ALLOCENTRIC VS. EGOCENTRIC SPATIAL MEMORY ENCODING: EVIDENCE OF A COGNITIVE SPATIAL MAP FROM VIRTUAL REALITY TESTING

ALLOCENTRIC VS. EGOCENTRIC SPATIAL MEMORY ENCODING: EVIDENCE OF A COGNITIVE SPATIAL MAP FROM VIRTUAL REALITY TESTING

By

CHRISTOPHE SEVIGNY, B.A.

A Thesis
Submitted to the School of Graduate Studies
in Partial Fulfillment of the Requirements
for the Degree
Master of Science

McMaster University

MASTER OF SCIENCE (2009)
(Psychology)

McMaster University
Hamilton, Ontario

TITLE: Allocentric vs. Egocentric Spatial Memory Encoding: Evidence for a Cognitive Spatial Map from Virtual Reality Testing

AUTHOR: Christophe Sévigny, B.A. (York University)

SUPERVISOR: Suzanna Becker

NUMBER OF PAGES: v, 43

ABSTRACT

Navigation is a very important area of spatial information research that presents researchers with a number of challenges. One of these challenges concerns the nature of spatial information encoding itself: is such encoding the result of a single mechanism system, a two-mechanism system or possibly a mixed system? One possible avenue of insight into this problem centers on the disorientation effect as described in Wang & Spelke (2000). A quick survey of basic findings, terminating with Waller & Hodgson (2006), indicates that there seem to be two systems at work. Moreover, the results obtained are based upon experiments carried out in actual reality. A virtual reality experiment was designed in an attempt to replicate the findings described in Waller & Hodgson (2006). The experiment is described in detail and its results are presented. These were found to be sufficiently reliable to justify pointing to a potentially rich field for future research, including such techniques as combining VR with fMRI to achieve more fine-grained results that cannot currently be obtained from the direct use of actual reality only. Underlying factors such as experimental control and data presentation are briefly described in the discussion section.

Key words: actual reality, allocentric system, egocentric system, geocentric system, hippocampus region, memory encoding mechanism, navigation, place cells, spatial memory encoding, true error, variable error, virtual reality

ACKNOWLEDGEMENTS

I would like to thank my committee members: Suzanna Becker, Patrick Bennett, Daniel Goldreich, and Hong-Jin Sun whose insightful comments and continual support made this work possible. I would also like to thank my family for their love and support during the research for and the writing of this thesis.

INTRODUCTION

General Overview

Environmental navigation is a pivotal element of survival for most species on earth. For animals, a reliable means of environmental navigation is critical since they must be able to return to the home or nest after foraging or hunting (Cheng, 1989). But how do animals accomplish this navigation without external referential systems such as maps? Presumably, they use an internalized system of navigation with possible access to external reference points that form part of the local ground topology, local points of reference, etc. This idea leads immediately to interesting questions concerning how these spatial memories are encoded and how they are stored and accessed. Early studies, according to Tolman (1948), fell into two general schools of thought. On the one hand, the strict behaviorists maintained that rat behaviour in navigating through, say, an unknown maze could be explained in terms of overlearned chains of stimulus-response bonds which became fixed through reinforcement, the other bonds being weakened and eliminated. Eventually, only the reinforced s-r bonds remained and these victorious chains of s-r bonds became the rat's solution to the maze problem. On the other hand, other behaviorists, whom Tolman labeled 'field theorists', felt that critical information was either being left out or ignored by the strict behaviorist approach. In a series of experiments summarized in his paper, Tolman showed that there was strong evidence for so-called 'latent learning', a term that, according to Tolman, was first used by Hugh Carlton Blodgett in 1929. Rats that had been run through the maze without reward for a certain number of days seemed unaware of their surroundings and made no progress in mastering the maze. However, their performance improved dramatically once rewards were introduced, showing that they had been storing information all along. Tolman did not deny that stimulus-response experience was happening, but he claimed that a cognitive spatial map of the environment, in this case the maze, also was being assembled in the rat's brain. As part of the support for this claim, he described one incident where rats were being put back at the start of the maze. However, two of them pushed their way up through the top of the maze cover and made their way directly to the food source. Tolman saw this not only as proof that a cognitive map had been constructed but also that it consisted of a wide strip as opposed to a narrow strip of the topology involved. Later research would show that a lot of this cognitive spatial information is stored in the hippocampus (O'Keefe & Speakman, 1987; Maguire et al., 1998; Bird & Burgess, 2008). The two generic approaches discussed in Tolman's paper eventually developed into an interesting research question: is spatial mapping a one-system or a two-system mechanism? The quest for an understanding of this basic question launched several decades of intensive neurological research into spatial memory dynamics. These were conceptualized and described as falling into one general scheme which itself yielded four generic approaches which turn out to be conceptual variations of one another although each subsystem places emphasis on different aspects (Allen & Haun, 2004). The general scheme says that spatial memory operates using a two-system mechanism of which one mechanism is called internal because it is based on internal adjustments and calculations made cognitively on the fly during navigation. It tends to be quite accurate but fades quickly once it is no longer active. The second, and complementary, mechanism is called external because it is largely dependent upon landscape points and their interconnections. These locations and reference points are memorized, resulting in the establishment of a fairly stable network of relations which yields a useful but approximate cognitive map of the environment. These two-system schemes are given different names in the literature: a) 'perception-action' versus a 'cognitive system' (Creem & Proffitt, 1998, 2001a, & 2001b); b) 'categorical processing' versus 'fine-grain information' (Huttenlocher et al., 1991; Allen, 1981; Newcombe & Huttenlocher, 2000; c) 'categorical adjustment modeling' versus 'coordinate coding' (Kosslyn, 1987, 1994; Kosslyn, Chabris, Marsole, & Koenig, 1992; d) 'taxon' versus 'locale systems' (O'Keefe & Nadel, 1978; Burgess, Jeffery, &O'Keefe, 1999; Burgess et al., 2002). This fourth category of subsystems is similar to the distinction between the 'stimulus-response' versus 'cognitive mapping' schemes. Just exactly how these two systems work together remains problematic and an object of continued research (Wolbers et al., 2008). The consensus for using a two-system scheme is almost unanimous but Allen and Haun (2004), argued that the two-system approach leaves gaps in representation capability when it comes to establishing verbal links necessary to enable us to describe the coordinated workings of these two systems. They suggested that it might be more plausible to set up a single system with internal transformation and unified representation of data modes. This would give the single-system much needed all-round homogeneity (Allen & Haun, 2004, pp. 57-59). On the other hand, the same authors point out that a 3-part scheme, similar to the system proposed by Piaget and Inhelder (1948/1967), consisting of a perception-action system, a categorical conceptual system, and a coordinate conceptual system, might provide a better account of current findings (Allen & Haun, 2004, pp. 59-60).

The research reported in this paper, assumes a two-system spatial representation of the sort just described. The emerging general consensus concerning this two-system scheme is well captured by the following quotation from Waller and Hodgson: ' Current theories of environmental cognition typically differentiate between an online, transient, and dynamic system of spatial representation and an offline and enduring system of memory representation' (Waller & Hodgson, 2006, p. 1).

Human Spatial Memory Encoding

In the wake of early results in animal spatial memory and animal spatial navigation, it was natural to produce spatial memory tests of human environmental navigation. Early studies of human spatial cognition were purely behavioural, whereas more recently, the advent of neuroimaging techniques has permitted investigation of the neural bases of human spatial abilities. Several early studies of human spatial memory emerged questioning the view of a geocentric human navigational system, and arguing in favor of a strictly viewpoint-dependent system (Diwadkar & McNamara, 1997). Moreover, evidence was also being accumulated that representations of different sized spatial layouts might also be orientation dependent (Sholl & Nolin, 1997; Roskos-Ewoldsen, et al, 1998). Finally, a solid challenge to the geocentric theory was formulated by Wang and Spelke (2000), who postulated that human spatial navigation is *dynamic* (not passive) and *egocentric*. They developed a very influential paradigm in which participants studied the locations of a set of objects and then pointed to those same items either with eyes open or with eyes closed and/or after disorientation. As a result of these experiments, Wang and Spelke developed the theory that navigation is based purely on an egocentric system which is consistently updated as humans are navigating their current environment, contradicting the then dominant geocentric view of spatial navigation. In the experiments conducted by Wang and Spelke, participants were asked to familiarize themselves with a novel environment (a lab room), a process which required that they memorize the locations of a variety of objects in the room. They were then asked to point to these objects. Over the course of the experiments, this pointing was done under differing conditions: a) with the participants having an unobstructed view of the objects in their vicinity and b) with the participants blindfolded. In their final experiment, Wang and Spelke had the participants sit on stools, put on a blindfold, and then rotate the stool for a minute (slowly enough that they would not become nauseous, but fast enough that they would become disoriented). Now came the test: participants were asked to point to each of the studied objects in the room from memory. Wang and Spelke observed a large increase in pointing error when the participants were disoriented as opposed to when the participants were not disoriented and attributed this to an impairment of the participants' transient, online, egocentric navigational system. These results formed the basis for the conclusion that heavy reliance was being put upon navigational information which was not geocentric in nature. Furthermore, the information appeared to be egocentric in nature. This view, i.e. that there is an egocentric encoding of spatial information, remained unchallenged for some 6 years or so. A new challenge came in the form of a compromise of the two commonly accepted systems: geocentric and egocentric. Waller and Hodgson (2006) argued that human navigation was not entirely dependent on a single egocentric or geocentric (now more commonly called allocentric) system, but rather on the combined results from both systems (Waller & Hodgson, 2006; Burgess, 2006). This very important paper is discussed further in the following section.

To summarize briefly what we have covered to date: research on mammals (rodents, primates, humans) has established that the hippocampus plays an important role in spatial navigation and encoding (O'Keefe & Speakman, 1987). Moreover, there is a strong body of evidence showing that human spatial navigation and spatial encoding also are dependent on the hippocampus (Maguire et al, 1998; Bird & Burgess, 2008). Furthermore, it is widely accepted that the hippocampus is used in the storage of long-term memories and that this allocentric, enduring memory is stored as a cognitive map in the hippocampus (Burgess, Becker, King & O'Keefe, 2001). Additionally, evidence is continuously accumulating that spatial memory consists of two interconnected systems, one of which is egocentric, dynamic and short-termed while the other is more stable and concerned with accumulating a cognitive map of the environment. As a participant begins navigating, the egocentric system maintains an accurate, continuously updated representation of the surrounding objects in order to maintain an accurate record of the current location, speed and direction within the environment. If the egocentric system cannot maintain such an accurate representation (e.g. after a very long delay and/or large viewpoint shift) then one switches to the more stable but less accurate allocentric (originally called geocentric) system. Finally, coordination of spatial information occurs mainly in the hippocampal region of the brain, especially the right hippocampus (Doeller & Burgess, 2007; Doeller et al., 2008; Hassabis et al., 2007a; Hassabis et al., 2007b).

Closer Examination of Dual Memory Systems at Work

So far, we have discussed evidence from purely behavioural experiments, which for the most part point toward a two-system representation in spatial memory. There is also a large body of evidence from electrophysiological recordings that supports this view. A series of experiments, measuring single unit activity in the rodent brain during a spatial task showed conclusively that there are cells in the hippocampal region which fire when the rodent is in a particular location, independent of the rat's heading direction or viewpoint (O'Keefe & Dostrovsky 1971; O'Keefe 1976; O'Keefe & Speakman, 1987; Hartley et al., 2003). This result permitted researchers to establish a definite correlation between the hippocampus and allocentric spatial memory encoding in the rodent brain. These hippocampus cells were named **place cells** since they became selectively activated when the rat found itself in a particular location. Additional studies indicated that this place representation was updated as the animal moved around using internal (e.g. vestibular, motor efference) and external cues (O'Keefe & Nadel 1978; O'Keefe & Burgess, 1996; Ekstrom et al., 2003). Part of the information being encoded includes direction, distance, pointing direction, local objects which become points of reference, speed of traveling. Many different cues have influence over a place cell's spatial selectivity, including the location and distance of local boundaries, the rat's heading direction, local objects which become points of reference, and the rat's speed of travel. This information falls into two general types: details, which depend upon sensory guidance and/or body movements and details, which depend upon local environmental cues. The first type has been labeled egocentric information because it is based upon the participant's position (rat or human) while the second has been labeled allocentric information because it is based on the environment which is outside the participant's body and/or position (King et al., 2002). As indicated above, there is convincing evidence that the construction of a cognitive map of the environment occurs primarily in the hippocampal region (O'Keefe and Nadel, 1978; King et al., 2002). Furthermore, this extra activity has also been linked with the build-up of grey matter as described in

(Maguire et al., 2000). There is also evidence that the hippocampal region allows data manipulations of various sorts such as flattening the data to render it more uniform in terms of modality. In other words, the hippocampal region seems to even out multimodal data to be consistent with a unimodal model, that is, a uniform point of view (King et al., 2002). It also seems responsible for effecting allocentric viewpoint-dependent representations in 3-D (O'Keefe and Recce, 1993). In addition to acting as a storage location for spatial information, there is strong evidence, though it is not well understood at the present, that the hippocampal system enables access to well-remembered positions by allowing searches through sets of retrieved information derived from these familiar positions (Burgess et al., 2001; Burgess et al., 2002). Similar evidence is provided by Bohbot (2007), who shows that the hippocampus is required to permit flexible navigation. Thus, the hippocampal region seems to act as a nexus region for gathering and manipulating spatial information. But it does not operate in isolation. Hassabis et al. (2007a, b) show that when individuals are asked to reconstruct new fictitious scenes, they activate a brain network which involves the hippocampus, parahippocampal gyrus, retrosplenial cortices, posterior parietal cortices, and ventromedial prefrontal cortex. These several systems bring in data from various senses. Most relevant to our study is the observation that patients with hippocampal damage are unable to construct multisensorial, fictitious scenes (Hassabis et al. 2007a, b). Hassabis et al. concluded that the hippocampus plays a critical role in the integration of past memories with current concepts, an idea that has been suggested by others (O'Keefe & Nadel, 1978; Burgess et al., 2001; Moscovitch et al., 2005; Eichenbaum, 2007; Hassabis et al, 2007a; Hassabis et al., 2007b). Finally, we note that Doeller et al. (2008, pp. 5918, 5919) point out that landmark-related learning obeys associative reinforcement, whereas boundary-related learning is incidental. They conclude that their findings strongly support the idea of parallel memory systems centered on the hippocampus and dorsal striatum. The hippocampal system deals with environmental layout (''place'' or ''locale'' learning), and striatal-dependent learning of responses is connected with individual stimuli (''response'' or ''taxon'' learning).

How the Two Systems Communicate.

An important set of questions concerns how the two levels of spatial representation interact: do we respond concurrently to both egocentric and allocentric levels of spatial representation, or do we choose to act based on just one or the other? That is, do they jointly contribute to our responses or do they compete? While this remains an unsolved question, empirical research is beginning to shed some insights. In general, it appears that the egocentric spatial updating system is the default system during active navigation, spatial orienting, and remembering the location of objects and landmarks around the observer. That is, when animals (including both rats and humans) are navigating through a space, be it a maze or a room, they tend to operate as much as possible by continuously updating their sense of their actual location based on a combination of sensory cues (visual, olfactory, auditory, etc.) and cues from their own motion. This sort of representation tends to be quite accurate provided that sensory information about the location of local landmarks, objects etc. is readily available. However, in the absence of sensory input (e.g. after prolonged obstruction of visual cues or in the dark) if disorientation or too great a change in location or heading direction occurs, there may be a switch from the calculation-intensive spatial updating system to the less computationally intensive but less accurate allocentric system. In the face of large changes from one's current position, for example when returning to a familiar environment from a novel perspective, reverting to an allocentric long-term memory representation, even if it must be re-aligned within the current context, may be the best or even the only option. At this point, we may begin to rely on boundary information recorded in our right hippocampal-based allocentric system. Using terminology from Doeller et al. (2008), we switch from landmark-based information to boundary-based information. Landmark-based information is derived from the participant's location through continuous calculations while boundary-based information may be based on a matching of distances to the nearest obstacle in all directions around the rat or the human participant. However, this potential explanation remains problematical because there are still unresolved issues in distinguishing landmarks from boundaries (Doeller et al. 2008).

Some of the earliest evidence of these dual spatial representations being used at different times came from Wang and Spelke (2000). They found that when people were oriented, even if blindfolded, they were highly accurate at pointing to the remembered locations of objects. When disoriented, they could still point reasonably well in spite of a drop in accuracy, but seemed to make consistent errors due to uncertainty in heading. The authors concluded that switching from responses based on an egocentric to an allocentric representation was triggered when participants became disoriented. Waller and Hodgson (2006) set out to replicate and extend the results obtained by Wang and Spelke (2000). They explored the conditions under which participants could point accurately to the remembered locations of objects. In a series of experiments, they controlled whether people were disoriented, and whether they were pointing to the actual remembered locations (egocentric pointing task) or pointing to remembered locations after an imagined viewpoint shift (a "judgment of relative distance" or JRD task). They found that disorientation was not required to produce an apparent "switch" from an egocentric to an allocentric strategy, but that a critical variable was the amount of rotation (real or imagined) that participants underwent. For rotations less than 135 degrees, people were consistent and highly accurate, whereas for larger rotations there was a sudden drop in accuracy, suggesting a switch in strategies or systems. Their explanation, while it differed from that of Wang and Spelke (2000), was quite straightforward: they argued that the egocentric system, based as it is on dynamic spatial updating, could operate with high accuracy as long as the calculations did not become too complicated. Their reasoning led them to posit that adjustments following a reorientation of 135 degrees or more would attain the critical threshold and trigger a switch from the egocentric to the allocentric strategy. Accordingly, they set up four sets of experiments and collected data. They obtained mixed results. In one set of experiments, participants did not perform as predicted and their navigational performance actually improved after 135 degrees instead of worsening. Waller and Hodgson (2006) discussed what this unexpected result might be indicating. According to them, system-switching is not an all-or-nothing event, but rather, the result of a continuous replacement process which gradually develops as the participant is less and less able to maintain an up-to-date adjustment of calculations because these are becoming either too numerous or too complicated. Thus, there might be a continuum of various mixed states of both systems being active simultaneously but in uneven proportions. The authors also pointed out that replacement rather than switching is also in part dependent upon what we could call familiarity factors. For instance, if participant A is more familiar with environment X or event Z, then adjustment calculations will be simpler and/or fewer for participant A than would be the case for, say, participant B who might not be nearly as familiar with the same situation. In such a case, participant A, to continue our example, would remain in egocentric mode and exhibit a lower increase in errors. Familiarity of task and/or location would seem to be an important bi-directional factor, bi-directional because it could be used as a signal that replacement, even complete switching was taking place or it could be used to indicate that no such replacement was taking place, indicating that this task was not particularly complicated for this participant.

Waller and Hodgson (2006) also discussed the concepts of transient online representation and an enduring memory representation. If a task is attempted and the participant's transient online system was not being affected by interference of some sort, then that stability would continue for a longer period of time than would be the case if interference had already begun to destabilize the participant. A second factor that might bear upon the issue is the possibility that there are differences between the self-reference system which codes participant-to-object spatial relations in body-centered coordinates and the object-to-object system which codes spatial relations among objects in environmental coordinates (Allen & Haun, 2004, p. 47; McNamara & Valiquette 2004). Such differences could account for the differing degrees of familiarity felt by participants and thus affect the switching or replacement of the online egocentric system by the allocentric system. Of course, in the long run that would also affect experiment results.

In order to further study the results reported in Waller and Hodgson (2006), we attempted to replicate their findings in a virtual environment. The goal was to show that experiments using a virtual environment obtain results that reliably mirror those obtained from experiments carried out in actuality and thus provide us with a reliable, additional tool for exploring complex events in neuropsychology.

METHOD

Participants were right-handed males with at least one year of experience playing first-person perspective video games, 14 of whom were McMaster University undergraduates recruited through the Experimetrix system, and an additional six who were non-McMaster participants recruited via email. Following a single calibration phase, during which participants pointed to visible objects, each participant completed four blocks of trials. Each block had two phases, a training phase during which participants learned the locations of a set of 6 objects, and a testing phase during which participants pointed to the remembered locations of objects. These phases are explained further below. Within each block, the two phases were repeated twice, so that there were a total of 12 pointing trials in each block (two testing phases with 6 pointing trials each). A test phase immediately followed the end of each training phase. During each test phase, memory for object locations was tested either from the same viewpoint or from a shifted viewpoint. Between blocks, the viewpoint shift was varied, so that in block 1 there was no viewpoint shift (imagined test perspective identical to reference perspective), while in blocks 2, 3 and 4 there was a non-zero viewpoint shift.

Virtual Reality Experiment

A publicly available version of the Unreal Software Engine was used to create the 3D environment and program the events in the experiment. Along with the freelydistributed graphics editor, the Unreal engine also comes with a proprietary scripting language that can be used to render objects in the environment at set times and to effect the flow of the experiment by triggering different events (e.g., making objects appear or disappear, cueing the participant with different tasks to perform, etc.). Using the graphics editor, we adapted a virtual reality environment which consisted of a grassy arena encircled by a stone wall. Beyond the grassy arena, were several different distal reference points: a sun, a dark mountain top and a snowy mountain top. Within the arena, we placed a set of six categorically related objects at different locations. In different phases of the experiment, either the participant could use the response buttons to freely navigate around the environment, or the participant was pinned to the centre of the arena and could press buttons to rotate his viewpoint, so that he was able to observe the objects from a variety of perspectives. We also used a second environment that we will refer to as the "pointing environment" in which the participant is placed within a cylindrical enclosure with stone walls and a grassy floor, again in an area which falls in a position central to the surrounding objects. His position puts him in the center of the objects. Within the pointing environment, visual cues appeared at the top of the screen, and a cylindrical compass with a red needle was displayed on the floor of the enclosure just ahead of the participant's viewpoint. The compass needle could be rotated by the participant using the arrow keys to indicate a pointing response.

To summarize: a single initial calibration phase was used to eliminate candidates who failed to qualify for the experiment. The experiment itself consisted of four blocks, each including a training phase and a testing phase. The viewpoint shift during the test phase was varied across blocks, as explained below.

Calibration Phase

The calibration phase consisted of a set of "eyes open" pointing trials. The "eyes open" pointing direction for each object was subsequently used as the true object direction, in order to calculate variable error in the testing phase.

Using the object the participant last faced as a reference object, the participant was given 9 seconds to rotate and look at another object, and then given 9 seconds to rotate back to face the reference object. For example, if the reference object was a balloon and the target object was a rook, the participant would start out facing the balloon, and would be cued to turn to face the rook and then turn back to the reference viewpoint facing the balloon. Then, the screen changed and the participant now found himself in the center of the pointing environment and was told that the pointer was currently pointing to the reference object, e.g. the balloon. He was then given instructions to point to a target object, e.g. the rook, using the left and right arrow keys on the keyboard to point the compass pointer in the direction where he remembered the target object, the rook, should be, and then to press the up arrow to enter his response. This constituted one pointing response for the target object, the rook. This double task, first in the grassy arena (eyes open condition) with the balloon as the reference object and the rook as the new target object and a second immediately following task in the pointing environment (eyes closed condition), was repeated until four pointing responses were recorded for each of the objects yielding 21 pointing responses per participant (reference object + ((target object x 4 repetitions) x 5 objects))

Training Phase

The training phase encompassed four stages: free navigation, "collect", "replace" and "central pivoting".

In the free navigation stage, the participant was initially randomly located within the grassy arena encircled by the stone wall as seen from a first person-perspective and asked if he was ready to commence. When the participant was ready, he navigated around the arena and studied the locations of six visible objects for 45 seconds. In order to mimic the Waller and Hodgson (2006) experiment, familiarity with the environment on the part of the participant was required. The navigation period occurred only before the first training phase. The participant was instructed to remember the locations of the six objects relative to distal landmarks and relative to each other. In other words, the participant was implicitly instructed to build an internal cognitive map of the objects.

In the collect stage, following the navigation stage, all of the objects were visible and the participant was cued to collect each object one at a time. This required that the participant navigate to the visible object and bump into it. On each trial, the participant was given a visible cue consisting of a small image of the object to be collected. The participant was told to navigate to the location of the cued object and bump into it, at which time the object would disappear. If the participant could not find the cued object after 15 seconds, he would be reoriented in the direction facing the object. The collect phase ended after the participant successfully completed the collect task for all six objects.

In the replace stage, following the collect stage, the participant was still randomly placed within the grassy arena, somewhere in the middle of the objects, none of which was visible for this set of trials. As for the collect trials, the participant was given a visible cue consisting of a small image of the object to be replaced. The participant was asked to replace the object, by going to the location at which he remembered the object being, at which point the object would reappear. If a participant failed to walk into an object after fifteen seconds, he was reoriented to its correct location. The object would appear in its correct location, blinking and bouncing up and down, so that the participant would thereby be reoriented to its correct location. The replace phase ended after the participant completed this task for all six objects from his initial, randomly assigned position in the grassy arena.

After completing the replace stage, the participant entered the central pivoting stage where he was placed in the center of the six objects while they remained visible. He was then asked to look around at the objects and study their locations for an additional 10 seconds. During this stage, the participant was able to rotate left and right but could not do anything else; he was unable to move forward or backward in the environment itself. This manipulation was designed to let the participant have one last chance to encode the locations of the objects relative to each other as well as to the distal landmarks before he began the pointing trials. After rotating about the center of the environment for 10 seconds, the participant was told to face a specific object (for example, the balloon) which became the participant's reference viewpoint (i.e., the last object he saw before he started the pointing trials).

Testing Phase

The participant then began the *pointing trials*. Each testing phase contained 12 pointing trials. In the first set of 6 pointing trials the participant did not undergo an imagined rotation, while in the second set of 6 pointing trials the participant was asked to point to each object from memory after undergoing an imagined rotation. In an imagined rotation, the participant is told to imagine he is facing an object in a direction other than his reference viewpoint, e.g. imagine facing the rook, where the reference viewpoint had been in the direction facing the balloon. There are two imagined rotation conditions: a low degree imagined rotation where the direction of the imagined object is 44-83 degrees away from the reference viewpoint, and a high degree imagined rotation where the imagined object direction is 107-131 degrees away from the original reference viewpoint. Just to clarify at this point: the participant himself is not told to imagine himself having rotated so many degrees. From his perspective, all he is asked to do is to imagine facing the object indicated, e.g. the peach. The rotation is the result of having the participant make mental adjustments in order to be facing the cued object. The category of this adjustment, low or high or zero imagined viewpoint rotation, is determined by the experimenter who controls the choice of the cued objects.

Each pointing trial begins on a blank grey screen upon which are shown two cues: a reference cue in the upper left of the screen and a target cue immediately to the right of the reference cue. For example, if the participant is to imagine facing the balloon (reference cue) and then point to the rook (target cue), then in the upper left of the screen there would be an image of an eye adjacent to an image of a balloon, then beside the balloon there would be an image of a hand with a finger pointing to an image of a rook to the right. After viewing this cue, the screen immediately changed and the participant found himself a) still in the center zone but in a slightly altered position in the zone, facing a compass pointer which he was instructed to use to make his response. (Note that in the zero degree condition this cued target object turns out to be the same as the reference cue, but that in the low or high condition, the target object will be a different object from the original reference object and the participant will have to make mental adjustments, simulating a certain degree of rotation in order to be facing this object). The participant completed this task once for each object with no imagined rotation and once for each object with an imagined rotation. This yielded 2 sets of 6 pointing trials per block.

In the latest version of the experiment, participants completed four blocks for a total of 48 pointing trials. The degree of rotation of the second set of 6 pointing trials changed in each block.

In the first block, the first set of 6 pointing trials always involved a zero degree rotation, meaning the last object they saw (the balloon, for example) also was the reference object. In the second set of 6 pointing trials, participants involved in version A of the program were cued with a low rotation (anywhere from 44-83 degrees) and participants involved with version B of the program were cued with a high degree rotation (anywhere from 107-131). The different versions of the program were created to counter-balance for any effects due to practice.

After having completed a block of the experiment, participants completed two more training phases which were exactly the same as before (a collect phase and a replace phase). The pointing trials were as before, but with the low and high groups switched around. Thus, by the end of the second block, all the participants had experienced two zero degree rotations (consisting of 12 pointing trials), one low degree rotation (consisting of 6 pointing trials) and one high degree rotation (consisting of 6 pointing trials).

Once the participants had completed the second block of the experiment, the experiment was repeated and they were asked to do the exact same set of tasks over again (free navigation, training phases, pointing trials, training phases, pointing trials). The calibration phase was not repeated before this repetition of the experiment.

True Error

True error was a signed value calculated based on the actual difference between the participant's pointing response and an object's true location. This error was only used as an exclusion criterion for the participants in the calibration phase and the zero degree pointing trials. Participant data was excluded if their average error on the calibration task was greater than two standard deviations above the mean (i.e. greater than 111.592 degrees) or if the average of the true error of any set of six pointing trials in the zero degree condition was greater than 35 degrees. These criteria were chosen because the participants had to have demonstrated a clear understanding of the task and an ability to demonstrate this by pointing to the locations of the objects and failing either of these criteria would show clearly that they had not. According to these exclusion criteria, seven participants' data were excluded.

Variable Error

During the calibration task described above, participants were asked to point to each object four times resulting in four pointing responses for each object. These pointing responses were averaged by object and this average was taken as the baseline pointing response for that object by that participant. Variable Error was defined in the same way as that used by Waller and Hodgson (2006). It was calculated by taking the standard deviation of the bearing differences between the pointing directions obtained in the calibration task and the pointing trials.

A one-way repeated measures ANOVA was used to test for differences between variable errors in the three separate rotation conditions (zero degree rotation, low degree rotation and high degree rotation). It was found that variable error differed significantly across the three rotation conditions ($F(2, 38) = 5.696, p = .007$). Post hoc comparisons showed (after Bonferroni adjustment for multiple comparisons) that variable error was significantly higher in the high rotation condition compared to the low degree rotation condition ($p = .036$) and compared to the zero degree rotation condition ($p = .010$). There was no significant difference between the zero degree rotation condition and the low degree rotation condition ($p = 1.000$). The average variable error plus or minus the standard error for each rotation condition being: zero degree rotation = 30.0116 ± 3.88519, low degree rotation = 35.5328 ± 6.02845 and high degree rotation = 64.8270 ± 11.53069. These results are summarized in Figure 1 below.

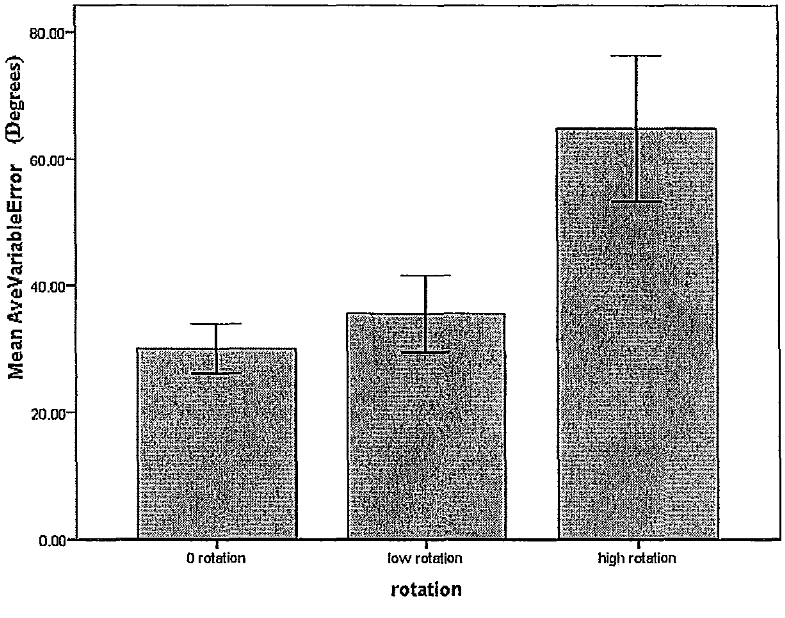

Figure 1 – Mean average variable error by rotation, rotation in the low rotation condition was 44-83 degrees and rotation in the high rotation condition was 107-131 degrees

Reaction Time

During the pointing trials, the time taken by participants to enter a response was also recorded. Reaction time is defined as the amount of time taken from the start of the pointing trial to the time when a pointing response is recorded. Figure 4 shows the raw reaction times for each participant with a curve of best fit. No outliers have been removed.

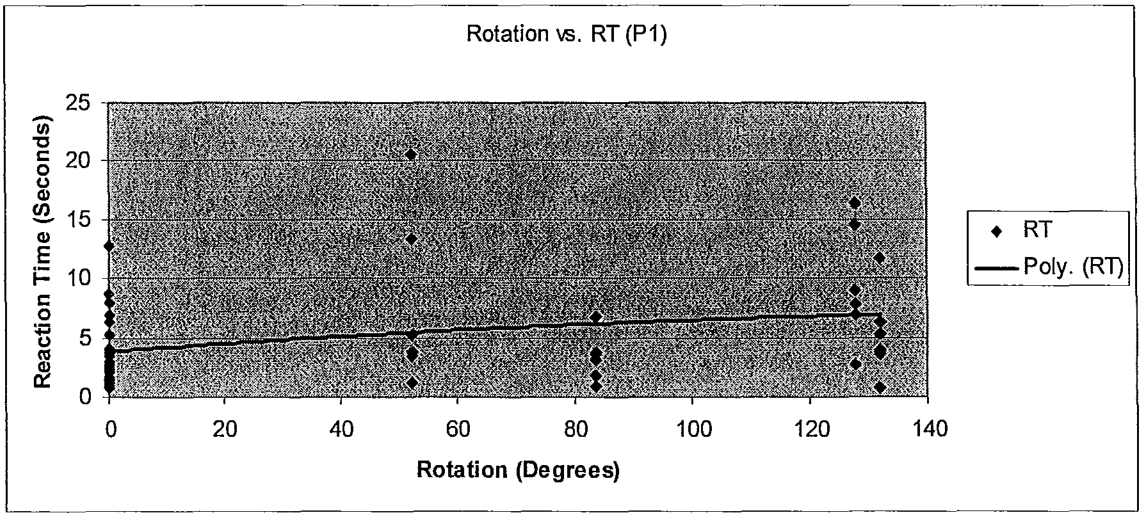

Figure 2 – Reaction times of participant 1 at each rotation condition

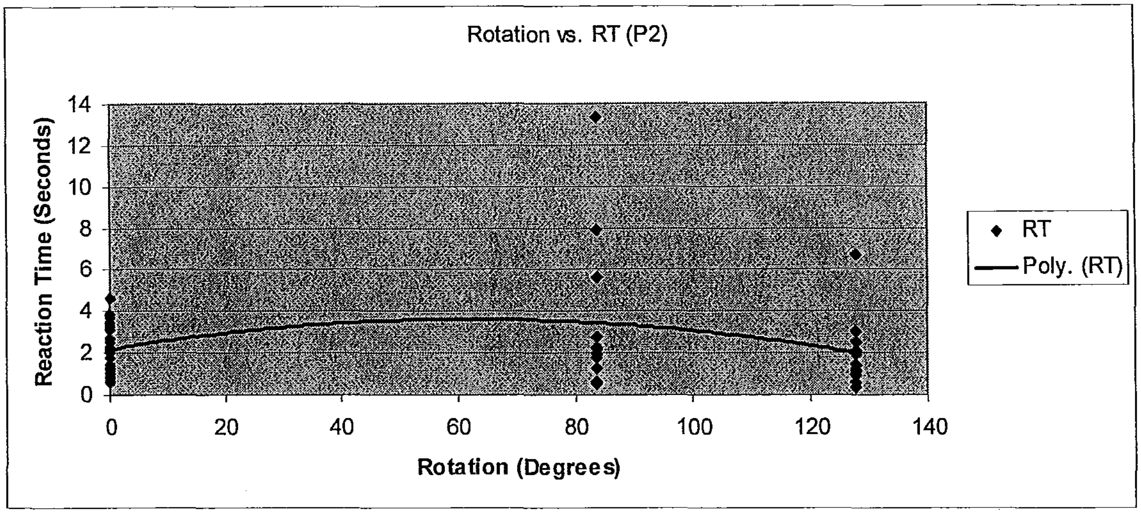

Figure 3 – Reaction times of participant 2 at each rotation condition

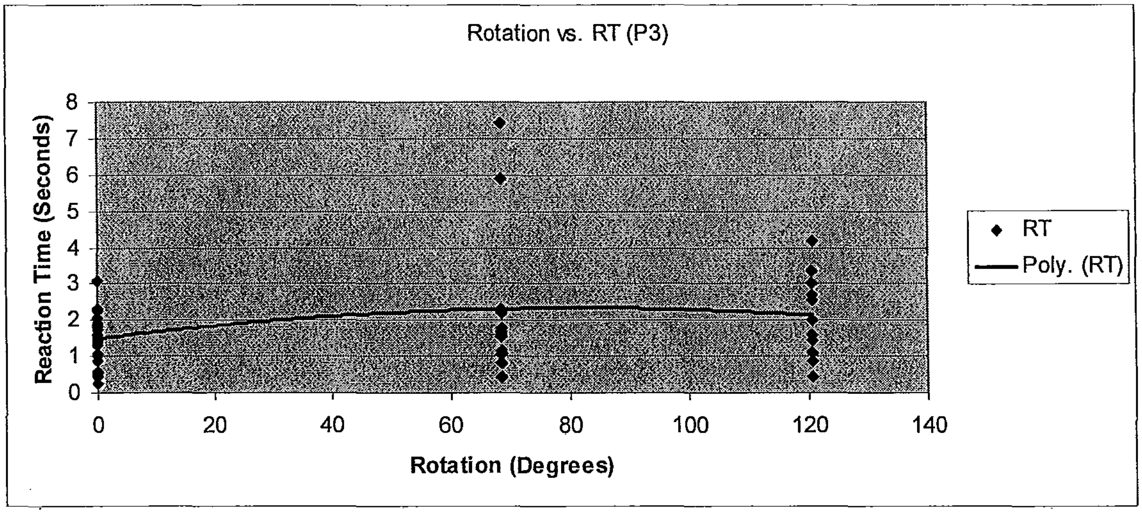

Figure 4 – Reaction times of participant 3 in each rotation condition

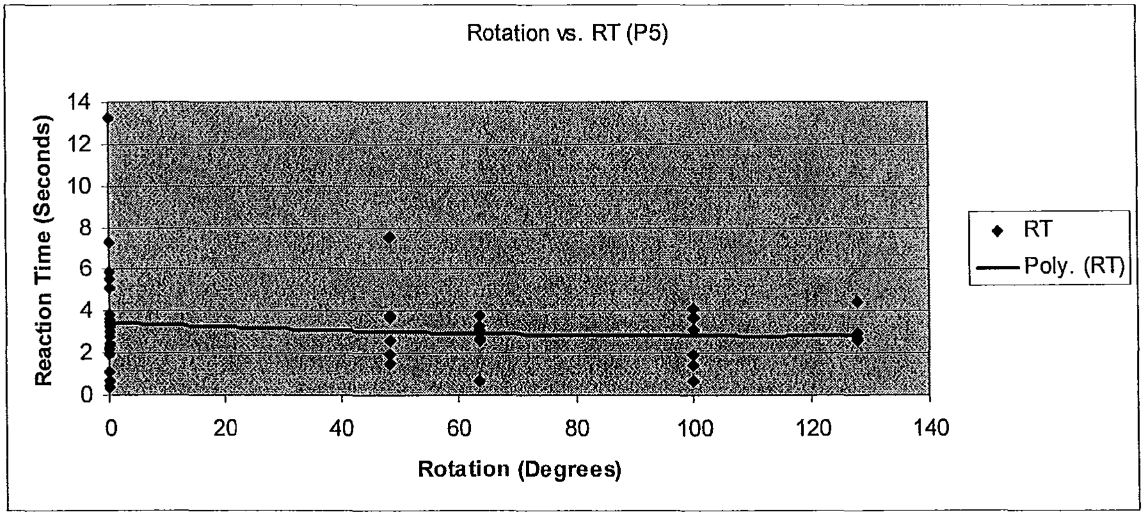

Figure 5 – Reaction times of participant 5 in each rotation condition

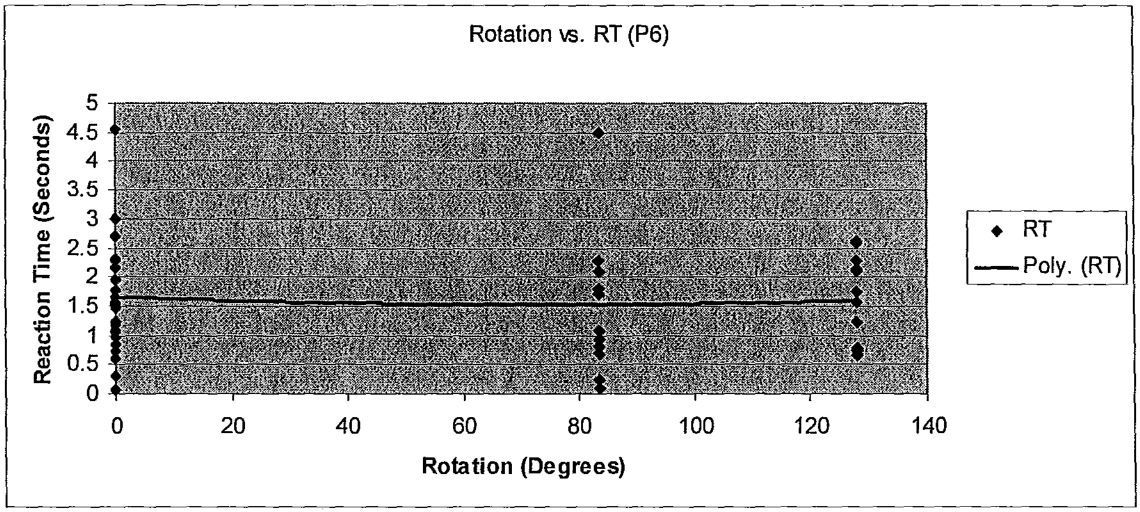

Figure 6 – Reaction times of participant 6 in each rotation condition

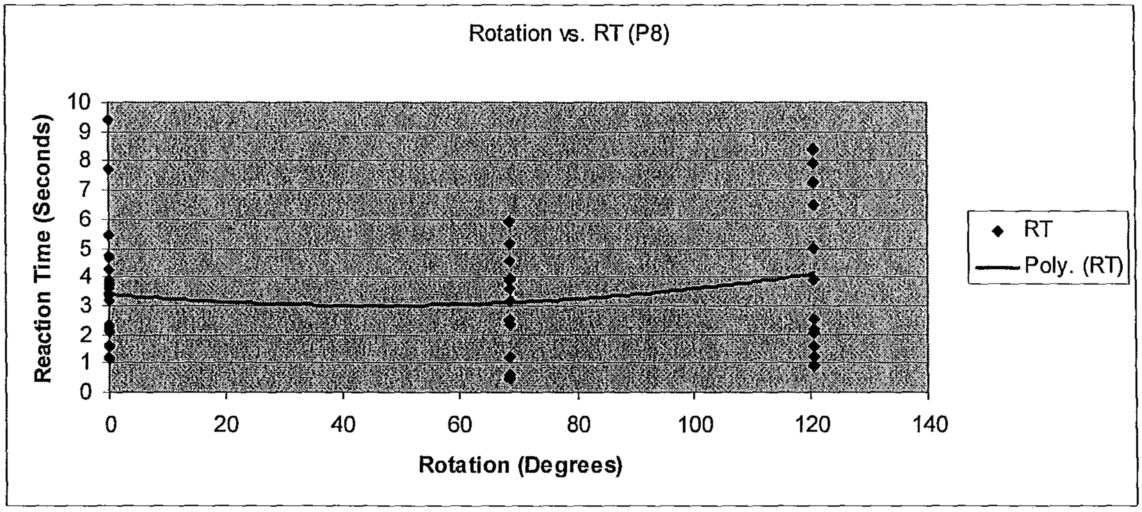

Figure 7 – Reaction times of participant 8 in each rotation condition

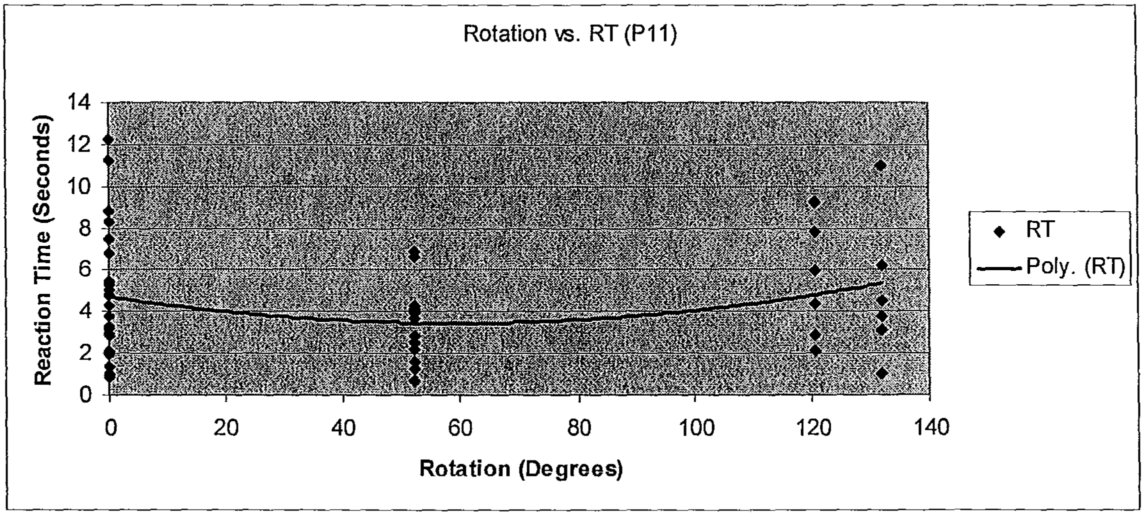

Figure 8 – Reaction times of participant 11 in each rotation condition

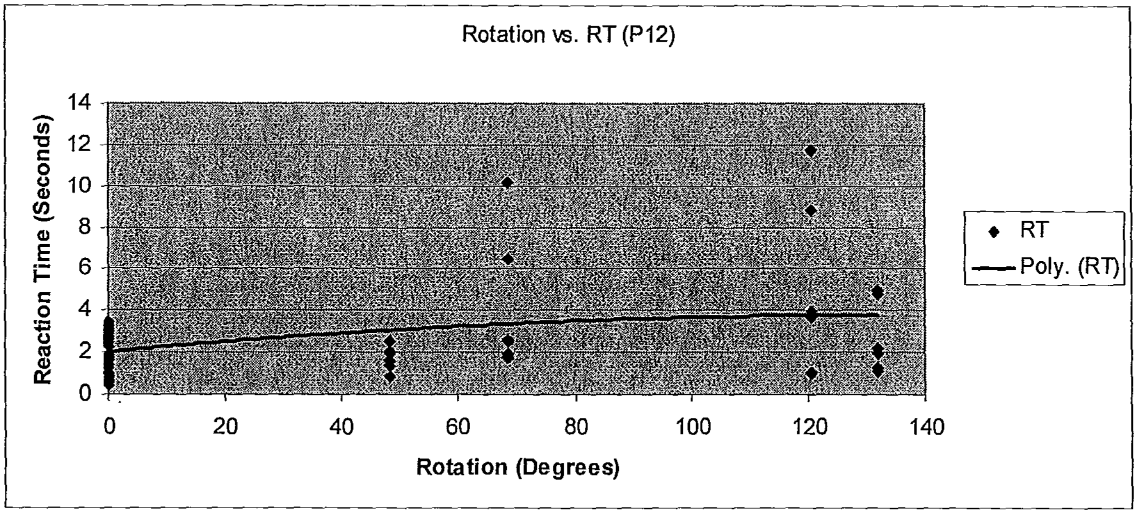

Figure 9 – Reaction times of participant 12 in each rotation condition

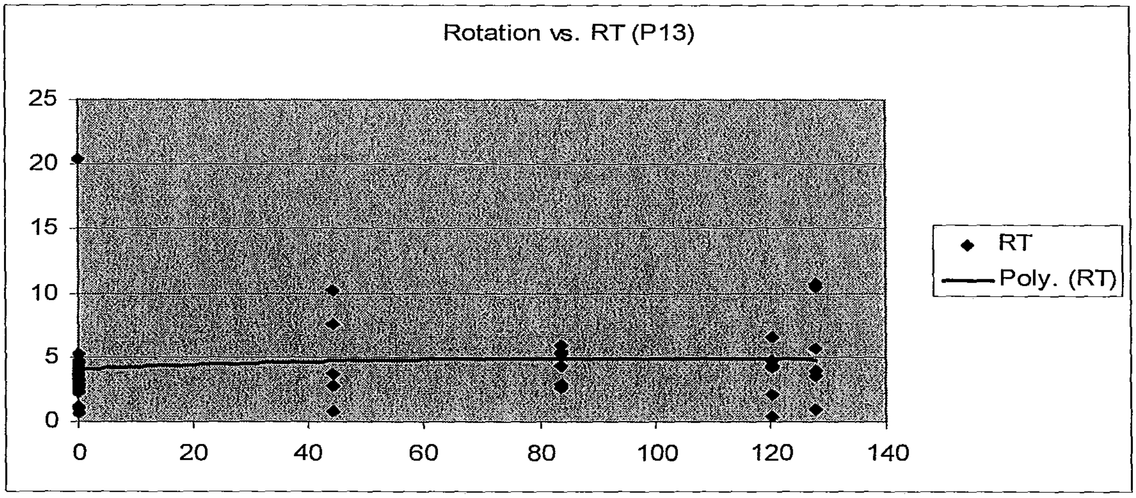

Figure 10 – Reaction times of participant 13 in each rotation condition

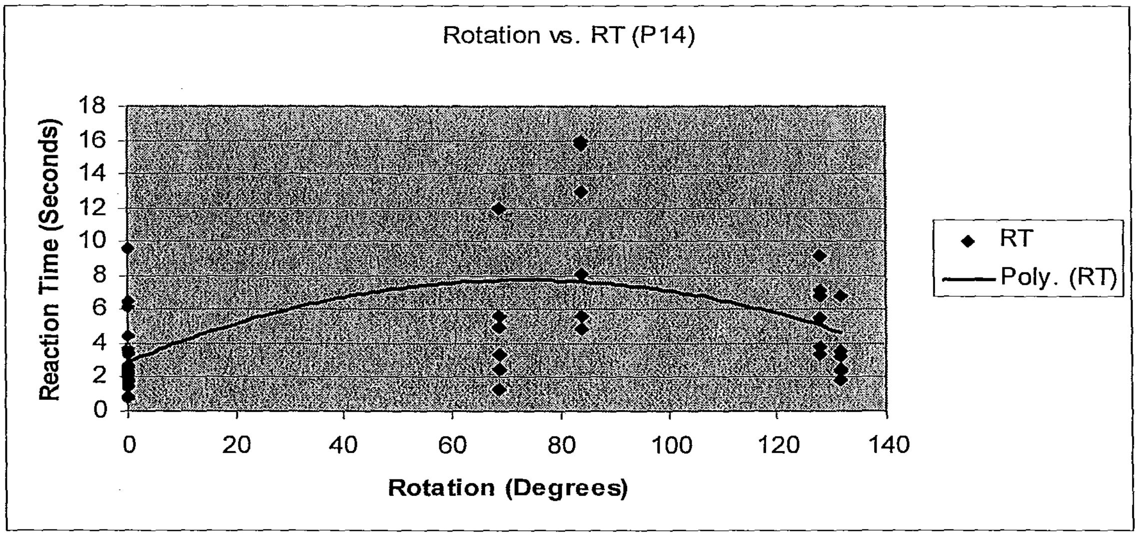

Figure 11 – Reaction times of participant 14 in each rotation condition

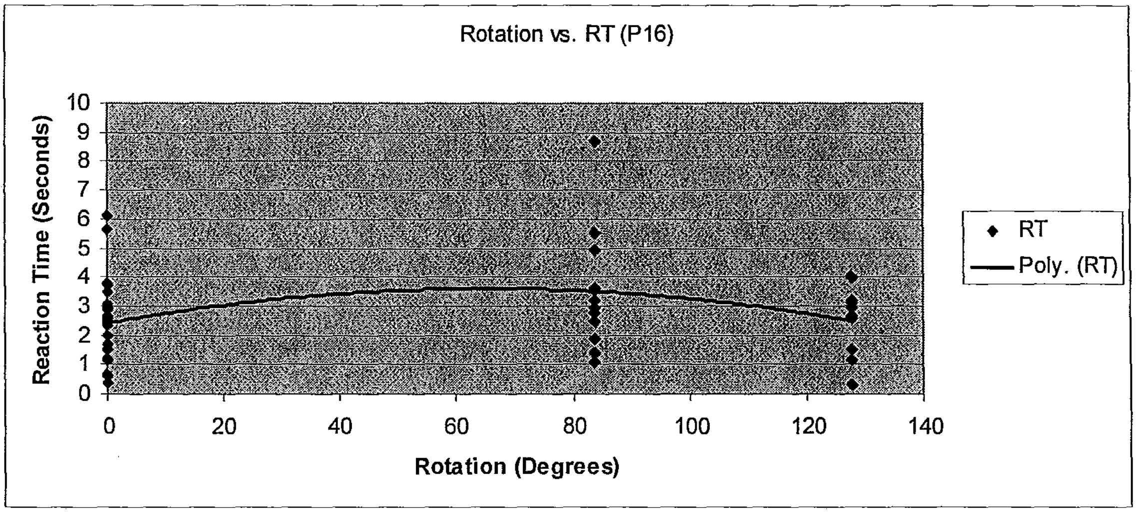

Figure 12 – Reaction times of participant 16 in each rotation condition

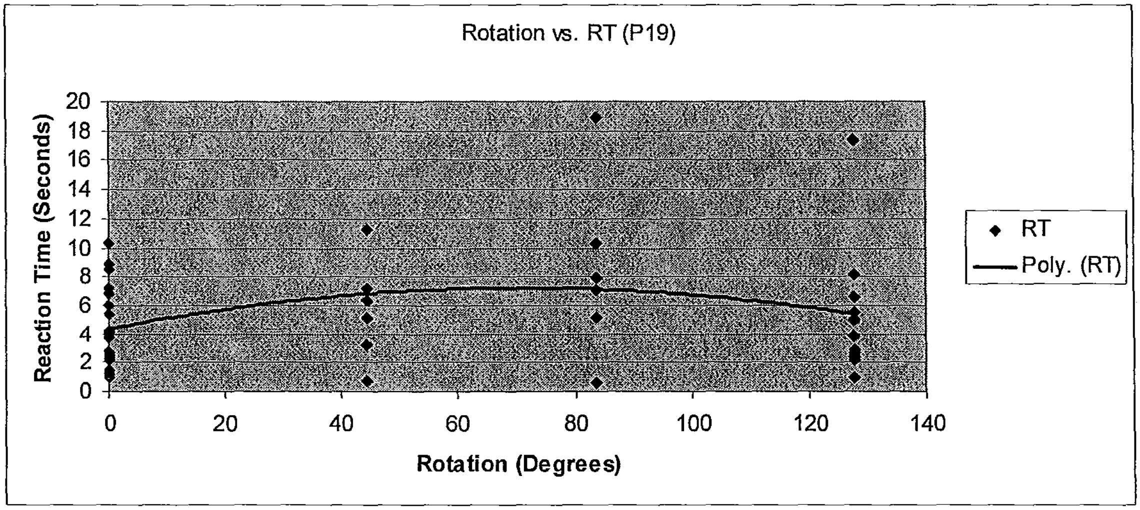

Figure 13 – Reaction times of participant 19 in each rotation condition

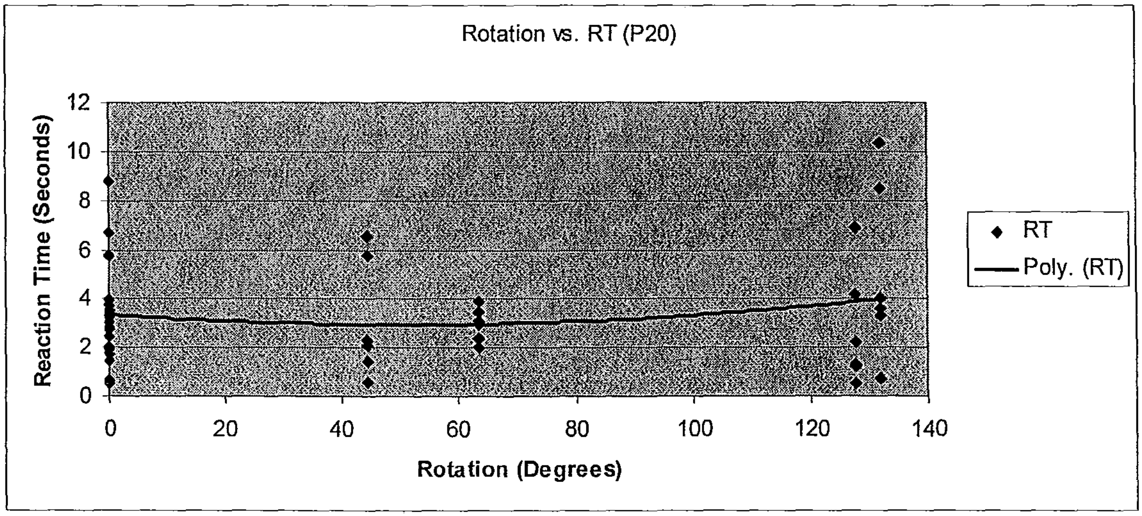

Figure 14 – Reaction times of participant 20 in each rotation condition

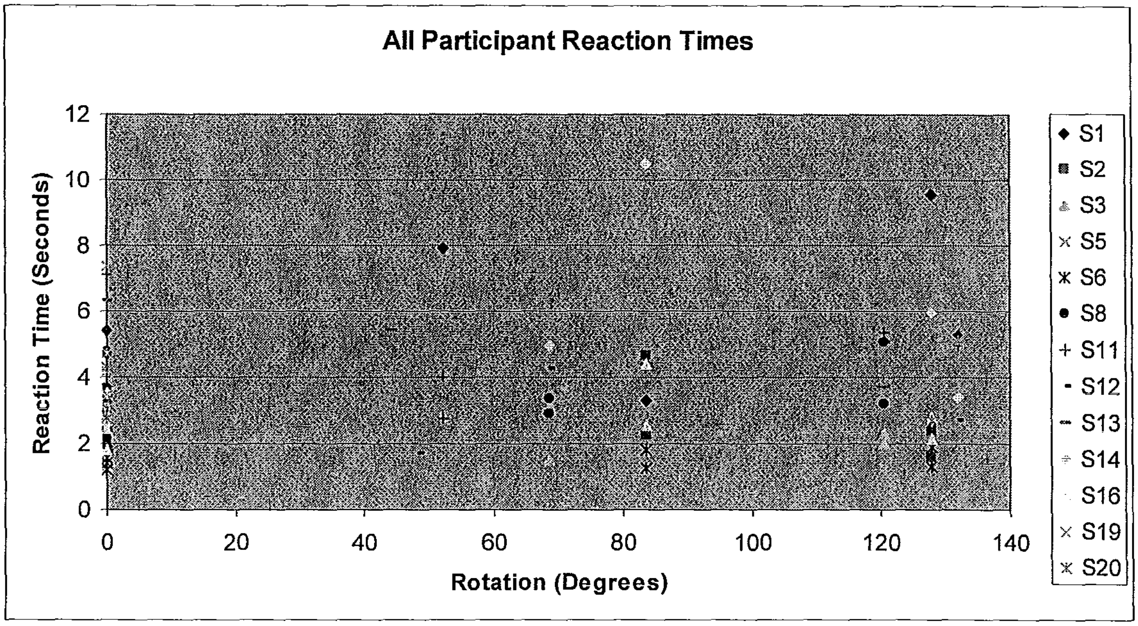

Figure 15 – Reaction times for all participants in each rotation condition done by the participant

A one-way repeated measures ANOVA was used to test for differences between reaction times in the three separate rotation conditions (zero degree rotation, low degree rotation and high degree rotation). It was found that reaction times differed significantly across the three rotation conditions, $F(2, 623) = 6.998, p = .001$. Post hoc comparisons showed (after Bonferroni adjustment for multiple comparisons) that reaction times significantly increased in the low degree rotation condition compared to the zero degree rotation condition, $p = .004$ and compared to the high degree rotation condition, $p = .014$. There was no significant difference between the low degree rotation condition and the high degree rotation condition $p = 1.000$. The average reaction time plus or minus the standard error for each rotation condition being: zero degree rotation = 3.065128 ± 0.136022, low degree rotation = 3.965128 ± 0.28471 and high degree rotation = 3.849231 ± 0.25038. These results are summarized in figure 16 below.

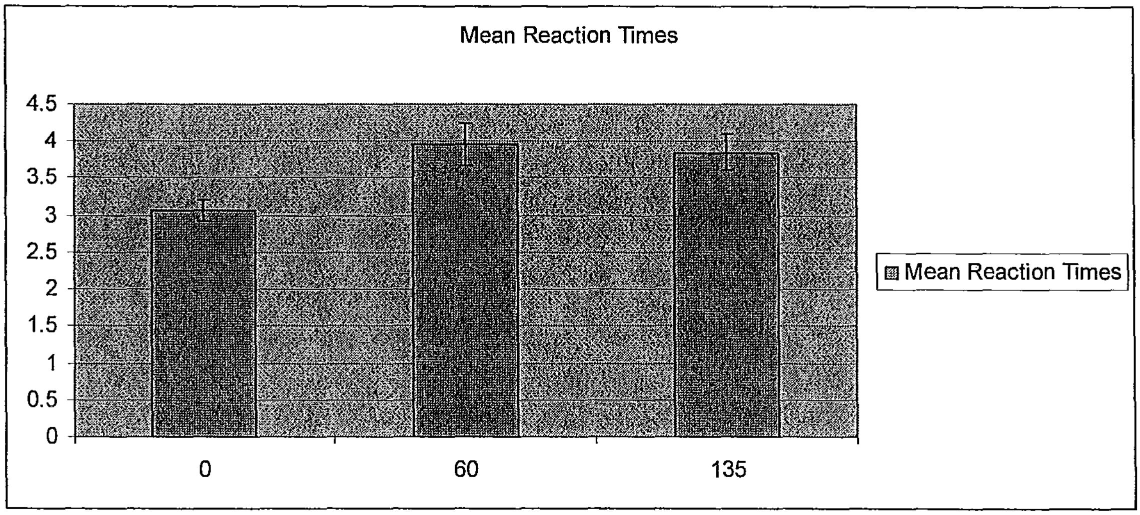

Figure 16 – Mean reaction times of all participants in each rotation condition

DISCUSSION

Underlying Goals

The research reported in this thesis addresses four related goals: 1) to create a virtual environment capable of effecting experiments which can be used to study the disorientation effect in spatial memory encoding; 2) to simulate, using our virtual environment, the methodology for studying the disorientation effect as described in Wang and Spelke (2000) and Waller and Hodgson (2006); 3) to replicate the results reported in experiment 4 by Waller and Hodgson (2006) especially as these pertain to the disorientation effect but basing these results on an experiment conducted in a virtual reality environment rather than in actuality; 4) to discuss the results obtained in terms of validating the use of virtual reality as a basis for further experiments specifically aimed at exploring spatial memory encoding as it pertains to the study of the effect of rotation on the disorientation effect.

1 A New Virtual Environment

Starting with an initial code developed by my collaborators (Sue Becker, John King and Christian Doeller), I created a script using the Unreal Engine 2 graphics engine to perform simulations in virtual reality. The virtual environment (a grassy field, a stone wall and a background) laid the foundation for an experimental set-up which could potentially transfer the study of the disorientation effect from experiments conducted in actual reality (as in Waller and Hodgson, 2006) to experiments conducted in virtual reality.

2 Realistic Simulation

In their fourth experiment, Waller and Hodgson (2006) prepared their participants by having them familiarize themselves with the location of 6 objects randomly located in a room. The participants walked around the room familiarizing themselves with the room and the locations of the objects. Then followed a series of tests involving self-rotations and pointing trials involving rotations of 0°, 45°, 90° and 135°. In their preparatory phase, participants learned four separate layouts. In other words, they had become quite familiar with the room and the locations of the objects before the actual pointing trials began. In the description of our method above, we indicate that we were able to closely imitate these preliminaries. To a significant extent, our participants also became quite familiar with our virtual environment before actual pointing trials began. All things being equal, if spatial encoding is affected by repetition and growing familiarity, our results should closely parallel those obtained by Waller and Hodgson (2006).

3 Results obtained

The results reported above show that we did obtain qualitatively similar results to those reported in Waller and Hodgson (2006). In fact, we did so more accurately than anticipated for we not only found that there is a sudden rise in errors after a certain amount of imagined viewpoint rotation, but that there is a main effect of rotation when we compare the low degree rotation with the high degree rotation condition, and no effect of rotation in the zero degree rotation compared to the low degree rotation conditions. In other words, our results suggest that viewpoint rotations under 107° have essentially no effect on accuracy of pointing. There is no evidence to support a gradual deterioration of the navigation system. At the same time, we saw equal reaction times for the low and high rotation conditions, suggesting that it was not just that the high rotation condition was more difficult but that perhaps a different strategy was being used, as will be discussed below. This supports Waller and Hodgson's criticism of Wang and Spelke (2000) who claim that there is only one system at work: the egocentric system which is highly accurate but becomes less so when disorientation becomes more severe. Waller and Hodgson (2006), on the other hand, claim that the increase of errors is due to a switching over from a purely egocentric system to an allocentric system. Their basic concept is simple: navigation in an egocentric system is done using constant updating calculations to maintain active knowledge of one's direction, current location and speed of movement, as well as active updating of the locations of objects around oneself by ongoing maintenance in visual-spatial working memory. If positional and/or rotational changes are slow enough and small enough, these calculations are not difficult to maintain. On the other hand, if the position or viewpoint changes too drastically then the required calculations for egocentric updating become too complex or difficult so one is more likely to resort to an allocentric strategy at this point. If people were using the same strategy of spatial updating for both low and high rotations, then one would predict gradually increasing reaction times along with gradually increasing large pointing errors as the viewpoint shift increased; we did not observe this. Instead, we observed relatively constant reaction times for both small and large viewpoint shifts, but a rather large decrease in accuracy for the large viewpoint shifts. Waller and Hodgson argue that Wang and Spelke are adhering too rigidly to a purely egocentric single-system argument. They argue that a single system should show continuous gradual deterioration as it loses the ability to update itself. What they see as problematic is the fact that the failure of the egocentric system becomes suddenly catastrophic rather than continuous. They also criticize the need to actually disorient participants in order to trigger the disorientation effect. They feel that merely updating spatial information from an imagined rotation of a sufficiently large degree should have the same disorienting effect. They show, in fact, that this is the case. In order to account for the lack of gradual disorientation and the ambivalent middle ground between no disorientation and high disorientation, they posit that two systems are actually at work: a highly accurate but somewhat unstable egocentric system and a less accurate but highly stable allocentric system. They claim that as the egocentric system becomes overburdened, there ensues a middle period of spatial confusion wherein both systems become simultaneously activated, with the egocentric system suddenly succumbing to complexity and being replaced by the allocentric system.

The results which we present in this thesis are consistent with Waller and Hodgson's interpretation which accounts for the fact that there is no effect of rotation on accuracy when we compare the zero degree rotation condition with the low degree rotation, but a large change in accuracy between the small and large viewpoint changes. Possibly, the increase of rotation combined with the lack of increase in variable error is evidence that the egocentric system is still at work and not yet overburdened in any significant manner.

So far, then, we have evidence which strongly supports Waller and Hodgson's criticism of Wang and Spelke (2000). Waller and Hodgson (2006) posit that there is a switch from the use of a purely egocentric system to an allocentric system for sufficiently large imagined rotations, or when a person is disoriented. Our results also support this conclusion.

The reaction time data are also somewhat surprising because they show a sudden decrease in reaction time after an imagined rotation. This could be taken as being consistent with a switch from a spatial updating mechanism for no rotation to an allocentric mechanism for rotations. If spatial information updating was being used for both low and high rotations, then one would expect the reaction time to increase steadily with increasing viewpoint shifts. This was not observed.

There are several possible explanations for our results:

1. The significant difference in reaction time between the 0 and low rotation conditions, in the absence of a significant increase in variable error, may be due to the spatial updating calculations presumably involved in employing a purely egocentric strategy. As the complexity of the task increases from zero to low rotation, spatial information updating is sufficient and hence variable error does not

increase, although reaction time does. As the complexity further increases with larger rotations, the egocentric system may become so inaccurate that it becomes more feasible to resort to an allocentric strategy and hence to switch systems. This would account for the sharp rise observed in variable errors. We still have to account for why there was no significant accompanying rise in reaction.

2. The lack of an increase in reaction time between the low-rotation and high-rotation conditions is surprising. One possible explanation for this finding is that, as the rotation steadily increases from low to moderate rotations, a critical threshold is reached, at which point the conflict between the two systems is resolved and the decision is taken to switch over to a constant-time lookup from allocentric long-term memory. Immediately, the conflict ceases, accuracy is dropped and reaction time remains constant. For increasingly large viewpoint shifts beyond that threshold, reaction time is not hampered by the additional calculations involved in spatial updating, but variable error is larger because of the switching over to a less accurate allocentric system. As we change from the low degree of rotation condition to the high degree of rotation condition, variable error becomes larger, but reaction time continues to remain fairly stable. According to this hypothesis, a) there are two systems in use and b) there is conflict between them which eventually leads to a switch from one to the other. To test this possibility, we would have to allow viewpoint rotation to vary continuously and run a great many more trials of this experiment. This is an important direction for future research.

3. There is also the possibility that system switching occurs at the start of the high degree phase of rotation. The large spike in pointing errors observed around the high

degree of rotation condition could be due to the double effects of a) increased reliance on the allocentric system and b) the 'release' of the predominantly accurate egocentric system. We tend to think of increased reaction time as reflecting an increase in workload, but there is also the possibility that such an increase might reflect the cost of switching between the two systems. It would be interesting to see if there are other thresholds beyond our high condition.

4 Validating the use of VR to carry out further experiments.

A further piece of evidence in favour of Waller and Hodgson's (2006) two-systems hypothesis are the results we observed in the low range of rotation. If, as Wang and Spelke (2000) claim, there is only an egocentric system at work, there should be a continuous more or less monotonic rise in variable error as the degree of rotation increases. In fact, this is not what our data has shown. What does happen is the variable error remains relatively stable during the zero degree rotation condition through the low degree rotation condition until quite suddenly the amount of variable errors increases rapidly. This does not accommodate itself to a single system theory very well for reasons discussed above and would be better explained by Waller and Hodgson's hypothesis that two systems are at work and that suddenly a switching mechanism is triggered to deactivate a more accurate system and activate a less accurate system. While the egocentric system is still active, variable error does not increase significantly but once that system goes offline and the allocentric system goes online, variable error increases significantly. These results are very encouraging because they show that the two-systems model can be used not only to obtain results of behavioural experiments in real environments, but also holds for behaviour in virtual environments, where there is the potential to have a much greater degree of control over events, concurrent reaction time measurements, timings, flow, activities, tasks, etc. This is a highly desirable form of control, given its accuracy and reliability (Christou and Bülthoff, 1999).

SIGNIFICANCE OF THESE RESULTS

Here we discuss the significance of the results obtained in this experiment with respect to: 1) reliability; 2) current limitations of fMRI experiments; 3) new experimental capabilities and 4) the variety of applications made possible using VR experiments such as ours.

1 Reliability of VR Experiments

As described in the introduction to this thesis, VR has been successfully used to carry out many experiments in psychology (e.g. Kosslyn et al., 1992; Christou and Bülthoff, 1999; King et al., 2002; Burgess, 2006). What we feel has been significant in the work reported here is the fact that it was possible to adapt a virtual reality engine created for another field i.e., gaming, to study spatial memory and imagery with enough precision and accuracy that it allowed us to simulate methods previously designed to be used only in actual reality experiments. Furthermore, not only did we manage to replicate these real world experiments in VR, but we also showed that the results were qualitatively very comparable. This degree of control and accuracy opens up a field which will enable us to generate virtual experiments which can simulate complex experiments that produce data with a degree of accuracy sufficient enough to become utilizable in their own right.

For instance, while it is still under debate whether the egocentric and allocentric levels of spatial representation constitute separate systems or a continuous representational hierarchy, mounting evidence makes it seem likely that one may respond based on either of the two levels. For example, one can close one's eyes and point to objects around the room, presumably a predominantly egocentric task; on the other hand, one can point to objects at a remembered location from a different remembered or even novel perspective, an activity which could involve egocentric and/or allocentric representations. This gets to the heart of a central question in memory research: what processes or mechanisms are involved when we engage in a spatial task from memory? If egocentric information is accessed/repeated, then are we actually performing mental navigation, mental rotation, or are we accessing a hippocampus-based cognitive map, using allocentric information, or are both egocentric and allocentric representations involved? Consider the case where the two shortcut-loving rats in Tolman (1948) scooted across the top of the maze and made their way directly to the food source: how did they manage that? Five plausible hypotheses immediately come to mind: 1) they smelled the food and followed their noses; 2) egocentric navigation; 3) allocentric navigation; 4) a combination of egocentric and allocentric navigation; 5) none of the above. Because we do not know the conditions governing the olfactory environment of the experiment, we cannot rule out the first possibility. The fifth alternative might involve navigation through extrapolation and re-routing. Given the multiple functions associated with the hippocampal region which we discussed briefly above, this hypothesis is not a candidate for immediate dismissal although testing for it might be problematic. The same considerations possibly apply to any spatial task based on remembered locations. This is one area where supporting evidence based on brain activity such as can be obtained through a series of fMRI experiments is needed to confirm or disprove hypotheses such as those we have just considered.

As discussed above, we believe that our data supports a two-system explanation of spatial memory after disorientation/viewpoint shifts. Further experiments in VR could be designed to pinpoint the conditions which trigger the system switch which may be caused by the disorientation effect. We also noted a surprising discrepancy between increase in variable errors and reaction time results. In future experiments, by systematically varying the degree of rotation in a continuous manner, the nature of the effect indicated above could be explored in increasingly fine detail.

2 Current Limitations of Neuroimaging

Currently, spatial memory experiments involving neuroimaging suffer from several severe limitations. Many of these limitations are due to the fact that body mobility is almost completely eliminated during these experiments. From the perspective of doing experimental research on movement, spatial information encoding and so forth, this imposed paralysis reduces the range of experiments which could be done using neuroimaging techniques such as fMRI since most (animal) navigating experiments and activities occur normally with body movement enabled (Maguire et al 2000;Wang & Spelke, 2000; Waller & Hodgson 2006). If we could set up virtual reality experiments which are known to yield reliable results, and the results we have obtained and reported in this thesis indicate that virtual reality experiments are remarkably accurate, then we could use neuroimaging techniques to confirm or contradict hypotheses underlying the disorientation effect. For instance, it would be very interesting to see if there is neurobiological evidence, in the form of separate brain systems being activated, corresponding to a switch between an egocentric system and an allocentric system when rotation becomes large enough to cause the disorientation effect and where, when and how this switch occurs. Equally interesting would be to map out activity regions and any interplay between these regions during such experiments as the one we used to establish the results reported in this thesis. Perhaps even more interesting yet is the exciting possibility of reversing the roles between virtual experiments and neuroimaging techniques. One promising possibility is to rethink the use of neuroimaging experiments. Instead of using neuroimaging results in a confirming role, it could very well be possible to use them as guides in designing and generating relevant new kinds of virtual experiments. This would greatly enhance the use of such experiments.

3 New Experimental Capabilities

The discussion above regarding the lack of predictable relations between variable errors and reaction time indicates that neuroimaging and virtual environment tools can combine to yield fine-grained experiments that show promise in resolving issues. As noted above, a series of experiments based on refining the rotation conditions we used, continuously or discretely, say, at every 5 or 10 degrees, could be used to investigate whether reaction times, systemic conflicts, systemic switching etc. continue to correlate, or whether unexpected variations begin to be triggered.

4 Variety of Applications Made Possible Using VR Experiments.

In this short section, we are interested in considering what these experiments could contribute to neuropsychological research in general. Combining the possibilities of using virtual reality experiments and neuroimaging techniques, we could set up experiments similar to Pine et al. (2002), but which would allow us greater freedom in mapping out brain activities related to the disorientation effect and in much greater detail than what was previously possible. The brain remains largely a black box for moral considerations and because of the simple fact of mortality. Mapping the activities of this black box has shown itself to be replete with difficulties mostly because of extreme complexity and complete lack of knowledge. The use of virtual reality experiments shows promise when seen more humbly as a viable and realizable new capacity to make progress in small steps accumulatively. As each experiment yields reliable information, the model becomes more accurate. Given that billions, perhaps even trillions of factors are involved, we need a tool which can deal with such quantities in realistic time intervals. Virtual reality may well open up that possibility. It may even be pointing to possibilities as yet unsuspected and as yet completely unknown, such as dynamic, multidimensional flow charts, or state charts which could be combined to yield dynamic models of brain activity in real-time.

BIBLIOGRAPHY

Allen, Gary L. (1981). A developmental perspective on the effects of subdividing macrospatial experience. *Journal of Experimental Psychology: Human Learnimg and Memory, 7*, 120-132.

Allen, Gary L. and Haun, Daniel B. M. (2004). Proximity and precision in spatial memory. In: *Human Spatial Memory.* Lawrence Erlbaum Associates, Pubs., 41-63.

Bird C.M., Burgess N. (2008). Insights from spatial processing into the Hippocampal role in memory. *Nature Reviews Neuroscience, 9*, 182-194.

Bohbot, Véronique D., Lerch, Jason, Thorndycraft, Brook, Iaria, Giuseppe and Zijdenbos, Alex P. (2007). Gray Matter Differences Correlate with Spontaneous Strategies in a Human Virtual Navigation Task. *The Journal of Neuroscience, 27:38*, 10078-10083.

Burgess, N. (2006). Spatial memory: How egocentric and allocentric combine. *Trends in Cognitive Sciences, 10:12*, 551-557.

Burgess, N., Jeffery, K. J., and O'Keefe, J. (Eds.) (1999*). The hippocampal and parietal foundations of spatial cognition.* New York: Oxford University Press.

Burgess N., Maguire E.A., O'Keefe J. (2002) The human hippocampus and spatial and episodic memory. *Neuron 35,* 625-641.

Burgess, N., Becker, S., King, J. and O'Keefe (2001). Memory for events and their spatial context: models and experiments. *Philosophical Transactions of the Royal Society of London B, 356,* 1493-1503.

Cheng, K. (1989). The vector sum model of pigeon landmark use. *Journal of Experimental Psychology: Animal behavior processes, 15,* 366-375.

Christou, C. G., and Bülthoff, H. H. (1999). View dependence in scene recognition after active learning. *Memory & Cognition, 27*, 996-1007.

Creem, S. H., and Proffitt, D. R. (1998). How do the cerebral hemispheres contribute to encoding spatial relations? *Current Directions in Psychological Science, 7*, 8-14.

Creem, S. H., and Proffitt, D. R. (2001a). Defining the cortical visual systems: "What," "where," and "how." *Acta Psychologica, 107*, 43-68.

Creem, S. H., and Proffitt, D. R. (2002b). Grasping objects by their handles: A necessary interaction between cognition and action. *Journal of Experimental Psychology: Human Perception and Performance, 27*, 218-228.

Diwadkar, V. A., and McNamara, T. P. (1997). Viewpoint dependence in scene recognition. *American psychological society, 8,* 302-307.

Doeller, Christian F. and Neil Burgess (2008). Distinct error-correcting and incidental learning of location relative to landmarks and boundaries. *PNAS, 105:15*, 5909-5914.

Doeller, C. F., King, J. A., and Burgess, N. (2009). Parallel striatal and hippocampal systems for landmarks and boundaries in spatial memory. *PNAS, 105:15,* 5915-5920.

Eichenbaum, H. (2007). Comparative cognition, hippocampal function, and recollection. *Comparative Cognition & Behavior Reviews, 2*, 47-66

Ekstrom, A. D., Kahana, M.J., Caplan, J.B., Fields, T. A., Isham, E.A., Newman, E.L., and Fried, I. (2003). Cellular networks underlying human spatial navigation. *Nature, 425*, 184-187.

Hartley, T., Maguire, E. A., Spiers, H. J., and Burgess, N. (2003). The well-worn route and the path less traveled: Distinct neural bases of route following and wayfinding in humans. *Neuron, 37,* 877-888.

Hassabis, Demis, Kumara, Dharshan, Maguire, Eleanor A. (2007b). *The Journal of Neuroscience, 27:52*, 14365-14374.

Hassabis, D. Kumara, D. Vann, S. D. and Maguire, E. A. (2007a). Patients with hippocampal amnesia cannot imagine new experiences. *PNAS, 104:5*, 1726-1781.

Huttenlocher, J., Hedges, L. V., and Duncan, S. (1991). Categories and particulars: Prototype effects in estimating spatial location. *Psychological Review, 98*, 352-376.

King, John A., Burgess, Neil, Hartley, Tom, Vargha-Khadem, Faraneh, and O'Keefe, John. (2002). The human hippocampus and viewpoint dependence in spatial memory. *Hippocampus, 12:6*, 811-820.

Kosslyn, S. M. (1987). Seeing and imagining in the cerebral hemispheres: A computational approach. *Psychological Review, 94*, 148-175.

Kosslyn, S. M. (1994). *Image and brain: The resolution of the imagery debate.* Cambridge, MA: MIT Press.

Kosslyn, S. M., Chabris, C. F., Marsolek, C. J., and Koenigh, O. (1992). Categorical versus coordinate spatial relations: Computational analysis and computer simulations. *Journal of Experimental Psychology: Human Perception and Performance, 18*, 562-577.

Maguire, E. A., Burgess, N., Donnett, G., Frackowiak, R. S. J., Frith, C. D., and O'Keefe, J. (1998). Knowing where and getting there: A human navigation network. *Science, 280*, 921-924.

Maguire, E. A., Gadian, David G., Johnsrude, Ingrid S., Good, Catriona D., Ashburner, John, Frackowiak, Richard S. J., and Frith, Christopher D. (2000). Navigation-related structural change in the hippocampi of taxi drivers. *PNAS, 97:8*, 4398-4403.

McNamara, T. P., & Valiquette, C. M. (2004). Remembering where things are. In *G. Allen (Ed.): Human spatial memory: Remembering where.* Mahwah, NJ: Erlbaum. 3-24.

Moscovitch, M., Rosenbaum, R.S., Gilboa, A., Addis, D.R., Westmacott, R., Grady, C., McAndrews, M.P., Levine, B., Black, S.E., Winocur, G. & Nadel, L. (2005). Functional neuroanatomy of remote episodic, semantic and spatial memory: A unified account based on multiple trace theory. *Journal of Anatomy, 207*, 35-66.

Newcombe, N. S., and Huttenlocher, J. (2000). *Making space: The development of spatial representation and reasoning.* Cambridge, MA: MIT Press.

O'Keefe, J., and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res., 34,* 171-175.

O'Keefe, J. (1976). Place units in the hippocampus of the freely moving rat. *Experimental Neurology, 51*, 78-109.

O'Keefe, J., and Nadel, L. (1978). *The hippocampus as a cognitive map*. Clarendon Press.

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

O'Keefe, J., & Recce, M. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus, 3*, 317-330.

O'Keefe J. and Burgess N. (1996) Geometric determinants of the place fields of hippocampal neurons. *Nature 381*: 425-428

Piaget, J., Infelder, B. (1967). *The child's conception of space.* (F. J. Langdon & J. L. Lunzer, Trans.) New York: Norton (Original work published in 1948).

Pine, D. S., Maguire, E. A., Burgess, N., Zarahn, E., Koda, V., Fyer, A., Szeszko, P. R., and Bilder, R. M. (2002). Neurodevelopmental aspects of spatial navigation: A virtual reality fMRI study. *Neuroimage, 15*, 396-406.

Roskos-Ewoldsen, B., McNamara, T. P., Shelton, A. L., and Carr, W. S. (1998). Mental representations of large and small spatial layouts are orientation dependent. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 24*, 215-226.

Sholl, M. J., and Nolin, T. I. (1997). Orientation specificity in representations of place. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 23,* 1496-1507.

Tolman, Edward C. (1948). Cognitive maps in rats and men. *The Psychological Review, 55:4*, 189-208.

Waller, D., and Hodgson, E. (2006). Transient and enduring spatial representations under disorientation and self-rotation. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 32,* 867-882.

Wang, R. F., and Spelke, E. S. (2000). Updating egocentric representations in human navigations. *Cognition, 77,* 215-250.

Wolbers, Thomas, Hegarty, Mary, Büchel, Christian, and Loomis, Jack A. (2008). Spatial updating: how the brain keeps track of changing object locations during observer motion. *Nature Neuroscience, 11:10*, 1223-1230.