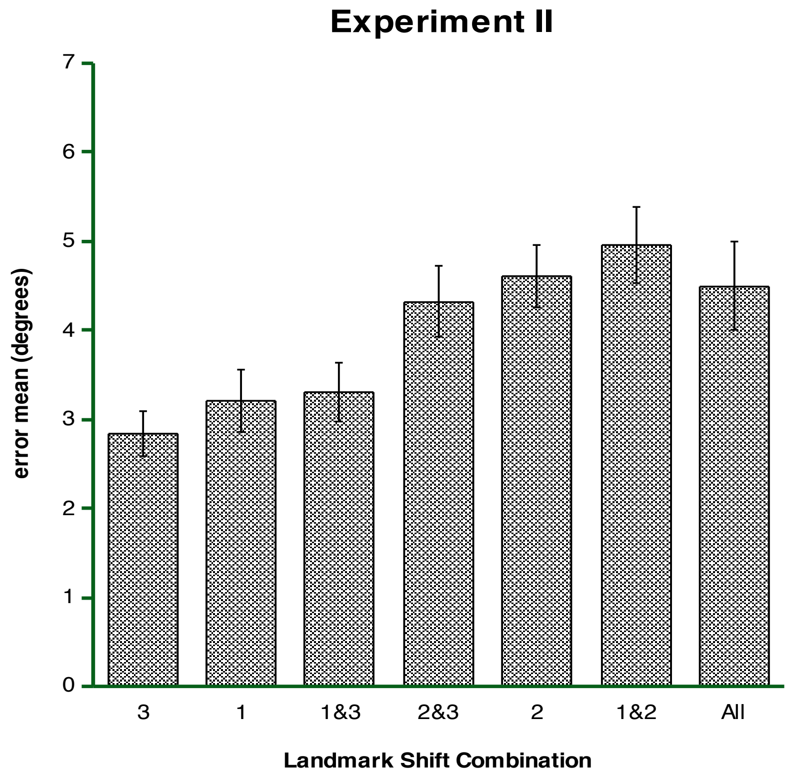

Figure 6 The above graph shows the overall error means for each Landmark Shift Combination in Experiment II - Local Only.

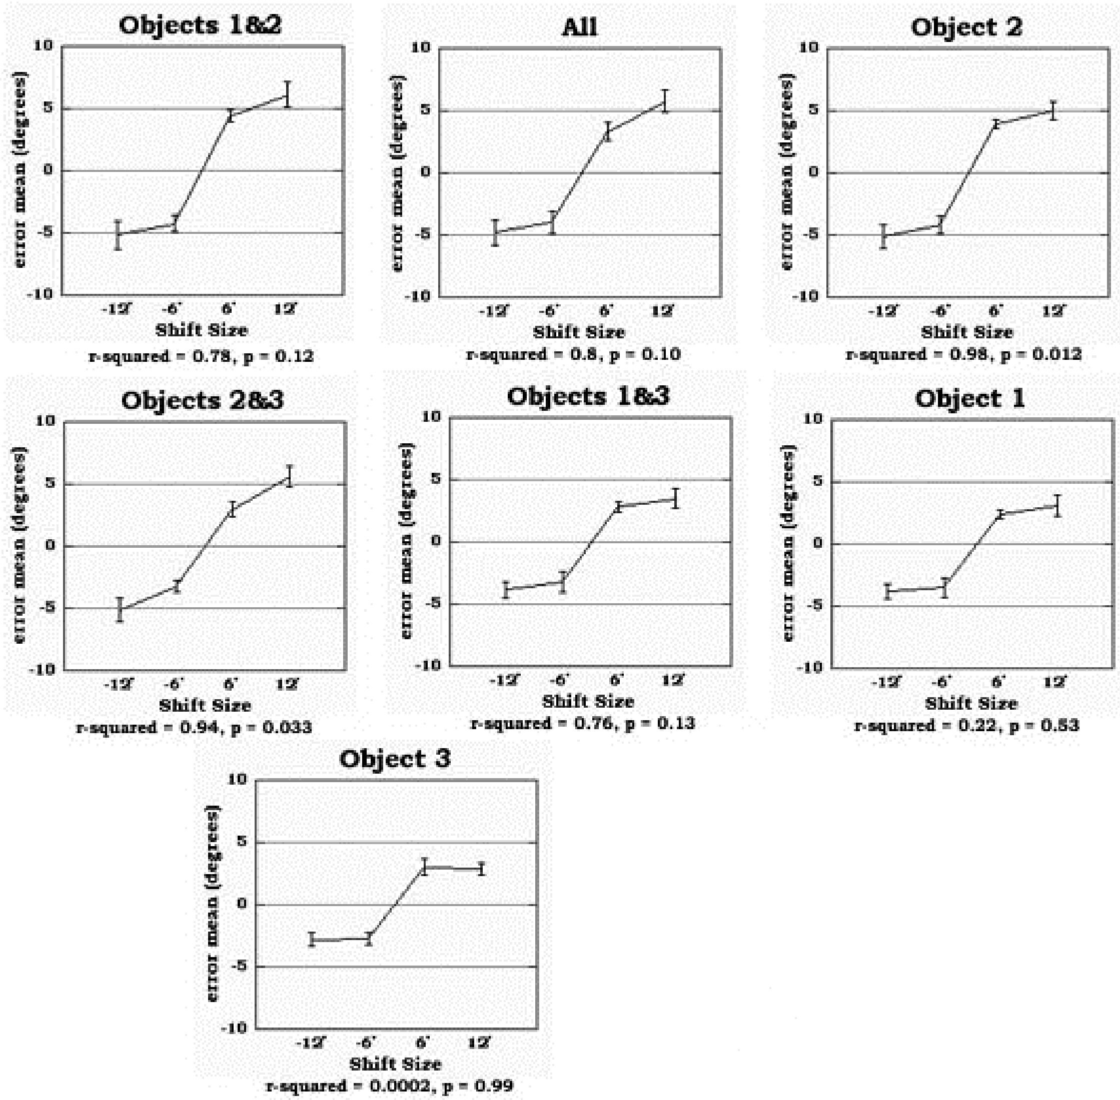

Figure 7 The above graphs show the error means in each Landmark Combination condition by each Shift Size in Experiment II. The correlation coefficients are included for each Landmark Combination/Shift Size cross.

3.4 DISCUSSION

In Experiment II only three local objects were available in the virtual arena and combinations of one, two, or three objects were shifted in each trial. The resulting errors away from the target pole were greatest when the object standing closest to the target pole, Object 2, was displaced. These errors were roughly the same size in all conditions that involved the movement of Object 2.

The farthest object, Object 3, was associated with the smallest errors. Therefore the farthest object was not used individually to find the target pole position. Object 1, the second closest object, also did not influence subjects' orientation when it was the only object displaced. The error mean increased slightly, but not significantly when Objects 1 and 3 shifted together.

Overall the results of this experiment indicate that only one object - the one closest to the target pole - was used as a visual landmark in this task. This outcome reveals an interesting inconsistency between this experiment and Experiment I. The inconsistency lies in the errors associated with local landmarks. In the previous experiment conditions in which all three local landmarks were displaced yielded significantly larger errors than conditions in which only Object 2, the closest object, was displaced. In the current experiment displacing Object 2 alone caused the same effect as displacing all three objects. The logical conclusion, therefore, is that the global objects included in the previous environment played a role in the subjects' performance after all. Although they were not used directly to determine the location of the target pole, they were evidently used as a visual context against which the configuration of local objects could be examined. When the local array was disturbed by a local object's displacement, the context provided by the global landmarks disambiguated the puzzle of which objects moved and which stayed in the original position. This additional information accounted for the error decrease in Experiment I when Object 2 moved by itself. However, it is important to note that the global context did not completely eliminate the error, showing that Object 2 must have had a very strong influence on subjects' directional decisions.

It is logical to assume that the proximity of Object 2 to the target pole made it the chosen visual reference relative to which the target pole position was gauged. However, the placement of Object 2 was such that another plausible explanation for its influence can be proposed. Relative to the subject's point of emergence into the arena, the position of Object 2 fell between the vector representing the subject's straight-ahead direction and the vector along which the target pole could be found. Therefore Object 2 was uniquely situated to move across the subject's visual field as he made the left turn in preparation to walk to the target pole. It follows, then, that the subject could potentially have used the movement of the object across the visual field to monitor the speed and duration of his rotation. Furthermore, on a purely practical level this object's "landmark" status could be earned simply because it fell within the subject's field of view for a longer total period of time than the other two objects.

Experiment II showed that a single object closest to the target pole was used as a visual landmark when walking to the pole. However, this experiment could not resolve whether this occurred because simply because it was the closest object or because it appeared in the subjects' line of sight more than the other objects. The next experiment was conducted to further investigate this issue.

4 Experiment III - Equidistant Objects

4.1 INTRODUCTION

Experiment II revealed that subjects were mainly relying on one local object, Object 2, to solve the problem of finding the target pole. This object was located closer to the target pole than the other two objects. It was also positioned in such a way that it fell directly within the arc that was traced by the subject's forward gaze as he turned to face the target pole. Therefore, unless the subject made many intentional head turns to look at the other two objects, Object 2 fell within his field of view for the longest period of time during each trial. This may have allowed the subject to extract more relevant information regarding his orientation in the environment from his spatial relationship with Object 2 than from his spatial relationships with any other objects.

Experiment III was designed to further investigate which strategy compelled subjects to use Object 2. In this experiment Object 1 was shifted in the original configuration so that it was now located the same distance away from the target pole as Object 2. However it was still standing to the left of the target pole while Object 2 remained to the right of the target pole if viewed by the subject from the point of emergence (POE). Hence, Objects 1 and 2 not only stood at points that were equidistant from the target pole, they *looked* like they were equidistant from the target pole when viewed from the POE.

If Object 2 was favored as a landmark in Experiment II because it was the closest object to the target pole then this time Object 1 would have just as much influence. This reason for using an object as a landmark will be referred to as the "Closest Object" strategy. If, however, Object 2 is favored because it maintains a prominent position in the subjects visual field as he turns and walks towards the target pole, then this object will still be weighted more heavily than Object 1. This reason for attending to Object 2 will be referred to, for a lack of a better term, as the "*En Route* Object" strategy.

This experiment will present another potential heuristic for encoding the configuration of Objects 1 and 2. Subjects will be able to regard the relationship between those two objects and target pole as one in which the target pole is in the center of the configuration created by the three elements. This will be referred to as the "Configuration" strategy. If this strategy is adopted in cases when Objects 1 and 2 are both shifted then the subsequent errors will equal the magnitude of the shifts. If this strategy is also used when the two objects are shifted individually, then the errors should look the same for each object. Specifically, they should turn out to be half the size of each object's shifts.

In order to provide a baseline against which all the errors in this experiment can be measured - another condition was introduced in which no shifts were made at all. This *No Shift* condition will help to define the extent to which accuracy decreases when object displacements occur. That is, it will show how well subjects perform the task when correctly positioned landmarks are available.

4.2 METHOD

Subjects

Ten adult subjects participated in this experiment for monetary compensation.

Environment

The environment in Experiment II was the same as the environment in Experiment II, with the exception that Objects 1 and 2 were now equidistant from the target pole. Both objects were 400 cm. away from the POE. Object 1 was 64° to the left of the POE (18° to the left of the target pole) and Object 2 was 28° to the left of the POE (18° to the right of the target pole). Objects 1 and 2 were separated by an angle of 36°.

Procedures

The procedures in this experiment were the same as the procedures in Experiments I and II.

Design

Two factors were crossed in this experiment: Shift Sizes and Landmark Shift Combinations. The shift sizes were : 10°, 6°, 0°, -6°, -10°.

Seven different landmark combinations were used: *All, Object 1, Object 2, Object 3, Objects 1&2, Objects 1&3,* and *Objects 2&3.* In the *All* condition all the objects were shifted together. The names of the other conditions state exactly which object or objects were displaced.

With 5 shift conditions and 7 landmark conditions there were 35 possible conditions in all. Each shift/landmark cross occurred in 4 trials. Along with the training trials that occurred every six trials there was a total of 169 trials.

4.3 RESULTS

A two-factor, within subjects ANOVA yielded significant main effects of both the Shift, $F(4,36) = 4,36$, $p = .000$, and the Landmark, $F(6,54) = 17.1$, $p < .000$, conditions (see Figure 8). The largest errors occurred in the *All* landmark combination condition (mean error = 6.2°), followed by the *Objects 1&2* (mean error = 5.5°), *Objects 2&3* (mean error = 4.65°), *Object 2* (mean error = 4.14°), *Objects 1&3* (mean error = 4.0°), *Object 1* (mean error = 3.13°), and *Object 3* (mean error = 2.1°) conditions. The *Object 3* mean error was the same as the mean error in the *NoShift* shift condition (2.1°).

There was also a significant interaction between the two factors, $F(24,216) = 3.9$, $p = .000$. There was a positive correlation between the shift sizes and the errors in the *All* ($r^2 = 0.98$, $p < 0.001$, slope = 4.0), *Objects 1&2* ($r^2 = 0.95$, $p = 0.005$, slope = 3.43), *Objects 2&3* ($r^2 = 0.95$, $p = 0.004$, slope = 2.3), *Object 2* ($r^2 = 0.88$, $p = 0.02$, slope = 1.73), and *Objects 1&3* ($r^2 = 0.99$, $p = 0.00$, slope = 1.6) conditions. The correlation between shift sizes and the *Object 1* condition approached significance ($r^2 = 0.74$, $p = 0.06$, slope = 0.77). There was no significant correlation between shift sizes and the *Object 3* condition ($r^2 = 0.4$, $p = 0.23$, slope = -0.14)(see Figure 9).

Matched-pairs t-tests showed that the errors in the *Object 2* condition were significantly higher than the errors in both the *Object 1*, $t(9) = 2.58$, $p = 0.03$, and *Object 3*, $t(9) = 4.32$, $p = 0.002$, conditions. By contrast errors in the *Object 2* condition were significantly lower than errors in the *Objects 2&3*, $t(9) = 2.4$, $p = 0.04$, *Objects 1&2*, $t(9) = 2.86$, $p = 0.02$, and *All* conditions.

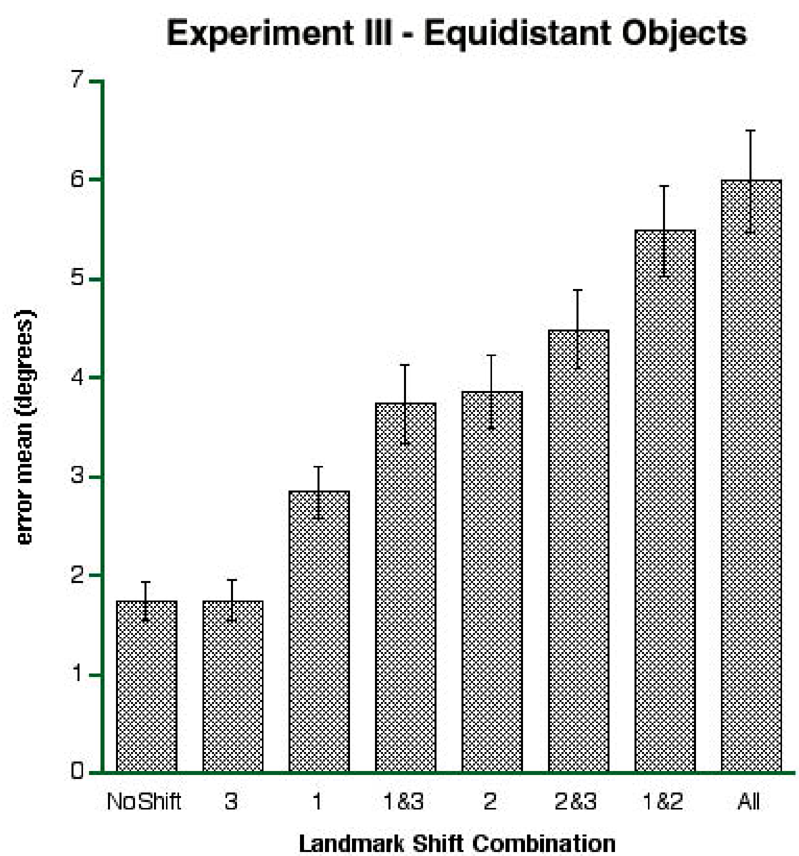

Figure 8 The above graph shows the overall error means for each Landmark Shift Combination in Experiment III - Equidistant Objects.

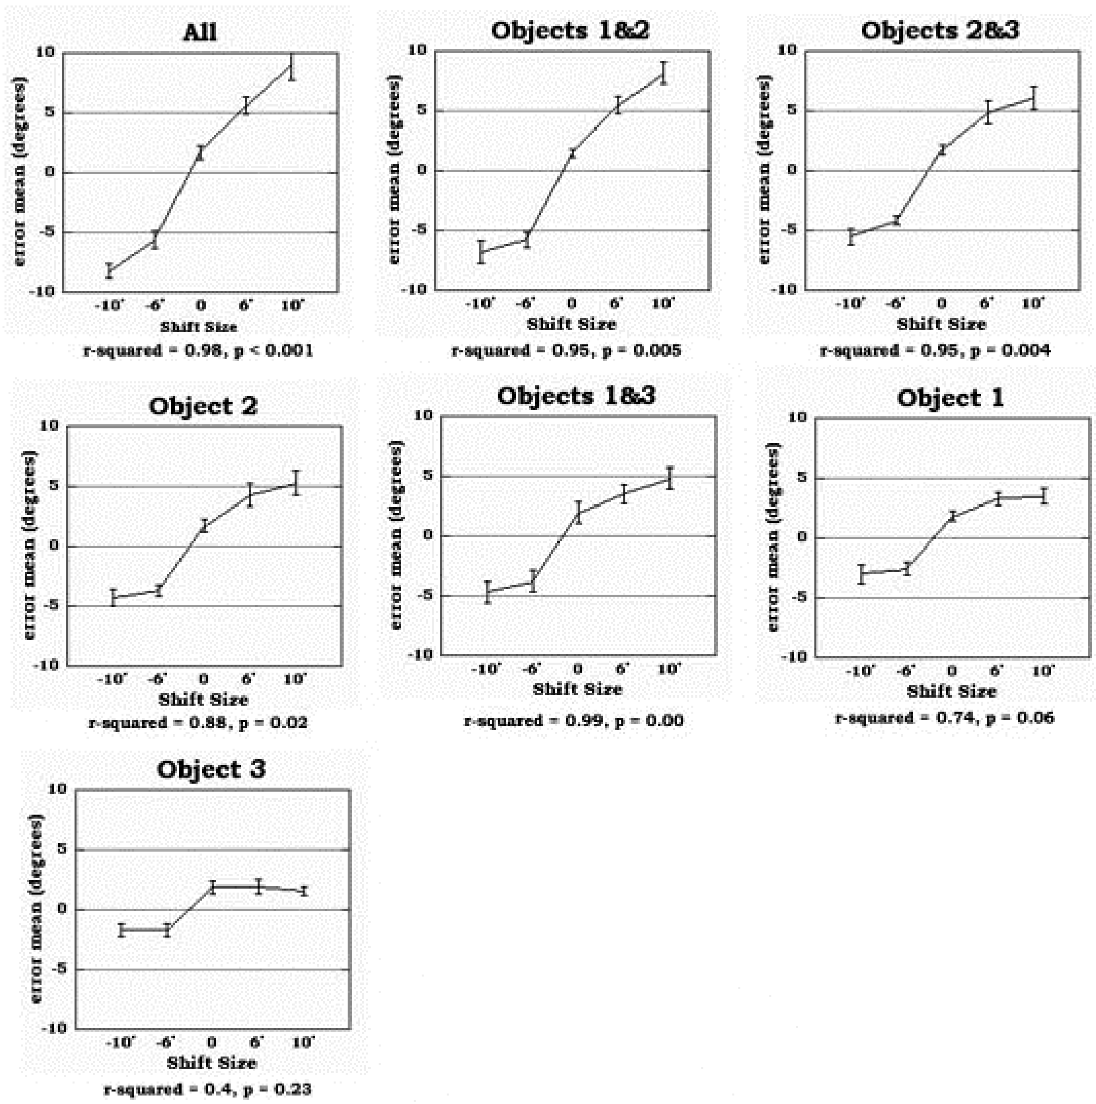

Figure 9 The above graphs show the error means in each Landmark Combination condition by each Shift Size in Experiment III. The correlation coefficients are included for each Landmark Combination/Shift Size cross.

4.4 DISCUSSION

In this experiment the *All* and *Objects 1&2* shifts caused the largest errors, with the *Objects 1&2* error mean being slightly smaller than the *All* error mean. *Object 2* shifts caused larger errors than *Object 1* shifts. However, this time *Object 1* shifts caused significantly greater errors than both *Object 3* shifts and *No Shifts* at all.

Shifting a combination of *Objects 2&3* yielded higher errors than shifting a combination of *Objects 1&3*. The displacement of Object 3 did not influence subjects' performance any more than when *No Shifts* were made. Therefore, on its own Object 3 was not used as a landmark. It did, however, seem to exert an effect when it was shifted in combination with other objects or when it was the only object that was *not* displaced. This indirect influence of object 3 will be discussed later in this section.

First, I'd like to address the question of whether the "Closest Object" strategy, the "Configuration" strategy, or the "*En Route* Object" strategy caused Object 2 to be favored as a landmark over the other objects. If the "*En Route* Object" strategy *was not* used then individual displacements of Objects 1 and 2 would have caused equal magnitudes of error. This is because both the "Closest Object" and "Configuration" strategies would have entailed equally weighting both objects as landmarks demarcating the target pole position. However the results show that Object 2 shifts still caused higher errors than Object 1 shifts, suggesting that the "*En Route* Object" or some related strategy must have contributed to performance.

Another noteworthy result in this experiment was that in contrast to Experiment II *Object 1* shifts now caused significantly larger errors than *Object 3* shifts or *No Shifts* at all. This means that making Objects 1 and 2 equidistant from the target pole elevated Object 1 to a "landmark" status along with Object 2, albeit not to as high a rank.

Evidence that the "Configuration" strategy may have been utilized in certain situations can be found in the results of the *All* and *Objects 1&2* conditions. Whereas in previous experiments the errors were, on average, less than half of the magnitudes of the object shifts, they were now nearly as large as the object shifts in these two conditions. This means that, for the most part, subjects were choosing a final position that maintained the same relationship with both objects, as did the target pole in the original configuration. That is, they were choosing a pole that was located in the center of the two-object configuration.

Taken together the above observations imply that in each situation subjects chose from a hierarchy of useful object/target spatial relations. When the two nearest objects maintained the original spatial relationship with each other subjects felt most confident looking for a target directly between the two objects. When this relationship was altered subjects switched to using Object 2 as the preferred landmark. This was because the position of Object 2 met the requirements of both the "Closest Object" and the "*En Route* Object" relationship relative to the target pole. Lastly, Object 1 was also occasionally used as a landmark because it too was a "Closest Object". Not surprisingly, when a second justification for using Object 1 was available, i.e., when it moved in conjunction with Object 3, then Object 1 was as influential as Object 2.

In order to ascertain the relative weights of the "Closest Object" and the "*En Route* Object" positions it will be necessary to create an experiment with an original configuration in which Object 1 is closer to the target pole than Object 2. Then Object 1 will be associated with the "Closest Object" strategy, while Object 2 will only meet the requirements of the "*En Route* Object" strategy. Such an experiment will make an interesting addition to a future study.

5 Experiment IV - No Landmark Control

5.1 INTRODUCTION

The results of the preceding experiments showed that subjects could not completely disregard the placement of objects even when reliable non-visual information was available to them. It may seem obvious that before going any further the following question had to be addressed. How well could the subjects have performed on the basis of path integration alone, without using any visual landmarks? After all, it would be foolish to consider the above findings remarkable if without visual landmarks people were completely lost in this task. Therefore a group of subjects was tested in a No Landmark Control experiment. As the name indicates there were no potential landmarks included in this environment. The outcome of this experiment would tell us whether or not subjects could have demonstrated an adequate level of accuracy in the previous Walking experiments had they not paid attention to the visible objects.

5.2 METHOD

Subjects

Ten adult subjects participated in this experiment for monetary compensation.

Virtual Environment

The environment in this experiment was the same as the environments in the previous experiments with the exception that no local or global objects were included.

Procedures

The procedures in this experiment were the same as the procedures in the previous experiment except that the subjects were not warned that any objects were going to shift positions…(as there were no objects to be seen).

Design

There were 50 trials in this experiment which contained only one condition.

5.3 RESULTS

The mean error for this group of subjects was 4.64° (0.16° s.e.) away from the position of the target pole. The graph in Figure 10 contrasts the results of this condition and the error means of the *No Shift* control condition in Experiment III.

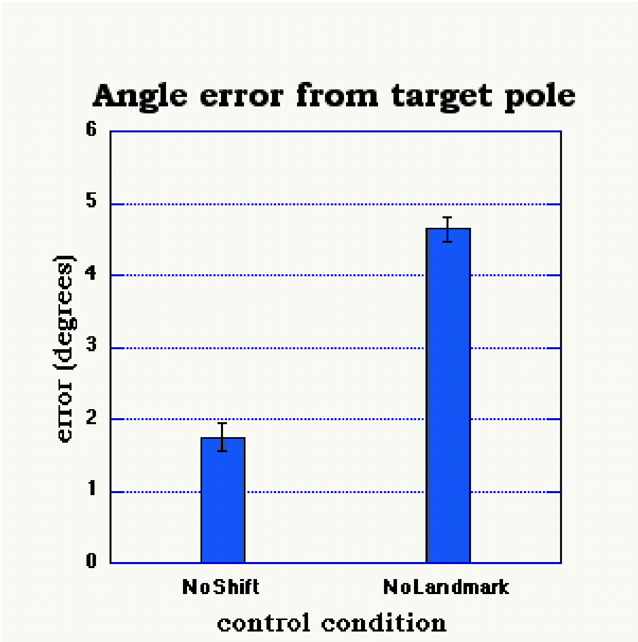

Figure 10 This graph shows the error means from the No Landmark condition (Experiment IV) and the No Shift condition of Experiment III.

5.4 DISCUSSION

Without any available landmarks subjects were able to perform the task reasonably well. When compared to the error mean of the *No Shift* condition of Experiment III, however, it is clear that people can perform the task with better accuracy when stable landmarks are available. Therefore it is a smart spatial learning system that automatically processes landmark information, as this type of information allows the traveler to find specific locations with better accuracy.

Although according to the mean error values it may seem that subjects performed no differently in most conditions of the preceding three experiments than they would have performed if no landmarks were available - that is not the case. The fact that the errors in some of the conditions were correlated with the sizes of the object shifts shows that the object shifts did, in fact, influence subjects' directional decisions.

6 Experiment V - Equidistant Objects With Global Objects

6.1 INRODUCTION

In Experiment I global objects provided a visual background relative to which the positions of local objects could be encoded. However, even with the availability of contextual information, subjects still erred when all local objects or one local object shifted to different positions. Their errors represented the extent to which they were relying on local objects to determine the position of the target pole.

The current experiment will recreate Experiment I, with the exception that it will now include the local object configuration that was used in Experiment III. That is, the target pole will now be half-way between objects 1 and 2. It would be interesting to see whether or not global objects will play a similar role when the target pole holds a simpler spatial relationship with two local objects. For instance, looking for a pole directly between two local objects may be such an easy heuristic that any other visual information such as global objects may be completely ignored. On the other hand, global objects may become an important back drop against which subjects can measure local inter-object distances.

As in Experiment III this experiment will also include a *No Shift* condition. Again, it will offer a baseline against which we can compare the results of other conditions. Furthermore, we will have another *No Shift* error mean which can be compared to that of Experiment III. This will give us a sense of how stable this measure is across different groups of subjects.

6.2 METHOD

Subjects

Eight adult subjects participated in this experiment for monetary compensation.

Environment

The environment was the same as the environment in Experiment I, with the exception that now Objects 1 and 2 were equidistant from the target pole.

Procedures

The procedures in this experiment were the same as the procedures in Experiments I, II, and III.

Design

Two factors were crossed in this experiment: Shift Sizes and Landmark Shift Combinations. The shift sizes were : 10°, 6°, 0°, -6°, -10°.

Seven different landmark combinations were used: *All, All Local, All Global, One Local,* and *One Global.* These landmark combinations were the same as the combinations used in Experiment I.

With 5 shift conditions and 5 landmark conditions there were 25 possible conditions in all. Each shift/landmark cross occurred in 4 trials. Along with the training trials that occurred every six trials there was a total of 121 trials.

6.3 RESULTS

A two-factor, within subjects ANOVA yielded significant main effects of both the Shift, $F(4,28) = 10.9$, $p = 0.000$, and the Landmark conditions, $F(4,28) = 16.4$, $p < .000$. As in Experiment I, the largest errors were found in the *All* landmark combination condition (mean error = 5.4°), followed by the *All Local* (mean error = 4.4°), *One Local* (mean error = 3.0°), *One Global* (mean error = 2.75°), and *All Global* (mean error = 2.7°) conditions (see Figure 11).

There was a significant interaction between the Shift and Landmark conditions, $F(16,112) = 3.6$, $p = 0.000$. This interaction can be understood by looking at the correlation graphs of Figure 12. The correlations between the shift sizes and error means approached significance in conditions *All* ($r^2 = 0.72$, $p = 0.07$, slope = 3.0), *All Local* ($r^2 = 0.65$, $p = 0.09$, slope = 2.1), and were significantly correlated in the *All Global* ($r^2 = 0.85$, $p = 0.026$, slope = 0.64) condition. They were not significantly correlated with the errors of the *One Local* ($r^2 = 0.23$, $p = 0.41$, slope = 0.86) or *One Global* ($r^2 = 0.005$, $p = 0.9$, slope = 0.05) conditions.

Matched-pairs t-tests showed that errors in the *All Local* condition were significantly higher than errors in the *One Local, All Global, or One Global* conditions. On the other hand, errors in the *All Local* condition were significantly lower than errors in the *All* condition.

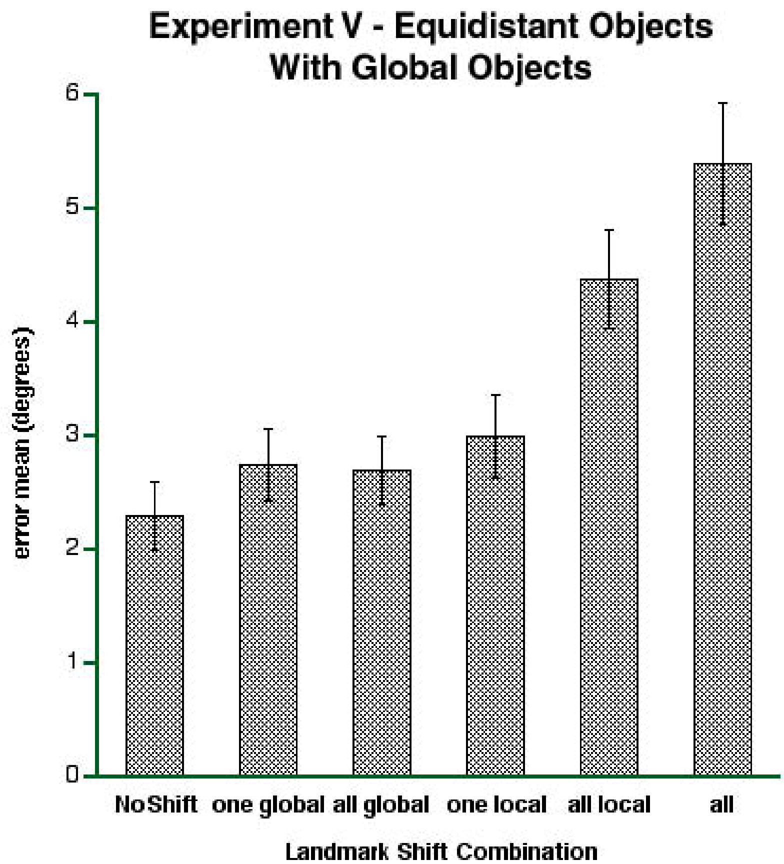

Figure 11 The above graph shows the overall error means for each Landmark Shift Combination in Experiment V .

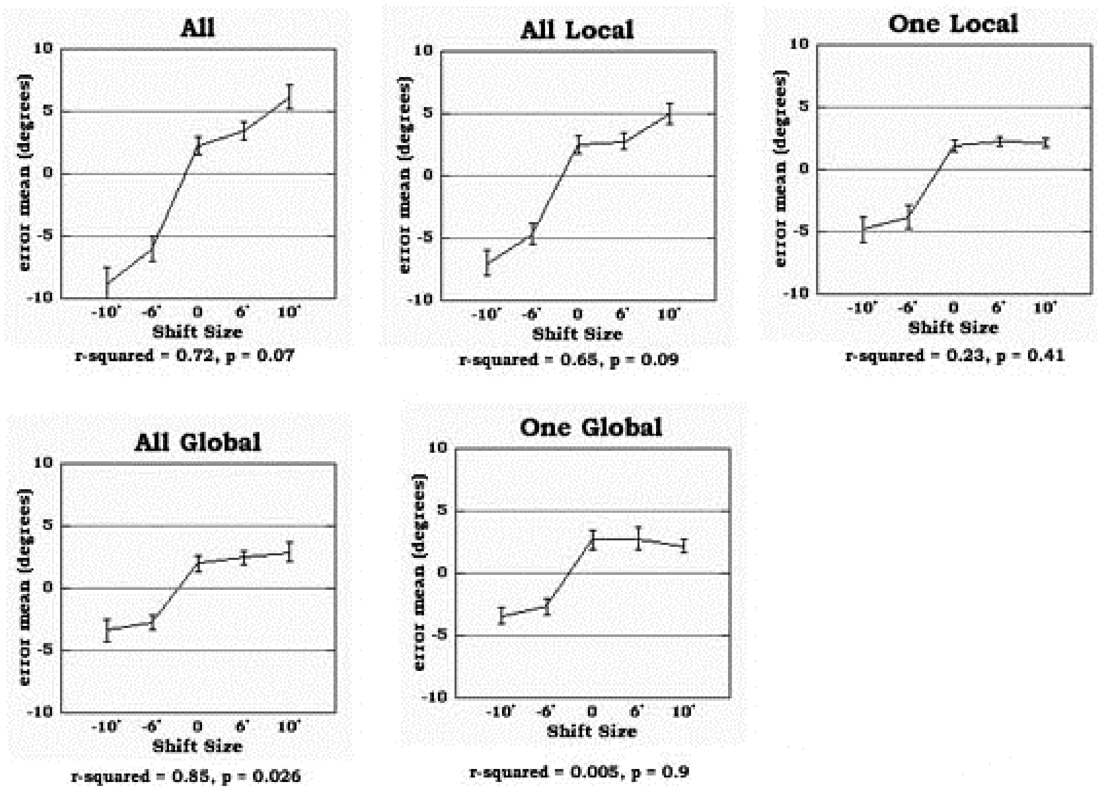

Figure 12 The above graphs show the error means in each Landmark Combination condition by each Shift Size in Experiment V. The correlation coefficients are included for each Landmark Combination/Shift Size cross.

6.4 DISCUSSION

Overall the results of this experiment followed the same pattern as did the results of Experiment I. The *All* and *All Local* conditions again yielded the highest errors, whereas the errors in the *One Local* condition were now not much higher than the errors in the *All Global* and the *One Global* conditions. Therefore, the presence of a global context was once again influential enough to reduce the effect of a highly weighted close landmark (Object 2) on the subjects' directional decisions. As in Experiment I the availability of a background took some of the subjects' attention away from the local objects and allowed subjects to make a more accurate assessment of whether or not one of the local objects had shifted away from its proper position. Therefore, a *One Local* shift was easier to spot.

In fact, it seems that the new configuration of the two closest local objects may have caused the global objects to be more influential than they had been in Experiment I. Although on average the errors in the *All Global* condition were not higher than the errors in the *One Local* condition, the errors in the *All Global* condition were significantly correlated with the shift sizes, whereas the *One Local* errors were not. Therefore, the simpler configuration of the two local landmarks may have made it easier to notice a breakdown of that configuration and impelled the subjects to rely more on the global objects.

Of course, according to the instructions it wasn't out of the realm of possibility that the one local object remained in its proper position while the rest of the objects shifted in the opposite direction. Even though that was not a type of shift incorporated into the experimental design the subjects had no way of knowing it. So the question still stands whether subjects chose to rely on the majority of objects because in that particular situation a visual majority was weighted more heavily than one object or because subjects really could tell that one object physically moved to a different location. This question will have to be answered in a future study.

7 Experiment VI - Colorswitch Control

7.1 INTRODUCTION

The following experiment was added to insure that the outcomes of the previous experiments could really be attributed to positional factors and not to some idiosyncratic features owned by different objects, such as colors. This experiment recreated Experiment III, with the exception that the colors of Objects 1 and 2 were swapped. This was a prudent control in light of the fact that subjects did on occasion refer to the objects by their colors. Therefore it was important to show that colors did not dictate the differential treatment of objects throughout the course of this study.

7.2 METHOD

Subjects

Two adult subjects participated in this experiment for monetary compensation.

Virtual Environment

The environment in this experiment was the same as the environment in Experiment III (Same Distance Landmarks) with the exception that the colors (and textures) of Objects 1 and 2 were swapped.

Procedures

The procedures were the same in this experiment as those in Experiments I, II, IV, and V.

Design

Two factors were crossed in this experiment: Shift Sizes and Landmark Shift Combinations. The shift sizes in this experiment were : 10°, 6°, 0°, -6°, -10°.

Seven different landmark combinations were used: *All, Object 1, Object 2, Objects 1&3,* and *Objects 2&3.* In the *All* condition all the objects were shifted together. The names of the other conditions state exactly which object or objects were displaced. Two conditions that were used in Experiment III were excluded from this experiment. These were the *Object 3* and *Objects 1&2* conditions. The *Object 3* condition was excluded because it was already established that this object did not influence performance when displaced on its own. It was also excluded because the main goal of the experiment was to see if objects 1 and 2 would be used differently than the way they were used in Experiment III. The *Objects 1&2* condition was excluded because it was considered to be effectively no different from the *All* condition for the purposes of this experiment.

With 5 shift conditions and 5 landmark conditions there were 25 possible conditions in all. Each shift/landmark cross occurred in 4 trials. Along with the training trials that occurred every six trials there was a total of 121 trials.

7.3 RESULTS

A two-factor, within subjects ANOVA yielded significant main effects of the Shift, $F(4,4) = 12.5$, $p = 0.016$, and the Landmark, $F(4,4) = 11.3$, $p = 0.02$, conditions (see Figure 13). The largest errors were made during the *All* condition (mean error =

6.2°), followed by the *Object 2* (mean error = 3.8°), *Objects 2&3* (mean error = 3.7°), *Objects 1&3* (mean error = 3.0°), and *Object 1* (mean error = 2.5°) conditions.

Once again, a significant interaction was found between the Shift and Landmark conditions, F (16, 16) = 7.0, p = 0.000 (see Figure 14). The shift sizes were positively correlated with the error means in the *All* (r^2 = 0.98, p = 0.001, slope = 5.0), *Object 2* (r^2 = 0.90, p = 0.01, slope = 3.18), *Objects 2&3* (r^2 = 0.83, p = 0.03, slope = 2.73), and *Objects 1&3* (r^2 = 0.90, p = 0.014, slope = 1.9) conditions. The correlation between the shift sizes and the *Object 1* condition was not significant (r^2 = 0.51, p = 0.18, slope = 1.3).

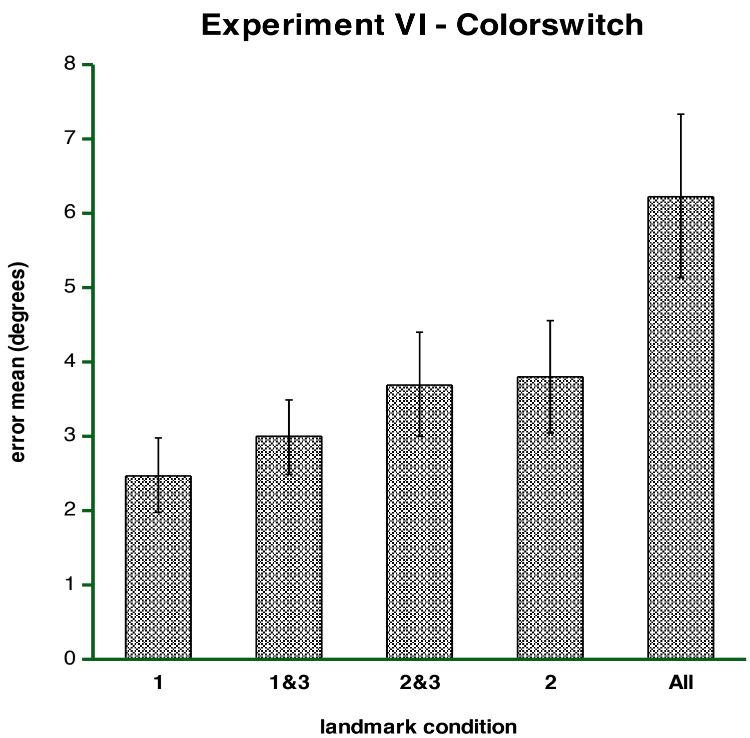

Figure 13 The above graph shows the overall error means for each Landmark Shift Combination in Experiment VI - Colorswitch.

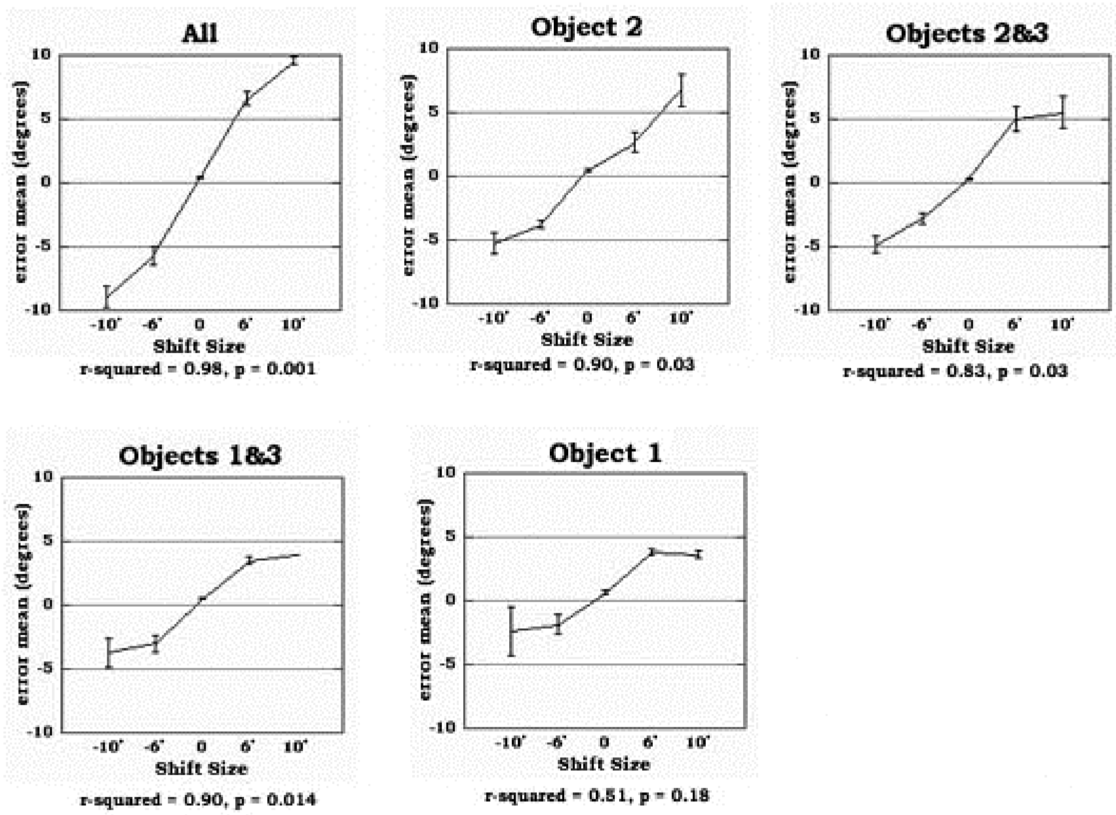

Figure 14 The above graphs show the error means in each Landmark Combination condition by each Shift Size in Experiment VI. The correlation coefficients are included for each Landmark Combination/Shift Size cross.

7.4 DISCUSSION

In this final experiment the colors of Objects 1 and 2 were swapped. This was done to make sure that Object 2 was not treated preferentially because of a unique feature such as its color. The results of this experiment showed that it was not the object's color, but it's position that influenced performance. As in Experiment III, Object 2 exerted a greater effect on directional decisions than Object 1.

It should be noted that the shapes of the objects were also unique. That is, the dimensions of the blocks varied slightly. For instance, Object 2 was a little shorter than Object 1. In light of the results of this experiment it seems unlikely, however, that the small differences in the objects' shapes or sizes caused their differential treatment.

8 General Discussion

In this study I set out to answer two questions: 1) If potential visual landmarks are known to be unreliable, can people ignore them and instead use only path integration to find a specific location? and 2) If they are compelled to use visual landmarks, how are those landmarks chosen? The preceding experiments showed that people were not able to completely ignore visual landmarks despite instructions to do so. When presented with an array of local and global objects people relied disproportionately on the object that stood closest to the target. Furthermore, they showed a preference for using the closest object that also fell within the path of sight as people made the necessary turn and translation to approach the target's location. The results of this study illuminated some of the principal strategies that are used by the human spatial learning system. These strategies will be discussed in the following sections.

8.1 USING DIFFERENT STRATEGIES

Apart from revealing people's "landmark" choices in a multiple-object environment this study also demonstrated people's flexibility in differentially weighting objects according to the configuration of the object array that was available. When one local object was located closer to the target than the rest they placed their bets on that one object to help them find the correct location. When two objects were equidistant from the target they adjusted by using both. However, they still assigned a heavier weight to the object that supplied them with more self-locomotion information as they moved towards the target. This object stood inside a space delineated by the direction they faced when they entered the arena and the direction in which they needed to turn to face the target pole. When large, distant objects could be seen in the background people accorded less influence to a single local object and used the background information to gauge whether or not local objects could be trusted in a particular trial.

The flexibility demonstrated by participants in these experiments is consistent with findings from a wide range of animal navigation studies. March, Chamizo, and Mackintosh (1992), for example, showed that the use of local and global information by rats could be manipulated by varying the reliability of the two types of cues. The authors trained two groups of rats to find food in a radial maze. One group could use both intra-maze and extra-maze visual cues to correctly locate the food source. The other group learned in an environment where only the extra-maze cues correctly defined the location of the food source. They found that upon being tested with only extra-maze cues the rats who initially had access to both performed significantly worse than the rats who learned with only extra-maze cues. Therefore when intra-maze cues were witnessed to be reliable they "overshadowed" the learning of features associated with extra-maze cues. In agreement with the current study the animals of March et al.'s study showed a preference for the closer visual information to guide their search when both near and far objects could be used.

Another example of flexible spatial learning behavior was described by Save, Poucet, & Thinus-Blanc, (1998). In a water maze task rats searched for a submerged platform that was located in a corner of a rectangular pool filled with an opaque liquid. To find the platform the animals initially appeared to rely on a colored cue card hanging on one of the walls of the pool. When the cue card was removed the rats immediately adjusted by using the shape of the pool. They only searched in either the correct corner or the corner diagonally opposite to the correct one. Furthermore, after several trials the rats showed a preference for only the correct corner, suggesting that they made another strategy shift that involved using some other static background cues that the experimenters had not controlled.

Single-cell recordings of hippocampal cells in the rat have confirmed that multiple categories of spatial information are being processed during exploratory activity (e.g., O'Keefe & Speakman, 1987). Recordings taken from complex spike cells, or 'place' cells, while the animal searched for food in a 4 arm maze showed that some cells responded to controlled local cues, some responded to uncontrolled global cues and some showed a firing pattern that was influenced by the interaction of the local and global information. The activity of these hippocampal 'place' cells presumably contributes to an animal's ability to use alternate information when the preferred cues are removed or displaced.

Some studies have shown that within groups of subjects, individuals can vary in their choices of orienting strategies when they are offered the same collection of visual cues. Huntingford and Wright (1989), for instance, conducted an experiment to see how spined sticklebacks used local and global landmarks. The fish were trained to discriminate between two different feeding patches in a tank. The feeding patches were labeled with distinctive local landmarks. Eventually, one of the feeding patches was equipped with a simulated predator which scared the fish whenever they entered that patch. Thus it was turned into a "dangerous" area. During the test phase the experimenters switched the local landmarks to see if the fish would mistakenly enter the dangerous food patch. They found that two thirds of the fish swam to the safe area despite the presence of the "dangerous" local landmark, thereby showing that they were using global landmarks outside the tank. One third followed the local signs and made the anticipated mistake. Therefore individuals in the same population of fish chose to use different strategies to orient themselves in the tank.

Correlates of such behaviors can also be found in human navigation research. In the introductory section of this paper I had mentioned a VR study conducted by Steck & Mallot (1997) which investigated the role of global and local landmarks in human navigation. In this study participants learned the layout of a virtual environment in which distinctive local and global objects coincided with various important road junctions. When local and global cues were manipulated so that they offered conflicting information the subjects showed varying personal preferences for landmark types. Some used only local objects, some used only global objects, and some alternated using one or the other. Interestingly, the latter group of subjects consistently chose the same pairings of a type of cue (were it local or global) with a particular junction. Therefore their choices were not completely arbitrary. Clearly, for these subjects, some distinctive property of each object made it the preferred landmark at that location in the environment.

The message that this and many other navigation studies conveys is that a mobile organism optimally exploits its environment by having the flexibility to utilize a wide range of cues. As suggested by Save et al.: "...spatial processing is not strictly confined to a precise category of environmental information but, rather, should be envisioned as an 'opportunistic system', using the most appropriate cues for spatial information encoding and efficient behavior" (Save et al., 1998, p. 121). I would like to caution that this type of description is correct to a point. It is true that humans and animals can easily adjust to using different types of spatial information. However, the current study demonstrates that the flexibility of the spatial processing system is limited. That is, there appear to be certain cues in the environment so important that their influence on navigational behavior can withstand the traveler's conscious effort to disregard them. Some of these cues and strategies will be described below.

8.2 SELECTING THE CLOSEST LANDMARK

In experiments I, and VI of this study both local and global objects were available to be used as landmarks. However, subjects relied heavily on one local object that was closest to the target pole. When the closest and second closest objects were placed at equal distances from the target pole the object formerly known as second closest exerted more influence than it did before (though still not as much as the original closest object). The inherent bias towards using the closest object as a landmark demonstrated in this study was observed in animal studies in the past. Of course, in the animal experiments the subjects did not expect the potential landmarks to randomly change positions from trial to trial. The fact that the humans in my study did expect objects to move about, thereby knowing that they would make unreliable landmarks, attests to the strength of this bias.

Cheng, Spetch and colleagues have conducted numerous studies looking at how birds and rodents use landmarks when searching for a goal (e.g., Spetch, Cheng, & Mondloch, 1992; Spetch et al., 1996; Spetch et al., 1997). In many of their experiments they have found that nearest landmarks are favored over others in these situations. Cheng (1989), for instance, trained pigeons to look for a small food well inside a box. On each side of the food well stood a rectangular wooden block. The blocks were placed so that one was closer to the food well than the other. When these blocks were shifted the pigeons adjusted their search in correspondence with the nearer landmark rather than the farther landmark. In another experiment (Cheng, 1986) rats showed the same tendency to attend to landmarks near a goal and disregard landmarks farther from the goal.

Reflecting on the findings of such experiments, Cheng (1989) remarked that there must be a great amount of flexibility in the way that landmark weights are assigned. After all, an animal can't afford to become disoriented in its environment when a particular landmark changes in appearance or disappears altogether. However, Cheng stated that assigning the largest weight to the nearest landmark can be safely regarded as a dominant principle.

An addendum to the 'closest landmark' principle is that absolute distance between the landmark and the goal does not seem to affect its weight. What is important is that it is closer to the goal than any other landmark. Spetch (1995) trained humans and pigeons to find a goal location that was defined by alternating pairs of landmark arrays. One pair of arrays contained a landmark that was the closest in one display but was "overshadowed" by an even closer landmark in another display. The other pair contained one landmark that was consistently the closest. This landmark and the overshadowed landmark were an equal absolute distance from the goal. However, when tested the participants did not respond to changes made to the overshadowed landmark as much as they responded to changes made to the consistently closest landmark. Therefore, *relative* proximity turned out to be the important factor in the task.

Another interesting question that was investigated in Spetch's landmark studies is whether or not the use of landmarks in spatial search transfers between a real environment and 2-D images of the environment. Lechelt and Spetch (1997) conducted a search experiment with pigeons in both a touch-screen task and an open field task. In this paradigm the pigeons learned to either peck at a feeder shown on a color monitor or search for it in an open field. Three distinctive objects stood at various distances from the goal. In situations when one landmark was either removed or shifted, search accuracy was disrupted most when the manipulation was performed on the closest landmark or the second closest landmark (which was located on the opposite side of the feeder). In the open field the closest landmark exerted much more control over the pigeons' behavior than the second closest landmark, but in the touch-screen task this difference did not exist. In both conditions performance was not affected when the farthest landmark was manipulated any more than when all landmarks were left in their original positions.

Lechelt and Spetch reasoned that the inconsistency between the open-field and 2D screen in terms of the treatment of the closer landmarks was probably due to the fact that on the screen the distances between the landmarks were "miniscule and variable" whereas in the open field they were much larger and must have appeared more stable to the birds. They also attributed the difference to the processing of 2-D relationships in the screen task and the processing of 3-D relationships in the open field task. I would like to carry this explanation further by adding that the pigeons were able to move about when learning the 3-D relationships, whereas their experience with the 2-D image was static. The crucial difference between the touch-screen condition in Lechelt and Spetch's experiment and the others I have described (including my own) may be that the touch-screen did not allow the subjects to experience their own motion relative to the potential landmarks as they approached the goal. Undoubtedly the landmark nearest to the goal is most useful because as the traveler moves towards the goal his distance and direction relative to the close landmark change faster than they do relative to farther landmarks. Therefore the closer landmark offers more precise information about the position of the goal. However, there may be another property of a landmark that is just as, if not more, important than its proximity to the goal and that is its position relative to the path that the traveler must take to reach the goal. That is, it in order for an object to be considered a useful landmark it seems to help that the object is easily visible as the traveler closes the distance between himself and the goal.

8.3 USING A LANDMARK THAT IS ENCOUNTERED *EN ROUTE* TO THE GOAL

This study showed that subjects were most influenced by an object (Object 2) that was closest to the target pole. However, it was also an object that remained prominently positioned within the subjects' view as they adjusted their heading to face the target pole and then walked towards it. Object 2 retained this role when the second closest object (Object 1) was moved to a location equidistant from the target pole. This is because Object 1 was on the other side of the target pole. Thus in order for the subjects to extract an equal amount of rotational information from Object 1 they would have had to keep turning past the target pole, perhaps an extra 20 or so degrees. If there was already an object available that offered enough information as the subject turned towards the goal then turning beyond the direction of the goal would have clearly been an inefficient, let alone unnecessary, strategy.

Viewed in this light the results of the current study agree with an interesting theory put forth by Collett and Collett (2000) regarding the role of path integration in landmark learning. In their paper Collett and Collett noted that ants did not seem to learn visual cues that were encountered when they were traveling away from their goal. For example, when ants were made to walk through a U shaped maze that ultimately led to their nest they only learned the cues that were found in the arm of the maze along which they moved *towards* the nest (Schatz, Chameron, Beugnon, & Collett, 1999). In another experiment it was found that landmark learning ceased when ants traveled beyond their nest (Collett, Collett, & Wehner,1999). The authors offered such observations as evidence that the path integration system may dictate which objects in the environment are treated as landmarks. The ant's path integration system, according to Collett and Collett, is a vector navigation system which constantly updates the vector representing the insect's distance and direction away from it's nest. Hence, they proposed, "The acquisition of visual cues may...only take place when the output of the vector navigation system is decreasing. Such a learning signal aids the rapid acquisition of routes because it allows reinforcement to act along the whole route on the very first trip home" (Collett & Collett, 2000,p. 255).

Of course, the experiments described in this study do not offer conclusive evidence that such a system plays a role in human spatial knowledge acquisition. The concurrence of one landmark with one rotation towards a goal does not provide adequate proof that people are noting a visual cue because it relates to an action that decreases the difference between their heading vector and the goal vector. A better test of this theory would involve a task such as the following: A series of distinctive objects should be arranged along a circuitous path that eventually leads to a goal. Subjects should train by walking several times along this path, which in certain segments forces them to move away from the goal. They should be told that their test will involve finding a straight path, or short-cut, to the goal in the absence of any visual cues. Subsequently the real test should involve some sort of recognition task revealing how well the subjects processed the various features of the objects. Presumably, if they use the 'vector navigation' strategy to select landmarks then their recognition of those landmarks will be better than their recognition of unattended objects.

8.4 AVERAGING

Another strategy demonstrated by the participants of this study was averaging the dictates of conflicting signals (for review, see Cheng & Spetch, 1998). In conditions in which the errors were correlated with shifts of particular objects these errors were, on average about half the magnitude of the actual object shifts. This implies that subjects were combining information from path integration and visual landmarks or from different landmarks and heeding a compromise between the two. Whether the veridical cues were non-displaced objects or internally-generated signals was not clear. It may have been a combination of both. However, both cases of averaging over the dictates of different visual cues and averaging over the dictates of different systems have been observed in previous navigation studies.

In the Spetch et al. studies (1996, 1997) pigeons showed a tendency to use a compromise between goal locations dictated by two landmarks that were shifted away from each other. The birds tended to search somewhere along a line that was parallel to the line segment that connected the two landmarks. This line along which they searched was also the same distance away from each landmark as was the goal in the original configuration. In the current study there were instances in which subjects may have been averaging between target locations that were implied by different landmarks. In Experiment V (Equidistant Objects With Global Objects), for instance, the errors in the *All Local* object combination condition were roughly half the sizes of the actual object shifts. Therefore, subjects may have been averaging the directions which the shifted local objects and the stationary global objects dictated. A similar response was found in the Experiment I (Walking) *All Local* condition.

Besides deciding between visual cues the subjects in the current study probably made compromises between the dictates of landmarks and path integration as well. As I mentioned in the General Introduction compromises have been observed between the path integration and piloting systems in both humans and other animals. Etienne, Teroni, Hurni, and Portenier (1990), for instance, trained hamsters to use a light to find their way back to their nest after feeding in the center of an otherwise dark arena. When the light source was rotated by 90° around the arena the hamsters were sometimes observed to search in a direction somewhere between the light source and the correct nest sight. Therefore, they were making a guess that incorporated both the visual signal and their path integration impulses.

In this study participants undoubtedly used the averaging strategy to strike a balance between the dictates of their internally-generated directional signals and the visual landmarks. For instance, the error means in the *All* landmark conditions of Experiments I (Walking) and II (Local Only) were about half the magnitudes of the shifts. Since all the objects moved together in those conditions participants could not have used other landmarks to counteract the dictates of the object array. Therefore they must have used their memories of performing the correct turn to resist fully complying with the object shifts. Interestingly, though, the error means in the *All* conditions of Experiments III (Equidistant Objects), V (Equidistant Objects With Global Objects), and VI (Colorswitch) were almost as large as magnitudes of the shifts. In these experiments the two local objects closest to the target pole were also equidistant from it. The size of the errors suggests that this new configuration added a greater weight to the visual landmark strategy. Presumably the combination of seeing correct intra-object spatial relations and being able to use the simple heuristic of walking in the middle of two close landmarks overruled the direction of the path integration system.

8.5 USING CONTEXTUAL CUES

In Experiments I (Joystick and Walking) and V (Equidistant Objects With Global Objects) the global objects clearly were not the primary landmarks used to specify the position of the target pole. However, they did affect participants' accuracy in a particular way. Evidently global objects were used to anchor the positions of the local objects, or, more often, one closest local object. This could be seen in the contrast between the results of Experiments I and II. In Experiment I the *One Local* (Object 2) shift did not influence subjects much more than the global objects. In that Experiment the entire array of local objects had a significantly greater effect on performance. However, when the global objects were removed in Experiment II the influence of the closest object (Object 2) greatly increased. Therefore, in Experiment I the array of global objects must have been providing a context by which subjects could judge the correct locations of the closer objects.

Furthermore, as it became clearer that Object 2 was weighted most heavily in the task, it also became apparent that the two peripheral local objects were sometimes playing an anchoring role as well. In Experiment IV, for instance, the local object farthest from the target pole, Object 3, was not used on its own to find the target pole. When it was shifted alone the errors were about as large as when no objects were shifted at all. However, the errors caused by Objects 1 or 2 when each was shifted in combination with Object 3 were higher than the errors found in the respective Object 1 or Object 2 single object conditions. Therefore Object 3 must have contributed to the directional decisions by tipping the "landmark" scale towards the close object with which it moved.

9 Conclusion

Through the use of a virtual environment this study was able to address one of the main questions in human navigation research: How do people select the objects that serve as visual landmarks during navigation? In the past this question has been a difficult one to answer because researchers were not able to manipulate the real-world environments in which humans travel. However, a virtual environment allowed us to systematically shift both large and small objects to see how such shifts would affect people's navigating abilities.

Displacing objects that were potential visual landmarks showed which objects were weighted more heavily when the objects were presented in various configurations. This study revealed the following guiding principles used in landmark selection and general navigation. 1) When both local and global objects were available subjects relied mainly on the local object configuration to find the target. However, the global objects influenced subjects by providing a visual background against which the positions of individual local objects could be gauged. 2) When only local objects were available subjects weighted most heavily the object closest to the target. 3) When two local objects were equidistant from the target the object that offered more information about the subject's relative motion as he or she approached the target was weighted more heavily as a visual landmark. 4) Finally, subjects relied on both path integration and visual landmarks to perform the place finding task as their final positions reflected a combined influence of displaced objects and their motion-based sense of the target position.

The above findings suggest that piloting and path integration are interdependent strategies. An object is initially selected to serve as a visual landmark because it offers more information about the traveler's movements relative to the target than any other available object. Since path integration is a strategy that relies on knowledge about self-motion, landmark selection seems to be motivated by path integration (Collett & Collett, 2000). Once the landmark has been selected and stored in memory the traveler can rely on a simple association between the landmark and a target location without having to perform the more laborious path integration operations to find the target. Thus, it would be interesting to further investigate how path integration dictates the selection of visual landmarks. Future studies using human subjects may find that objects acquire the special "landmark" status only if the traveler sees the object as his or her distance and direction from a target location decreases.

References

Cartwright, B. A. and Collett, T. S. (1982). How honeybees use landmarks to guide their return to a food source. *Nature, 295,* 560-4.

Cartwright, B. A. and Collett, T. S. (1983). Landmark learning in bees. *Journal of Comparative Physiology A, 151,* 521-43.

Chance, S. S., Gaunet, F., Beall, A. C., and Loomis, J. M. (1998). Locomotion mode affects the updating of objects encountered during travel: The contribution of vestibular and proprioceptive inputs to path integration. *Presence, 7* (2), 168-178.

Cheng, K. (1986). A purely geometric module in the rat's spatial representation. *Cognition, 23*, 149-78.

Cheng, K. (1989). The vector sum model of pigeon landmark use. *Journal of Experimental Psychology: Animal Behavior Processes, 15,* 366-75.

Cheng, K. and Spetch, M. L. (1998). Mechanisms of landmark use in mammals and birds. In *Spatial Representation in Animals,* pp. 1-18, (ed. S. Healey), Oxford: Oxford University Press.

Christou, C. G., and Bulthoff, H. H. (1999). The perception of spatial layout in a virtual world. Max-Planck Institut fur biologische Kybernetik, Technical Report No. 75.

Collett, T. S. (1996). Insect navigation *en route* to the goal: Multiple strategies for the use of landmarks. *The Journal of Experimental Biology, 199*, 227-235.

Collett, T. S., Cartwright, B. A. and Smith, B. A. (1986). Landmark learning and visuo-spatial memories in gerbils. *Journal of Comparative Physiology A, 158*, 835-51.

Collett, M. and Collett, T. S. (2000). How do insects use path integration for their navigation? *Biological Cybernetics, 83,* 245-259.

Collett, M., Collett, T. S., and Wehner, R. (1999). Calibration of vector navigation in desert ants. *Current Biology, 9*(18), 1031-1034.

Collett, T. S., and Kelber, A. (1988). The retrieval of visuo-spatial memories by honeybees. *Journal of Comparative Physiology A, 163,* 145-150.

Collett, T. S., and Land, M. F. (1975). Visual spatial memory in a hoverfly. *Journal of Comparative Psychology, 100*, 59-84.

Darken, R. P., Allard, T., and Achille, L. B. (1998). Spatial orientation and wayfinding in large-scale virtual spaces: An introduction. *Presence, 7* (2), 101-107.

Edelman, S. and Bulthoff, H. H. (1992). Orientation dependence in the recognition of familiar and novel views of three-dimensional objects. *Vision Research, 32*(12), 2385-2400.

Etienne, A. S. (1992). Navigation of a small mammal by dead reckoning and local cues. *Current Directions in Psychological Science, 1*(2), 48-52.

Etienne, A. S., Berlie, J., Georgakopoulos, J., and Maurer, R. (1998). The role of dead reckoning in navigation. In *Spatial Representation in Animals,* (ed. S. Healey), pp. 54-69, Oxford: Oxford University Press.

Etienne, A. S., Maurer, R., and Seguinot, V. (1996). Path integration in mammals and its interaction with visual landmarks. *The Journal of Experimental Biology, 199,* 201-209.

Etienne, A. S., Teroni, V., Hurni, C., and Portenier, V. (1990). The effect of a single light cue on homing behavior of the golden hamster. *Animal Behavior, 39,* 17-41.

Flach, J. M., and Holden, J. G. (1998). The reality of experience: Gibson's way. *Presence, 7* (1), 90-95.

Fujita, N., Klatzky, R. L., Loomis, J. M., & Golledge, R. G. (1993). The encoding-error model of pathway completion without vision. *Geographical Analysis, 25*(4), 295-314.

Fukusima, S. S., Loomis, J. M., and Da Silva, J. A. (1997). Visual perception of egocentric distance as assessed by triangulation. *Journal of Experimental Psychology: Human Perception and Performance, 23,* 86-100.

Gibson, J. J. (1979). The ecological approach to visual perception. Boston: Houghton Mifflin.

Huntingford, F. A. and Wright, P. J. (1989). How sticklebacks learn to avoid dangerous feeding patches. *Behavioral Processes, 19,* 181-9.

Kearns, M. J., Warren, W. H., Tarr, M. J., Duchon, A. P. (submitted). Path integration from optic flow and body senses in a homing task. *Perception.*

Koh, G., von Wiegand, T.E., Garnett, R. L., and Shinn-Cunningham, B. (1999). Use of virtual environments for acquiring configurational knowledge about specific real-world spaces: I. Preliminary Experiment. *Presence, 8* (6), 632-656.

Lechelt, D. P. and Spetch, M. L. (1997). Pigeon's use of landmarks for spatial search in a laboratory arena and in digitized images of the arena. *Learning and Motivation, 28,* 424-445.

Loomis, J. M., Klatsky, R. L., Golledge, R. G., Cicinelli, J. G., Pellegrino, J. W., and Fry, P. A. (1993). Nonvisual navigation by blind and sighted: Assessment of path integration ability. *Journal of Experimental Psychology: General, 122*(1), 73-91.

Mallot, H. A., and Gillner, S. (1998). View-based vs. place-based navigation: What is recognized in recognition-triggered responses? Max-Planck Institut fur biologische Kybernetik, Technical Report No. 64.

March, J., Chamizo, V. D., and Mackintosh, N. J. (1992). Reciprocal overshadowing between intra-maze and extra-maze cues. *Quarterly Journal of Experimental Psychology, 45B*, 49-63.

Margules, J., and Gallistel, C. R. (1988). Heading in the rat: Determination by environmental shape. *Animal Learning and Behavior, 16*, 404-410.

Mittelstaedt, H. and Mittelstaedt, M.L. (1982). Homing by path integration. In *Avian Navigation* (ed. F. Papi and H. G. Wallraff), pp. 290-297. Berlin: Springer Verlag.

O'Keefe, J. and Speakman, A. (1987). Single unit activity in the rat hippocampusnm during a spatial memory task. *Experimental Brain Research, 68,* 1-27.

Papi, F. (1992). *Animal Homing*. London: Chapman & Hall.

Ramos, J. M. J. (2000). Influence of the shape of the experimental room on spatial learning in rats. *Physiology & Behavior, 70*, 351-357.

Riecke, B. E., van Veen, H. A. H. C., and Bulthoff, H. H. (2000). Visual homing is possible without landmarks. Max-Planck Institut fur biologische Kybernetik, Technical Report No. 82.

Ruddle, R. A., Payne, S. J., and Jones, D. M. (1998). Navigating large-scale "desk-top" virtual buildings: Effects of orientation aids and familiarity. *Presence, 7* (2), 179-192.

Ruddle, R. A., Payne, S. J., and Jones, D. M. (1999). Navigating large-scale virtual environments: What differences occur between helmet-mounted and desk-top displays? *Presence, 8* (2), 157-168.

Save, E., Poucet, B., and Thinus-Blanc, C. (1998). Landmark use and the cognitive map in the rat. In *Spatial Representation in Animals*, pp. 119-133, (ed. S. Healey), Oxford: Oxford University Press.

Schatz, B., Chameron, S., Beugnon, G., and Collett, T. S. (1999). The use of path integration to guide route learning in ants. *Nature, 399,* 769-772.

Scholkopf, B., and Mallot, H. A. (1995). "View-based cognitive mapping and path planning" *Adaptive Behavior, 3,* 311-348.

Simons, D. J., and Wang, R. F. (1998). Perceiving real-world viewpoint changes. *Psychological Science, 9,* 315-320.

Spetch, M. L. (1995). Overshadowing in landmark learning: Touch-screen studies with pigeons and humans. *Journal of Experimental Psychology: Animal Behavior Processes, 21,* 166-81.

Spetch, M. L., Cheng, K., and MacDonald, S. E. (1996). Learning the configuration of a landmark array: I. Touch-screen studies with pigeons and humans. *Journal of Comparative Psychology, 110*, 55-68.

Spetch, M. L., Cheng, K., MacDonald, S. E., Linkenhoker, B. A., Kelly, D. M., and Doerkson, S. R. (1997). Use of landmark configuration in pigeons and humans: II. Generality across search tasks. *Journal of Comparative Psychology, 3* (1), 14-24.

Spetch, M. L., Cheng, K., and Mondloch, M. V. (1992). Landmark use by pigeons in a touch-screen spatial search task. *Animal Learning and Behavior, 20,* 281-92.

Steck, S. D. and Mallot, H. A. (1997). The role of local and global landmarks in virtual environment navigation. Max-Planck Institut fur biologische Kybernetik, Technical Report No. 63.

Tarr, M. J., and Pinker, S. (1989). Mental rotation and orientation-dependence in shape recognition. *Cognitive Psychology*, *21*(28), 233-282.

Teroni, E., Portenier, V., and Etienne, A. S. (1987). Spatial orientation of the golden hamster in conditions of conflicting location-based and route-based information. *Behavioral Ecology and Sociobiology, 20,* 389-97.

Tinbergen, N. (1972). *The animal in its world.* Harvard Press, Cambridge, MA.

Waller, D., Hunt, E., and Knapp, D. (1998). The transfer of spatial knowledge in virtual environment training. *Presence, 7* (2), 129-143.

Warren, W.H. (1998). Visually controlled locomotion: 40 years later. *Ecological Psychology, 10*(3-4), 177-219.

Zahorik, P., and Jenison, R. L. (1998). Presence as being-in-the-world. *Presence, 7* (1), 78-89.