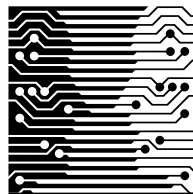# THE SPATIAL LEARNING METHOD: FACILITATION OF LEARNING THROUGH THE USE OF COGNITIVE MAPPING IN VIRTUAL REALITY

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF COMPUTER SCIENCE,
FACULTY OF SCIENCE
AT THE UNIVERSITY OF CAPE TOWN
IN FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
MASTER OF SCIENCE

By
Cathryn Johns
February 2002

Supervised by
Edwin Blake

Abstract

When moving through an environment, people unconsciously build up a mental image, or cognitive map, of that environment. When later planning a trip or giving directions to someone else, they can mentally walk through the environment, remembering features relevant to their current task. This process of building up a cognitive map of a place and using that map later is called cognitive mapping.

This dissertation presents the novel idea of using the cognitive mapping process to teach relationships between data items, called the spatial learning method. By creating a VE where the buildings or rooms represent data items, and the paths between the buildings or rooms indicate the relationships between the data items, visitors exploring the VE would not only be building up a cognitive map of the environment, but also learning the relationships implied by the layout.

To investigate the feasibility of the spatial learning method, such a VE was created. Three studies using this VE were run concurrently on a single set of 26 participants. The first study investigated whether visitors to a VE can in fact build up an accurate cognitive map of it, and studied the effect of the VR display system on the cognitive mapping process as well as the effect of having provided participants with a map of the VE. The second study investigated whether data relationships can be inferred from the cognitive map, looked at the effect of display type and having a map on learning, and compared learning via the spatial learning method with that via a conventional lecture (presented to a separate group of 7 participants). The third study examined the relationship between various psychological factors (emotions such as enjoyment, interest, and distress, as well as the sense of presence) and cognitive mapping and learning.

Study 1 found that while most participants did build up a cognitive map of the VE used in the study, the maps were generally of low quality. Study 2 showed that the learning of the underlying data set varied greatly between participants, with some remembering almost all of the data points and the relationships between them while others could barely answer the most rudimentary questions about the data set. Study 2 also showed that participants who attended the conventional lecture performed significantly better at the learning test than participants who were taught via the spatial learning method. Study 3 found that the participants who attended the VR sessions did not experience any more positive emotions that those who attended the lecture, and also showed that emotions and the sense of presence were unrelated to both cognitive mapping and learning. Studies 1 and 2 also showed that using an immersive, head-tracked VR system as opposed to a desktop system did not affect either cognitive mapping or learning.

With these findings it is difficult to recommend the use of the spatial learning method as a teaching tool. However, some of the results are encouraging, and it may be possible to improve the method at least to the point where it could be used a teaching aid to supplement conventional methods rather than replacing them.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

While Virtual Reality (VR) is primarily used for manufacturing, entertainment and training (such as training shipboard firefighters in VR before progressing on to training in real fire situations [75]), it has also been used for education in schools [25]. This type of education has mostly focused on simulating a particular environment and allowing children to explore the consequences of various actions — for example, Jackson *et al* [37] used VR to create a Virtual Environment (VE) in which children could change world parameters (such as increasing the emission of greenhouse gases) and see the consequences of their actions (such as the effect on the global temperature). Osberg *et al* [57] has investigated the educational technique of allowing children to create their own VEs, which requires them to research the subject material and to understand it well enough to use it to create a VE, rather than passively absorbing the same information through conventional means such as a classroom lecture. The key element in using VR in education is that it allows students to interact directly with information [57], creating an experience that is more engaging than traditional teaching methods.

This dissertation presents a new way of using VR for education which combines direct interaction with physical exploration. This concept is called the spatial learning method.

1.1 The Spatial Learning Method

With the use of Virtual Environments, data can be presented in truly 3-dimensional form, allowing users to enter into the data and study relationships close up. By representing a data set in such a spatial form, with spatial relationships between objects in a VE representing the relationships between data items, we can represent almost any set of data in VE whether it is spatially based or not.

For example, a database of famous composers could be spatially represented within a VE by representing each composer with a statue, with the distance between composers indicating the degree of similarity of their music styles. By using spatial hyperlinks, or *teleports*, different sets of spatial relations can be created for the same data set — so in the composers example, teleports could be set up to link statues in a temporal relationship in addition to the similarity of style relationship. In such a VE, the geography and environmental features are based on data items and the relationships between them.

This could be a useful way of presenting data to be learned. By setting up data relationships as spatial relations, we make use of the brain's well-developed spatial abilities. Just by exploring such a VE, the visitor is subconsciously acquiring knowledge about the layout of the environment — and thus also knowledge about the data items relate to each other.

This occurs because when people move around new environments, whether they are real or virtual, they subconsciously build a mental image of the space they are in [24]. This mental image is encoded in the hippocampus [4], and is called a "cognitive map" [41, 53, 76]. It helps people find their way in environments that they have visited before, and also helps them remember the structure of the place, for example if they are asked for directions [24]. This process is called cognitive mapping.

Making use of cognitive mapping to teach data relationships by creating a VE with the layout based on the data set could be called *spatial learning*. This dissertation presents a set of exploratory studies into the feasibility of using this novel approach (described in more detail in Chapter 3) to teach data sets.

1.2 Aims

This research explores three areas that need to be investigated before the use of the spatial learning method can be recommended: those issues related to cognitive mapping in Virtual Environments (VEs); those issues related to the feasibility of the spatial learning method; and the psychological impact of VEs and spatial learning on the user. The ultimate aim is to either recommend the spatial learning method and provide recommendations as to what equipment will be needed and what other factors need to be considered, or to eliminate the method as a feasible learning option.

1.2.1 Cognitive Mapping in VEs

It needs to be established whether a visitor to a VE can form a cognitive map of that environment within the limited amount of time that is normally available to the visitor to explore the environment. It must also be established whether immersive equipment (such as head-trackers and Head Mounted Displays) is necessary to form a cognitive map of the VE, or whether a keyboard-and-monitor-based desktop system is sufficient. It must be determined whether it is necessary to provide visitors with a map of the structure of the virtual environment, or whether this causes them to pay less attention to their surroundings (thereby detracting from their ability to form an accurate cognitive map of the area). Finally, the effect of a visitor's spatial abilities on their ability to form a cognitive map of the VE must be studied.

1.2.2 Feasibility of the Spatial Learning Method

It is important to determine whether it is possible for a visitor, having built up a cognitive map of a VE, to absorb and comprehend the underlying data set not as a map, but in terms of the actual data points and the relationships between them. Again, it must be established whether a desktop system supports the desired result, or whether a immersive system is necessary. Also, the effect of a participant's spatial abilities on their ability to absorb the

underlying data set must be determined. Finally, having established that the underlying data set can be absorbed, the extent to which this occurs should be compared against the amount information retained during a conventional lecture-style presentation.

1.2.3 Psychological Impact

The impact the use of a VR system (whether desktop-based or immersive) on a visitor's emotions must be investigated, and it must be determined whether these effects are due simply to the novelty of the situation or whether they are also felt by those with more experience in VR. The emotions evoked by the VR system used in the spatial learning method should be compared with the emotions evoked by a conventional presentation of the data set. In addition, another aspect of the visitor's psychological experience, namely that of *presence*, should be considered. It would be useful to determine whether a stronger sense of presence helps in forming a cognitive map of a VE, and whether it affects the ability to translate this map into an understanding of the underlying data set.

1.3 Overview of Methodology

Three exploratory studies have been designed to investigate the areas of interest outlined above. These studies were run concurrently on single set of 33 participants, who were divided into five groups:

- Those using an immersive Head Mounted Display (HMD) system with a map of the VE (4 participants);

- Those using an immersive HMD-based system without a map (4 participants);

- Those using a desktop system with a map (9 participants);

- Those using a desktop system without a map (9 participants);

- Those attending a lecture on the data set (7 participants).

The number of participants in the HMD conditions was low because of the high drop-out rate due to simulator sickness. Approximately one in two participants felt dizzy and nauseous during the HMD session (see Sections 4.4.6 and 8.1 for more details). Because of these low sample sizes, nonparametric statistics were used to analyze the data (see Section 4.8).

1.3.1 Study 1 — Cognitive Mapping

Study 1, reported in Chapter 5, investigated the issues surrounding cognitive mapping in VE. The study consisted of two parts. The first part was an experimental study which tested two hypotheses: firstly, that the VR display system and input device used would affect cognitive mapping; and secondly, that providing participants with a dynamic You-Are-Here map showing position and orientation would affect cognitive mapping.

The second part of Study 1 was a relational study which investigated the relationship between several other factors and cognitive mapping, namely: the percentage of the environment that was explored; whether or not the experimental task was finished; previous

VR/gaming experience; the age of the user; the gender of the user; the primary travel mode of the user; the spatial abilities of the user; the percentage of time that a map of the VE was visible; and psychological factors such as emotional aspects and presence.

1.3.2 Study 2 — Learning the Data Set

Study 2, reported in Chapter 6, investigated the feasibility of the spatial learning method. Again, the study consisted of two parts. The first section was an experiment study examining two hypotheses: firstly, that the presentation style (HMD VR system, desktop VR system, and conventional lecture) would affect the learning of the data set; and secondly that providing VR participants with a dynamic You-Are-Here map showing position and orientation would affect their learning of the data set.

The second part of the study was a relational study investigating the relationship between several other factors and learning — specifically, the extent to which a cognitive map of the environment was formed; psychological factors, such as emotions (enjoyment, fear, surprise, shyness, interest, distress) and the sense of presence; the percentage of the environment that was explored; the amount of time spent in the VE; the age of the user; the gender of the user; the spatial abilities of the user; and the percentage of time that a map of the VE was visible.

1.3.3 Study 3 - Psychological Impact

The final study, Study 3 (reported in Chapter 7), investigated whether the VR display system influenced a visitor's emotions and sense of presence. Again, it consisted of two parts. The first part was an experimental study, with two hypotheses: firstly, that presentation style (HMD VR system, desktop VR system, and conventional lecture) would affect the emotions of the participants, and that within the VR conditions the display type would affect the sense of presence of the participants; and secondly that providing VR participants with a dynamic You-Are-Here map showing position and orientation would affect their emotions and sense of presence.

The second part of study 3 was again a relational study which examined the relationship between the participant's emotions and sense of presence and various other factors, namely: the extent to which a cognitive map was formed of the VE; the extent to which the data set was learned; the amount of time spent in the VE ; the percentage of the environment that was explored; whether or not the experimental task was completed; previous VR/gaming experience; gender; and age. In addition, the relationship between the emotions felt by participants and their sense of presence was investigated.

1.4 Outline of Dissertation

Chapter 2 presents background material relevant to this research. The various types of VR display systems that support navigation are discussed, along with the different types of input devices. In terms of the psychological impact of VR on users, the concept of presence is presented along with a discussion on its measurement.

The concept of a cognitive map is presented, along with some of the controversy about the exact form that this mental representation takes. The formation of a cognitive map is

discussed, along with the consequent implications for VEs, and the types of distortions that occur in cognitive maps are reviewed. The various methods of measuring cognitive maps are also presented.

The concepts behind navigation and wayfinding are presented, along with a discussion of navigation issues particularly relevant to VEs. The various types of tools that may aid in navigation are discussed, and guidelines for designing navigable VEs are presented. The importance of an individual's spatial abilities is discussed, as are the factors which influence these abilities.

Chapter 3 presents the theory behind the spatial learning method, and several examples are given. Several related learning methods are discussed. The data set requirements are reviewed, and some potential applications are presented. Finally, the concept of non-linear spatiality is presented, along with some potential applications.

Chapter 4 begins by reviewing the aims of this research, along with the three studies designed to investigate these aims. The rest of the chapter describes the methodology used, which is common to all three studies. The data set used in these studies is described, as is the VE that was designed to present the data set. The conventional presentation is also discussed, along with the experimental procedure for both the VR and conventional presentations. The methods of data collection are described as well as the measures used in the studies. Finally, the methods of analysis are discussed.

Chapter 5 describes the first of the three studies, which investigates the effect of various factors on cognitive mapping. The aims of the research specific to this study are reiterated, followed by the results of the study. These results and their implications are then discussed.

Chapter 6 describes the second of the three studies, which tests the spatial learning method by comparing the results of the VR presentation and the conventional presentation, and also examines the effect of several factors on VR learning. The aims of the study are reiterated, followed by the results of the study. These results and their implications are then discussed.

Chapter 7 describes the last of the three studies, which investigates the psychological impact of the VR learning experience. The aims of the study are reiterated, followed by the results of the study. These results and their implications are then discussed.

Finally, Chapter 8 summarizes the conclusions drawn from the three studies, and considers their combined implications for the spatial learning method. Recommendations as to the use of the spatial learning method are presented, along with an analysis of the necessary further research into the method.

There are several appendices to this dissertation. Appendix A shows some screenshots of the digital questionnaire presentation (described in Section 4.6.1). Appendix B presents the Differential Emotions Scale used to assess the emotional impact of the experience (described in Section 4.7.1), while Appendix C presents the Presence Questionnaire used to measure the sense of presence felt by the VR participants (described in Section 4.7.1). Similarly, Appendix D presents the Everyday Spatial Abilities Test (described in Section 4.7.2) used to measure participants' spatial abilities. The last of the appendices, Appendix E, presents the test given to participants in order to measure how of the data set they had learned (see Section 4.7.6). The correct answers are also included.

Chapter 2

Background

The concepts behind the spatial learning method are drawn from a range of disciplines, such as Virtual Reality, environmental psychology, and urban planning. This chapter presents background information from these fields that is relevant to this research, such as the various VR display systems and input devices (Section 2.1.1, the sense of presence, in terms of its psychological effect on participants (Section 2.2), cognitive mapping theory (Section 2.3), navigation and wayfinding theory (Section 2.4), particularly with regard to VEs (Section 2.4.4), and spatial abilities research (Section 2.5).

2.1 Virtual Reality

2.1.1 VR systems and Input Devices

There are many different types of Virtual Reality systems. Some systems, such as interactive workbenches, are used primarily for examining and working on a single object. Other systems, such as Head-Mounted-Displays and CAVEs, are used primarily to act as an interface to a Virtual Environment, which users can enter and be surrounded by. This second type of system is the kind which is referred to in this dissertation.

There are three main types of systems used to display Virtual Environments:

1. Desktop systems, where the display is a monitor screen. Stereo may or may not be available, depending on whether some type of stereo glasses are used.

2. Head-Mounted-Displays (HMDs), where the display is two small monitors directly in front of the eyes. By displaying slightly offset views to each eye, a stereo effect can be created.

3. CAVEs, where the display is projected on to three walls and the floor (sometimes all four walls and the floor) and the user stands in the middle. Stereo glasses are used to create a stereo effect.

Each of these types of systems require the user to be able to move around inside the environment being displayed, and there are many different devices used to accomplish this. These devices can be divided into two categories: direct, and indirect. Direct devices map the user's physical body movements directly into the virtual environment, and include

treadmills and trackers. Indirect devices interpret commands made by the user (such as those via a keyboard, mouse, or joystick) into movements within the virtual environment.

Direct input devices provide more physical cues for the user. For example, turning one's head in the real world causes the virtual display to change as it would as if it were real, and so the physical feedback such as muscle tension and inner ear stimuli correspond to the virtual movement [35, 86]. With indirect devices, however, the only physical feedback is the sensation of a finger hitting a key or pressing a button.

Unfortunately, even the best display and input systems fall far short of the real world, Even direct devices cannot map every body motion into virtual motion, as every portion of the body would need to be tracked and processed, and even the best display systems have limitations. Low display resolutions, accuracy of colours, limited field-of-view, and lack of peripheral vision all decrease the system fidelity (the extent to which the Virtual Environment is indistinguishable from a real environment) [66, 78, 86, 87]. Issues such as display update rate and latency also affect fidelity [86], as does the lack of positional sound (and the lack of normal, everyday noises, as it is practically impossible to include all of these in a VE). Lack of haptic cues [44], the fact that not all bodily actions can be translated into VE actions, and the limited number of interactions available within a VE all serve to make the virtual experience more impoverished than a real experience. This must be remembered when studying complex phenomena such as navigation and wayfinding in Virtual Environments.

The fact that the amount of time spent in Virtual Environments is in terms of minutes or hours also affects research into navigation, wayfinding, and cognitive mapping, as with results from real environments the amount of time spent is normally measured in terms of days, months, and years.

2.2 Presence

Much research in Virtual Environments (VE's) is concerned with the concept of presence, that of being in an actual place when one is in a VE. Presence has been linked to greater knowledge transferability [70] as well as enhancement of learning and performance [89]. More generally, presence is informally used as a measure of how "good" a virtual environment is. This research investigates whether the sense of presence affects cognitive mapping and the learning of the data set underlying the layout of the VE — if the participant feels that they are in a real place, they may react to it as real in terms of cognitive mapping (and thus learning) as well as navigation and wayfinding.

2.2.1 Defining Presence

The general understanding of the concept of presence is that it is the sense of *being in* a virtual environment. Witmer & Singer [89] define presence as "the subjective experience of being in one place or environment, even when one is physically situated in another". In the case of teleoperators, presence is the sense of being at the remote site, rather than at the operator's console. In the case of Virtual Environments, presence refers to experiencing the simulated environment, rather the physical surrounding one. Slater *et al* [70, 71] take this concept further, defining presence as a state of consciousness, the psychological sense of being in the virtual environment, together with corresponding modes of behaviour. Users

who experienced a high sense of presence should remember the experience as having visited a place, not as having seen images. Behaviour in the VE should be consistent with behaviour in similar situations in everyday life.

Witmer & Singer [89] believe that presence is a normal awareness phenomenon, and is based on the interaction between sensory stimuli, environmental factors which encourage involvement, and internal tendencies to become involved. They say that presence in a VE depends an attention shift from the real world to the VE. However, it does not require the total removal of attention from the real world. The degree to which attention is shifted away from the real world determines the amount of presence felt by the user. Thus, presence is a matter of focus. However, they also present an alternative view, namely that presence may be similar to *selective attention*. Selective attention describes the tendency to focus on only relevant or interesting information. The argument is that experiencing presence in a VE requires the ability to focus on one set of relevant stimuli, those of the VE, to the exclusion of irrelevant stimuli from the real world. In this model, both involvement and immersion are necessary to experience presence.

Slater *et al* [70] hypothesize that presence is an increasing function of two orthogonal variables. The first is the extent of the match between the displayed sensory data and the internal representation and world model employed by the participant. The second is the extent of the match between proprioception and sensory data, i.e. the consistency between changes to the display must be consistent with changes caused by the participant's movement and location.

2.2.2 Measuring Presence

Presence is most often measured by means of a self-report questionnaire given to participants after they leave the VE. Generally, a numeric scale is used, a technique known as magnitude estimation. However, this technique does have some possibly serious flaws. One of these is the "range effect", whereby participants' numerical ratings are strongly influenced by the range of the physical stimulus to which they are exposed. Another possibly important flaw is the "anchor effect", whereby the value which participants assign to a given condition depends on the conditions to which it is compared [59]. These effects make it difficult to compare data on presence between separate experiments, but fortunately seem not to have an effect within a single experiment.

Slater *et al* [70, 71] use a 1 to 7 scale to measure subjective presence. The six questions are based on three basic determinants, namely

- The sense of "being in" the VE;

- The extent to which there were times when the virtual world seemed more the presenting reality than the real world; and

- The sense of having visited a place rather than seeing images.

Two sample questions from their questionnaire read as follows:

1. Please rate *your sense of being there* in the virtual reality.

2. When you think back about your experience, do think of the virtual reality more as *images that you saw*, or more as *somewhere that you visited*?

Prothero *et al* [59, 60] also use a 1 to 7 scale questionnaire. Samples of items on their questionnaire:

1. In sharkworld, I felt like ... (1 = I was standing in the laboratory, wearing a virtual reality helmet; 7 = I was in some sort of ocean, near a shark-infested shipwreck).

2. Did the virtual world seem more like a picture, or more like a scene looked at through a window? (1 = like a picture; 7 = like looking through a window).

Witmer & Singer [89] developed a presence questionnaire (the PQ) based on four categories of factors which influence presence. Items on the PQ are also measured on a seven-point scale. Some sample questions from the PQ are as follows:

1. How inconsistent or disconnected was the information coming from your various senses?

2. How closely were you able to examine objects?

3. How distracting was the control mechanism?

4. How proficient in moving and interacting with the virtual environment did you feel at the end of the experience?

Witmer & Singer's questionnaires appear particularly attractive as subjective measures of presence, as they seem to be aware of reliability and validity issues, and have attempted to show that their questionnaires are both reliable and valid. A recent study by Johns *et al* [38] suggests, however, that the PQ is not suitable for comparing presence levels across different Virtual Environments.

2.3 Cognitive Maps

2.3.1 Definitions and History

Essentially, a cognitive map is a mental image of a place, a network of representations of both places and the relationships between them [53]. There are many definitions of cognitive maps and cognitive mapping (see [41]), all similar but with different emphases. The major reason for the differences is that this field of research is multi-disciplinary, and thus has a wide base of knowledge and viewpoints, but no strong united philosophical and theoretical base.

A distinction is generally made between the terms "cognitive map" and "cognitive mapping" [41]. Cognitive mapping is traditionally defined as a process composed as a series of psychological transformations by which someone acquires, stores, recalls and decodes information about the relative positions and attributes of the items in their everyday spatial environment (Downs & Stea, 1973a in [41]), and is thus the process by which cognitive maps are formed and used. A cognitive map is a network of representations coding both places and the sequential relations among them [52]. Tolman, who first used the term "cognitive map", said that there is a map like representation in the nervous system which is used to guide our everyday movements, and that is represented as a cartographic map which gains Euclidean properties with repeated exposure to the environment [41].

The idea behind cognitive maps was first developed by Edward Tolman in 1948, as an attempt to explain the learning behaviour of rats in mazes [76]. At the time, there were two main schools of animal psychologists. One group believed that the maze behaviour of rats was a simple case of stimulus-response connections. According to this "stimulus-response" school, learning is a matter of strengthening some stimulus-response connections, and weakening others. A rat in a maze is exposed to a range of external stimuli (e.g. sights, sounds, smells) as well as a range of internal stimuli coming the muscle and viscera of the rat itself. The rat's central nervous system was seen as a switchboard, with incoming external stimuli being linked to the outgoing messages to the muscles, with the linkages being created so as to traverse the maze most accurately. In this system, the learning process is simply the process of strengthening the linkages that lead to accurate turnings and other approved behaviours. Bad results (for example, taking a long time to get to the food in the maze, thus staying hungry for longer) will reduce the likelihood of that particular set of linkages being used again, while good results (e.g. getting to the food quickly and reducing hunger) will strengthen that particular set of connections. These reactions are known as negative and positive reinforcements, respectively.

The second group of animal psychologists, to which Tolman belonged, were known as "field theorists", and believed that in the course of learning, something like a map of the environment gets established in the rat's brain. They saw the brain more as a map room than a switchboard, where incoming impulses are elaborated into a "tentative, cognitive-like map of the environment" [76]. According to Tolman, these maps could vary between being narrow and strip-like, and being broad and comprehensive (nowadays referred to in terms of "route" knowledge and "survey" or "configuration" knowledge, respectively, in [86, 81, 17, 67, 52], amongst others). The narrower and more strip-like the map, the less successfully it can be used when the environment changes — for example, if a previously clear path is blocked off. In a strip-map the starting position and the goal are connected only by a single, relatively simple path, whereas in a comprehensive map a wider arc of the environment is represented. Narrow strip-maps can be induced by:

- a damaged brain;

- an inadequate array of environmental cues;

- an overdose of repetitions on the original path;

- the presence of too strongly motivational or too strongly frustrating conditions (when conditions are too strongly motivational, it becomes difficult to relearn the environment when the original path is no longer correct; when conditions are too frustrating, learning stops once a single usable route has been found).

Tolman's paper and ideas were largely ignored at the time, but were revived in the early 1970's (Allen, 1985 in [41]). The concept has expanded and current theories are that cognitive maps don't just denote relative positions. They also contain attributive values and meanings, and are not independent of role, function, and subjective impressions and feelings (Wood & Beck, 1989 in [41]). They encode resources, memories, aspirations, as well as factual information about layout (Spencer et al, 1989 in [41]).

2.3.2   Controversy over Cognitive Maps

There is a certain amount of controversy over the term "cognitive map". Some (e.g. Buttenfield, 1986, in [41]) say that a cognitive map is the external product of measurement, rather than an internal representation of an environment, while others (such as Gatrel, 1983, in [41] and Liben, 1981, in [41]) say that the term "cognitive map" refers to internal thinking processes, and that the external form which is evoked is a "spatial product".

There are four different ways of thinking about the term "cognitive map":

1. *Explicit statement: A cognitive map is a cartographic map.* O'Keefe & Nadel (1978, in [41]) hypothesized that the hippocampus is a map which is a 3D, Euclidean model of the world, and that it has rigid geometrical properties. This is based on rat studies, where the position of the rat in a maze correlated with neural activity in specific parts of the hippocampus [44]. The main arguments against this approach is that there are too few neurons to cover all experiences, and adding new data would require the rebuilding of each neuron's identity.

2. *Analogy: A cognitive map is like a cartographic map.* In this approach, there is no physical region in the brain that stores a cognitive map. The mind's eye inspects the cognitive map in the same way that the physical eye inspects a graphical map, and there is a correspondence between input and output behaviours of the storage and retrieval functions of the two representations. Effectively, the term map denotes a functional analogue, and Downs and Stea (1973, in [41]) argue that it is an analogy to be used, not believed.

3. *Metaphor: A cognitive map works like a cartographic map.* The reasoning behind this approach is that we behave as if we have a map in our head (Kaplan, 1973a and Graham, 1976, both in [41], and thus it is reasonable to act as if the cognitive map is a cartographic map. The problem with this approach is that it is very easy to slip into believing that there really is a map physically in the brain (Downs, 1976 in [41]), but this metaphor is useful because it makes relationships explicit and provides guidance for using imagery to remember abstract conceptual relationships that are not easily imaged ([83]). The main argument against this approach is that non-Euclidean properties have been discovered in external representations (Kuipers 1983 in [41]).

4. *Hypothetical Construct: a convenient fiction.* In this approach, the term "map" does not have a literal meaning, and "cognitive map" refers rather to covert, non-observed processes and organizations of elements of knowledge. This explanation was inspired by the finding of spatial products with non-Euclidean properties, such as intravisity (A is seen as further away than B, B is seen as further away than C, but C is seen as further away than A), and non-communicativity (the distance from A to B is not seen as the same distance as from B to A) ([77], Baird *et al*, 1982, in [41]). In this approach, whether the internal entities are like or work like a map is immaterial [41].

This last approach seems the most practical for those who are interested in using cognitive maps to study the amount of information that an individual has about an environment, as in this research, rather than in the actual process of cognitive mapping. While it is important to the theorists to know the true internal form of cognitive maps, it is possible to

use them in research regardless of whether the cognitive map is a real map or simply works like a map.

2.3.3 Forming a Cognitive Map

Cognitive maps are formed during purposeful activity (Spencer *et al*, 1989 in [41]), and the process involves integration of images, information and attitudes about the environment (Spencer & Blades, 1986 in [41]). In large-scale environments, the space is outside immediate perceptual field of user, and so cognitive activity is needed to build a cognitive map [7]. Gärling, Böök, and Lindberg [26] describe the five stages of forming a cognitive map:

1. Information is received by the senses;

2. The information is stored in the sensory registers;

3. The information is transformed or re-coded;

4. The transformed information is stored in short-term memory;

5. The information in short-term memory is stored in long-term memory.

Certain factors are thought to aid in the cognitive mapping process, such as frequency of access to the environment, the distinctiveness of the environment (either visually or through background knowledge of the place), and the frequency with which certain landmarks are thought about and used in planning of routes. Previous knowledge of environments of the same type can also aid in cognitive mapping, even if one has never been in that particular environment before [40]. When it is particularly important for visitors to build up an accurate cognitive map of an environment, as in the spatial learning method, these factors must be borne in mind when designing both the environment and the visitor's introduction to it.

Research suggests that once the basic structure of the cognitive map has been formed, further familiarity with the environment expands the extents of the cognitive map, adds detail, and makes the cognitive map more complex, but does not alter the basic structure [23]. In terms of the spatial learning method, it is thus important to make sure that the visitor's initial cognitive map is correct.

Physical Cues and Cognitive Mapping — Implications for Virtual Environments

Research conducted by Held & Rekosh in 1964 (reported in [73]) indicates the importance of motor experience and sensor motor interaction to the perception of the environment. The less motor-environment interaction available, the less accurate the perception of the environment. Stea [73] comments that these findings may have implications for transport methods used to explore a new environment. Passive systems, such as busses, may lead to cognitive maps that have less detail and are less accurate than active systems, such as walking or driving.

This idea has been confirmed by Holahan [33], who reports that drivers produce more accurate sketch maps, and Cohen & Cohen [16] who report that those who use mainly public transport tend to have the most inaccurate maps. Appleyard [1] also found that those who use public transport drew the lowest quality sketch maps, and those who drive

a car produced the highest quality sketch maps. Carr & Schissler [11], however, found no difference in cognitive mapping between drivers, passengers, and commuters

This has deep implications for navigation and cognitive mapping in Virtual Reality, as many VR systems do not support any motor-environment interaction beyond that afforded by pressing a key or holding down a button.

While little research has been done into the effect of proprioceptive and motor cues on cognitive mapping in Virtual Environments, several studies have been conducted into their effect on navigation, and since the cognitive map underlies and supports navigation and wayfinding [79], these studies indirectly reflect on cognitive mapping as well.

Chance *et al* [12] found that direction estimates were more accurate when participants walked naturally than when they used a joystick for movement, regardless of whether turning was a tracked, natural movement or joystick controlled. Iwata & Yoshida [35] quote two studies that suggest that proprioceptive feedback is important in navigation (Ware & Slipp, 1991; Witmer, Bailey, Knerr, & Parsons, 1996), and describe a study by Bakker, Werkhoven & Passnier (1998) in which it was found that the highest accuracy of estimation of rotation occurred when participants actually used their legs to turn themselves around. Their own study found that distance estimations were more accurate when using a treadmill than a joystick for movement [35].

However, a study by Witmer & Kline in 1998 [86] found no difference in distance estimation between using a treadmill, joystick or teleports for movement. Similarly, Koh *et al* found no differences in bearing and range estimations between the real world, desktop VR with keyboard-based movement, and immersive VR conditions with tracked movement. Waller [80] mentions that few studies have found differences in spatial perception between HMD and desktop displays, and himself found that display type (immersive versus desktop) had little effect on distance estimation. This finding is supported by Ruddle *et al* (Ruddle (1996) in [64]), who found that no significant route-finding differences were found between desktop and immersive displays, and although the data were not entirely conclusive, Ruddle *et al* found in a later study on a seascape VE [66] that direction estimates were not affected by the display type, although relative straight-line distance estimates were more accurate when using the HMD than a desktop display.

In summary, then, while the theory suggests that navigation and wayfinding performance should differ between the real world (which has all the physical cues available), immersive VR (which has some physical cues) and desktop VR (which has no physical cues), experimental results are mixed. The first study described in this dissertation investigates whether participants to a VE can build up a cognitive map of the environment (see Chapter 5).

2.3.4   Using a Cognitive Map

In wayfinding and navigation, one generally wants to travel from a "start" position to a "goal" position. If both these positions can be represented within the cognitive map, then the wayfinding process becomes a process of mentally searching the cognitive map for a path connecting the two positions. If there is no direct connection, then subgoal positions are used. These subgoals may become more refined, with intervening subgoals being added between the original set of subgoals [40]. Routes are often not optimal in terms of distance, and are a sequence of approximations in the right direction that is guided by these subgoals [23].

This process does not have to be a conscious one, although it can be [41]. Essentially, having a good cognitive map of an environment allows you to travel the route in your head, planning your trip, recognizing places as you get there, and predict what comes next on the route.

When there are choices along a route, having a good cognitive map of the place is more useful than having a learned route [52]. This is similar to Tolman's concept of narrow strip maps and broad comprehensive maps [76], and more recently, route and survey knowledge. If only a single route to the destination is known, the narrower and more strip-like the map, and thus the less successfully it can be used when the environment changes — for example, if a previously clear path is blocked off.

Having a broad cognitive map of the environment that is being explored is particularly important in the spatial learning method, as this implies that the visitor knows more of the interrelationships between the data points.

2.3.5   Distortions in Cognitive Maps

Cognitive maps exist in psychological space, as the places that they represent exist in physical space. The mental image reflects some of the metric features of the physical space, but distances, locations, and other geometrical properties of the psychological space may not have a one-to-one correspondence with the physical space that is being represented. The symmetry and reflexivity axioms of Euclidean space may not hold, and directions may not be tied into any comprehensible coordinate system. Any correspondence with the physical space that does exist may be purely topological [52].

Distortions are not due to retrieval failures, but to distorted representations. They are a result of faulty perception or of errors in the encoding process [13]. Holahan & Dobrowolny [34] suggest that errors and distortions are not random, but are related the behaviour of the individual in that particular environment. Tversky [77] denies that distortions are failures in processing so much as they are normal consequences of the heuristics used when processing environmental data. Distortions may thus occur during storage or during recall.

There are many types of distortions that may occur in cognitive maps, and many reasons for each type of distortion. When designing a VE for use in the spatial learning method, it is important to be aware of the various types of distortions that may occur, as well as the reasons for them. This not only allows the VE designer to reduce the number of distortions, but also to make use of certain distortions (for example, roads with lots of angles are generally perceived as being longer than straight roads of equal length (Sadalla & Magel, 1980, in [4], and [9]). Two objects can thus be placed close together, but joined via a winding route to give the impression of a greater distance.

Some types of distortions are discussed below, but of necessity these discussion are brief overviews only.

Distances

By far the most common type of distortion, distance distortion has also the most research dedicated to it.

Cognitive distance is a factor in three types of spatial decisions: it affects the decision to stay or go, it affects the decision of where to go, and it affects the decision of which

route to take [10]. Distance is judged not just in terms of physical distance, but also in terms of biological criteria (e.g. whether or not a destination is within walking distance) [5]. Cognitive distance differs from perceived distance in that perceived distances are those which can be seen and judged in one glance, while cognitive distances are those which need to be thought about in the abstract, without actually being seen [10]. Holahan [33] defines perceived distance as the distance between you and a visible object, and cognitive distance as the distance between you and a non-visible object.

Cognitive distance is affected by direction: the length of journeys away from the center of town are overestimated (Lee, 1970, in [10]), although some studies have shown the opposite effect (Golledge *et al*, 1969, in [10]. Routes which travel along a slope, whether up or down it, are seen overestimated in length [4].

Cognitive distance is also affected by the attractiveness of the destination; the more desirable it is to reach the goal, the shorter the distance appears [10]. Distance judgements also improve with emotional investment [33].

The accuracy of perceived distance is also important. If the distance is perceived inaccurately, then it will be remembered inaccurately and the cognitive map incorporating it will be inaccurate. The characteristics of the path itself affects the perception of distance. Roads with lots of angles are generally perceived as being longer than straight roads of equal length (Sadalla & Magel, 1980, in [4] and [9]), and routes containing more information and landmarks are seen as longer than more impoverished routes (Milgrim, 1973, in [4].

Equal distances are judged as shorter when they are within the same region than when they are in different regions, such as neighborhoods, or regions defined by clusters of landmarks (Hirtle & Jonides, 1985, in [4]). Straight line distances tend to be overestimated between neighborhoods, and when neighborhoods are separated by a physical barrier such as a railroad or river [13].

Golbeck [28] states that distance distortions vary with age (younger children overestimate distances) and with the actual distance between the two points under consideration — distances of around 3 feet tend to be overestimated, while distances of around 5 feet tend to be underestimated.

Distance distortions are tempered by familiarity with the environment, as the more familiar an individual is with a particular environment, the more accurate their distance estimations are [10, 28].

Incompleteness

Bechtel [4] notes that the most common distortion in cognitive maps is that of leaving something out. Incompleteness occurs when an environmental feature or object that exists in the physical environment is not included in an individual's cognitive map [33]. This is a fairly straightforward distortion, as objects in the environment that are not noticed cannot be included in a cognitive map [4].

Augmentation

Augmentation occurs when incorrect items are added to a cognitive map. This can occur when two places become confused and objects from one place are remembered as being from the other [4], or when knowledge of other, similar types of environments indicates that a

feature should be present, even though it was not present in that particular environment [33].

Size

Size distortions are also very common [4]. Features that are better known, and that are more familiar, will tend be remembered as bigger than features that are unfamiliar. The size of area is also seen as bigger the more often it is used, especially when it is used for social purposes or for communication [52].

Rectilinear Normalization

Rectilinear normalization refers to the tendency to impose a North-South-East-West grid on environments. Intersections that do not meet at 90^o tend to be forced into a rectilinear N-S-E-W grid [84], even when the true direction of intersection is as large as 45^o [13]. The same tendency causes buildings to be recalled as clusters oriented to the cardinal directions [84].

Displacement

Related to distance distortions, displacement refers to the tendency to remember neighborhoods, buildings and landmarks as being closer to environmental features than they actually are. Displacements are generally towards any outstanding feature or landmark, such as a river, railroad, or distinctive part of town [13].

Alignment and Rotation

Rotation and alignment effects are a result of the heuristics used to encode location when it is difficult to remember the exact placement and orientation of object. In these cases, the object is remembered in terms of the surroundings rather than in and of itself. This heuristic may be applied on storage of the image of the environment, or during retrieval (or reconstruction) of the image, for example during questioning [77].

The rotation effect is caused by objects or environments inducing their own natural axis and imposing this induced axis on the memory of its position and orientation. There are 3 types of axes which can be induced by visual images of an environment [77]:

- axes of symmetry/balance (e.g. a road bisecting a region);

- main-line axis (e.g. a road and a river running in the same direction);

- landmark axes (e.g. a road connecting two landmarks).

When the background (such as the office in which an object is placed) has its own natural axis, the two axes tend to be drawn together, and are remembered as being more aligned than they actually were (compare Figures 1 and 2). In addition, the natural axis of an are or object is often rotated in memory so as to align to either a horizontal or vertical axis (compare Figures 1 and 3), so for example rivers and roads are often remembered as running more directly north-south or east-west than they actually do [13, 77, 78].
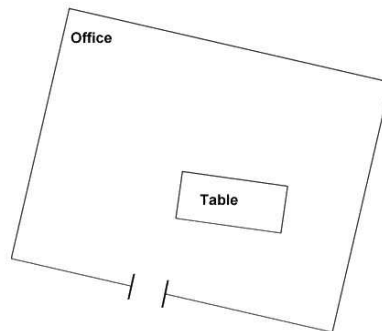
Figure 1: A table in an office. The office is rotated 13^o to the right; the table is rotated 8^o to the right.
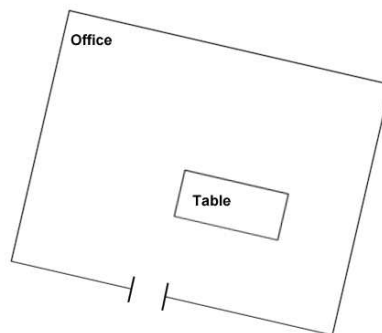
Figure 2: A remembered view of the office in Figure 1. The table has been rotated to align with the office.
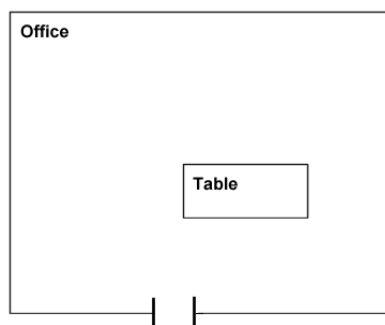
Figure 3: Another remembered view of the office in Figure 1. In this case, the office and table have been rotated to align with a horizontal and vertical axis.
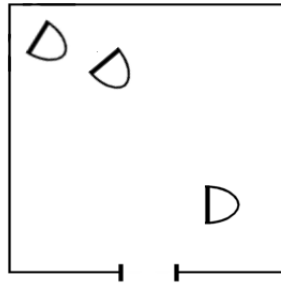
Figure 4: A scene in which the two objects near the top left corner will be grouped together in memory.
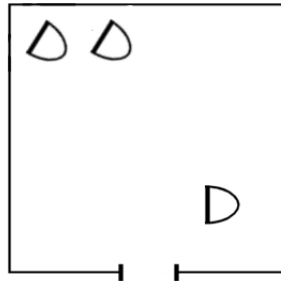
Figure 5: The two objects in Figure 4 have been lined up relative to each other, and have been given the same orientation. The object in the right hand bottom corner was not seen as part of the group.

The alignment effect occurs when two or more objects in a scene are positioned close together. Memories of the scene tend to group the objects together, to orient them towards each other, and to line them up relative to each other [13, 77, 78]. For example, the scene in Figure 4 may be remembered and sketched as in Figure 5, where the two objects close together are lined up together, and are both oriented in the same direction.

This effect also affects distance judgement. In a study by Coren and Girgus (1980, in [77]), participants were asked to judge distances between pairs of dots. When the two dots in the pair were perceived as part of the same perceptual group, the distance between them was estimated as being smaller than when the two dots were perceived as being part of separate groups.

Hierarchical Errors

Hierarchical errors are those caused by the nature of the storage of a cognitive map. McNamara [50] states that cognitive maps are stores as a hierarchy, which contains nested levels of detail. Information is organized under the relation of containment, and such a structure can be expressed as a tree in which terminal nodes correspond to objects, and nonterminal nodes correspond to various clusters of objects. For example, the city "Seattle" would be a terminal, while the state "Washington", which is a collection of cities, would be a non-terminal. The errors that people make when judging spatial relations lends some

support to the hierarchical representation theory. For example, Seattle is often judged to be south-west of Montreal, instead of its actual relation of north-west. Presumably, this kind of judgement is based on the spatial relationships of the larger, containing area — the US is south of Canada, so Seattle must be south of Montreal. This type of effect is also reported by Tversky [77] and Chase [13], who call them "part-whole" errors.

Straightening Effect

Straightening affects both the shape of features and locations. Edges (such as rivers, roads, or town outlines) which have many turns and bends are remembered as being less complex than they actually are, with only the general outline being remembered. Smaller turns and bends are forgotten. For example, the USA/Canada border is remembered as being straighter than it actually is [77].

Distortions in Virtual Environments

Because there is some uncertainty as to whether the cognitive mapping process is the same in a virtual environment as in a real environment, it is important to note that distortions also occur in cognitive maps of virtual environments. Some of these distortions are the same as those that occur in the real world, while some seem to be specific to VEs. Most studies are concerned with distance distortions and direction distortions, and how these differ between the different types of virtual environments (Section 2.1). Many of the results are contradictory, and while it is not possible to discuss all these studies and findings, a representative sample is presented.

Witmer & Kline [86] compared distance judgements between participants using a treadmill to control movement in a VE and those using a joystick. They did not find any significant differences in distance estimation.

Iwata & Yoshida [35] also found no difference in distance estimation of straight paths when using a two-dimensional treadmill as opposed to a joystick, but did find significant differences when the route included bends and turns. In both cases, they found that females were more accurate at distance estimation than males.

Koh *et al* [43] found that bearing and range estimations were equally low in real world, desktop VR, and immersive VR conditions. Henry (1992, in [43]) found that judgements of room size were smaller in simulated environments than in real ones, and that angle estimations had a wider variance in the simulated environment.

Waller [81] found that distances between objects were generally underestimated when using either an head-tracked HMD or a mouse-controlled desktop system. Providing feedback on the distance estimates significantly improved their accuracy, and a geometric Field-of-View (the visual angle depicted in the virtual scene) of between 50^o and 80^o resulted in the greatest accuracy.

Colle & Reid [17] found that after using a desktop and keyboard based system, participants estimated the angles between two objects more accurately if the objects had been in the same virtual room rather than in two different rooms.

Witmer & Kline [86] found that with regard to perceived distance between the participant and an object, both VR and real world estimates were low, but that VR estimates were less accurate than the real world estimates. They also found that the size of the virtual object affected the accuracy of the estimates, but floor textures and patterns did not. The

error in the distance estimates was also affected by the actual distance between the participant and the object — estimates were less accurate at shorter distances. With regards to traversed distance, Witmer & Kline found that VR estimates were again lower, but were more accurate than those of perceived distance. Distance cues (one beep per 10 feet of movement) improved distance judgements, but using a treadmill rather than a joystick or teleportation did not affect the accuracy of traversed distance judgements.

Ruddle *et al* [66] found that participants who used a HMD estimated straight-line distances more accurately than those who used a desktop system, but also found that there was no difference in direction estimation accuracy.

In a different study, Ruddle *et al* [64] found that participants who followed a route with more turns (although no choice points) made less accurate direction estimations, but that FOV did not affect direction estimation.

Chance *et al* [12] report that direction estimates were more accurate when participants could control their virtual movement by walking normally (with their body position and heading being tracked to update their virtual position and orientation) than when they used a joystick to control movement and turning.

Witmer *et al* [87] found that while both VE and real world distances were underestimated when walking blind to a previously seen target, errors in the VE condition were twice the magnitude of those in the real world.

2.3.6 Measuring Cognitive Maps

As a cognitive map is an internal construct, it has to be transformed into an external product before it can be assessed and measured [52]. There are many different ways of accessing an individual's cognitive map, each with its own difficulties, advantages, and disadvantages. Some of the more common methods are discussed briefly, together with their strengths and weaknesses.

Direction Judgement

One method of measuring spatial knowledge is by using direction judgements. This usually involves asking participants to point in the direction of a certain landmark or environmental feature. This method has been used in Virtual Environments by Ruddle *et al* [64] amongst many others. Asking participants to point to the same landmark from several different locations allows for inference about the location of the landmark through triangulation [55]. Direction judgements only access a single aspect of an individual's cognitive map, however, and should be used in conjunction with other methods.

Distance Judgement

Distance judgements are another way of measuring spatial knowledge that is more popular than direction judgements. Participants can be asked to estimate distances in absolute or relative terms, or in terms of travel time [33, 55]. The method of direct magnitude estimation may also be used, whereby a standard distance is given for the distance between a pair of objects, and participants are required to give the distances between other pairs of objects in terms of this distance [10]. Again, distance judgements only access one aspect of an individual's cognitive map, and should be combined with other methods.

Multi-dimensional Scaling

Multi-dimensional scaling essentially involves many sets of paired-comparison distance judgements. Participants are asked to scale the distance between all possible pairs of places in an environment. A multi-dimensional scaling algorithm is then run on the output, and a configuration of the places is obtained. This configuration represents the structure inferred by the pairwise judgement procedure [29]. A map can be produced from this set of configuration data, and this constructed map can then be compared to a cartographic map [55]. For an environment with a significant number of places, however, this can be a very long procedure for the participants. Newcombe [55] warns that this type of task requires considerable information-processing capability, and requires skills that are not well-practiced and may be beyond the capabilities of some subjects.

Verbal Descriptions

Asking participants to give a verbal description of the environment is another way of accessing their spatial knowledge [33]. The description could be written down or spoken aloud, and participants could be asked to describe a route (a technique used by Lynch [47]), or they could be asked to give a general description of the environment [33].

This type of data can be very difficult to analyze, as there can be tremendous variability in the content of individual responses [29].

Reconstructions

Yet another method of accessing an individual's cognitive map is to ask them to create a model of the environment under question using blocks provided by the researcher. This has the advantage of producing spatial representations directly, with points related to each other simultaneously in two or three dimensions [55]. Howard, Chase & Rothman(1973, in [55] asked participants to place scale models of eight buildings, with the distance between two of them provided, and found that the judged distances correlated very highly with the actual distances.

While modelling may be a suitable method for small environments, larger environments with non-trivial numbers of landmarks may be more difficult to represent. It may become more difficult to place blocks correctly in relation to the blocks already placed and to show the connections between these blocks, and if the blocks are not labelled or otherwise visually identifiable it may be difficult for both the participant and the researcher to keep track of which landmark each building represents.

Accuracy of Wayfinding

It is also possible to measure an individual's knowledge about the environment by setting them several wayfinding tasks and analyzing their performance [88], and this is one of the methods used in this research (see Section 4.7.5).

There are various ways to analyze wayfinding performance. Koh *et al* [43] give several measures used to analyze wayfinding performance: the number of wrong turns made, the route traversal time, the number of misidentifications of the destination, and the distance traveled. Rovine [61] rated wayfinding performance on three criteria:

1. whether or not the most direct route taken;

2. the distance traveled to reach the destination;

3. the number of turns taken to reach the building;

Waller *et al* measured wayfinding ability by means of a blinded walk, and then counted the number of wrong turns taken, and measured the total time taken to cover the route.

Measuring an individual's wayfinding performance appears to be one of the least artificial methods of accessing cognitive maps, as this is the type of activity which a cognitive map is used for in real life situations.

Sketch Maps

The second method of measuring cognitive mapping used in this research is that of analyzing sketch maps. This is the most commonly used method, but also the most controversial, and consists of giving the participant a sheet of blank paper and asking them to draw a map of the environment under question. Generally, further instructions are not given, so as to allow the participant to draw their own conclusions about what is required [42]. Sometimes, however, participants are told to draw the map as if to help guide someone who is new to the environment [34], which is the approach taken in this research (see Section 4.7.5).

Problems with Using Sketch Maps   Wood & Beck [90] point out some problems with using sketch maps in research. They caution that some participants may never have drawn a map before, and may suffer from a form of graphophobia. Holahan [33] warns that the reliability of many of the measures used to analyze sketch maps has not been assessed, and also points out that the drawing ability of the participants may affect the validity of the results. This last is a common criticism of sketch maps, being voiced by, in addition, Golledge [29], Siegel & Cousins [69], Darken *et al* [22] and Newcombe [55]. Newcombe also points out that freely drawn maps are difficult to classify and study quantitatively. Downs [24] agrees that scoring, coding and analysis of sketch maps are idiosyncratic and problematic, but does not deny their usefulness as data, and concludes that while the concept of a map is useful for studying an individual's representation of a space, the production of a map is a problematical technique.

In Defense of Using Sketch Maps   Despite the criticisms leveled by Wood & Beck [90], they agree that there are certain advantages to using sketch maps. The mapper can communicate large amounts of information to the researcher, and can also easily communicate the interrelatedness of environmental features — it is easier to draw three items than to explain their relative positions. Holahan agrees that sketch maps are easy and efficient and provide data in a vivid and qualitatively rich format, and warns that while other techniques (such as verbal descriptions and distance judgements) are more reliable, they are also more artificial and thus less valid [33]. Rovine [61] comments that many of the techniques used as a replacement for sketch maps require complex and tedious processes for data collection and data analysis, and quotes three studies (Howard, Chase & Rotham, 1973, and Rothwell, 1974 and 1976) that suggest that sketch maps are relatively reliable. These studies found that for adults, graphic skill accounted for less than 2% of the variance in sketch
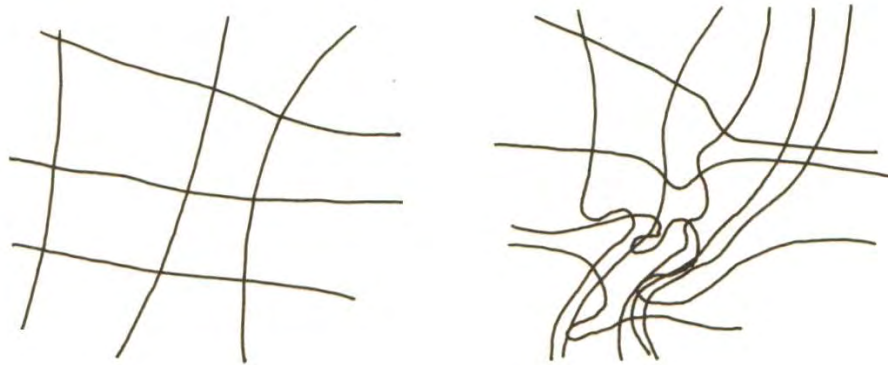
Figure 6: Two examples of using the grid transformation to analyze sketch maps. The sketch map which was the basis for the grid transformation on the left was quite accurate, as the lines are only slightly distorted. The sketch map on which the grid transformation on the right was based, however, was quite badly distorted. These figures are taken from a study by Beck & Wood [5].

map accuracy. Newcombe [55] cites two more advantages of sketch maps, in that they allow participants to produce spatial representations as spatial products directly, and that they allow points to be related to each other simultaneously in 2D or 3D. Newcombe also quotes two studies that show that sketch maps are of higher quality than those produced by multi-dimensional scaling techniques, and concludes that while sketch maps do depend on drawing ability, they are nevertheless and easy-to-understand everyday task that provides context for distance judgements, and they provide a natural means of representing orientation and distance information simultaneously.

As this research uses adult participants, the major problem with using sketch maps is eliminated, especially in terms of the findings of Howard, Chase & Rotham, and Rothwell.

Analyzing Sketch Maps There are many different ways of analyzing sketch maps, the majority of which focus on understanding the cognitive mapping process or finding out the cognitive mapping skill level of the participant (such as the map classification schemes used by Appleyard [1]. These types of measures are not appropriate when the aim of the sketch map is to assess the environmental knowledge of the participant (as is the case in this dissertation).

A measure that does assess the accuracy of sketch maps in order to determine how well individuals know the locations of objects in the environment is used by Beck & Wood [5]. Each individual sketch map is transformed onto a regular geometric grid. Displacement of points and lines from their true positions show up as wavy lines on the grid, as the grid lines have to bend out of their way to pass through all the points that should appear on that line on a standard map [5]. An example of this method is shown in Figure 6. While this method discovers systematic distortions in a particularly vivid way, it does not produce quantitative data and thus the results can be almost as difficult to compare and assess as the original sketch maps.

Appleyard [1] analyzed the accuracy of sketch maps by determining the number of

structural errors made in terms of the zones of the city under study. He categorized errors into five different types:

1. faults (a break between zones);

2. bends (the relationships between zones were spatially distorted);

3. mislocations (zones placed in the wrong position on the map);

4. zone reversals (the position of two zones were reversed);

5. distortions (the road system within a zone was distorted).

Rovine [61] derived three types of measures from sketch maps: frequencies of landmarks, path segments, and intersections; complexity of the sketch map (sequential or spatial, as described in [1]); and accuracy of the sketch map in terms of topological accuracy of placement of target buildings. Buildings were judged to be accurately placed if the building was in the appropriate sequence with respect to the buildings before and after it on the sketch map, and if the path connecting the buildings accurately reflected the turns needing to be made when traveling from the preceding to the following building. This method of measuring local accuracy defines accuracy only in terms of the information included on an individual's map rather than the information available in the environment. This is similar to the approach taken in this research, described in Section 4.7.5.

2.3.7 Children and Cognitive Mapping

While this research uses only adult participants, it is important to realize that children differ from adults in terms of cognitive mapping abilities, and thus that the results of this research cannot automatically be carried over to children. Not only is their perception and processing of the environment different to adults, so is their understanding of and responses to methods of eliciting cognitive maps.

A child's understanding of the spatial layout of an environment develops through three stages of cognitive reference systems [51]. The first stage is that of an undifferentiated egocentric reference system, where the only elements included in a representation of the environment are those of particular personal significance. The elements are mostly all of the same type, and there is no differentiation between points of view. Elements and landmarks are not organized in terms of a spatial whole, and positions are mainly sequential and topological, without concern for left-right, before-behind, or distance relationships [51]. This stage lasts until about six years of age, and during this stage children locate objects in the environment relative to their own body [13].

The second stage is that of differentiated and partially coordinated subgroups based on fixed references. At this stage, the child organizes the environment on the basis of groups of elements, where the groups are not systematically related to each other, but the positions of elements within each group are correct. Each group shows a particular point of view or journey [51]. This stage lasts between the ages of about 7 to 9, and during this stage children use a fixed coordinate system in which objects and their own body are oriented relative to landmarks and other fixed points in the environment [13].

The third stage is when the representation is operationally coordinated and hierarchically integrated. At this level, an overall coordinate system exists, and left-right, before-behind and relative distance relationships are correct. New subgroups can be added to the model, with the representation being changed to incorporate the new group so that it is correctly related to the existing groups [51]. This adult-like representation occurs at about 11 or 12 years of age, and children are now able to orient themselves and objects in the environment relative to an abstract coordinate system such as the cardinal directions [13].

Elicitation of the child's model of the environment must match their capabilities [29]. Young children have limited vocabularies and limited graphical skills, and this must be taken into account when asking them to draw a sketch map or describe a route [42]. Asking children to arrange blocks to create a model of the environment has been used with some success, and this may be a better method of accessing a child's cognitive map than using sketch maps, provided that the model is smaller in scale than the original environment (i.e. miniaturization is required rather than expansion) [55].

The way in which children use cognitive maps also differ from adults. In a study on navigational strategies, Volbracht *et al* found that children remembered routes on the basis of specific landmarks, particularly those which were strongly linked to their everyday lives (for example, a sweet shop). Children also had difficulties in gaining an overall picture of the environment, and although a 2D map can assist in this, children of elementary school age may have equal difficulties in using a map [79].

2.4   Navigation and Wayfinding

2.4.1   Definitions and Basic Concepts

There is an important distinction between navigation and wayfinding. Navigation is the physical movement through an environment, as well as a cognitive element known as wayfinding [22]. It consists of extracting information about the environment, forming mental representations, and using these representations for route planning and moving about in the environment, whether real or virtual [18]. Navigation also involves updating your position and orientation during travel with respect to your destination, and re-orienting yourself and re-establishing travel towards your destination on becoming lost [45].

Wayfinding is the ability to learn and remember a route through an environment (Blades, 1991 in [41]), as well as the ability to find a particular location, and recognize it when you reach it [33]. It requires cognitive mapping [4, 22], route planning, and distance estimation [22], and consists of a decision making stage, and decision execution stage [4].

2.4.2   Landmark, Route and Survey Knowledge

There are three basic types of knowledge that one can have of an environment. *Landmark* knowledge is the most basic knowledge one can have about an environment. At this stage, knowledge consists of a disconnected set of landmarks [81]. An individual who has landmark knowledge of a place orientates themselves exclusively via highly visible landmarks such as buildings and statues, but cannot navigate from one place to another using the landmarks as guides at decision points [84]. Landmark knowledge can be used to recognize the location in which one may find oneself [14]. Landmark knowledge can be acquired either

by directly viewing objects in the environment, or by viewing indirect representations of the environment such as photographs [14].

Route knowledge, which is also known as procedure knowledge, is a slightly more advanced form of knowledge. It is normally gained during personal exploration of an environment [67], is acquired by associating navigational actions to landmarks ("turn until you're facing the statue") [33, 84], and may require sensorimotor experience of the routes [15]. Route knowledge is characterized by sequentially organized information about particular routes [64]. A person who has only route knowledge of a place can't take advantage of shortcuts [33], and won't know how to navigate around blockages in their route. When navigating via route knowledge, people tend to visualize the route in the mind's eye from a first person perspective [14].

Someone who has *survey* or *configurational* knowledge of a place, by contrast, has an overall picture of the place and know where things are in relation to each other. They can infer routes [14, 67] and see the place from a third person, all-encompassing perspective [14]. Survey knowledge represents the configuration relations among locations and routes in an environment, encoding the topographical properties of the place [14], although it may also require knowledge of the distances between locations [86]. Survey knowledge is gained by multiple explorations of an environment using multiple routes [67], or by the study of maps or other media [14]. It appears as though it is quite likely that survey knowledge will eventually develop from route knowledge on repeated exposure to the environment (Moeser, 1998, in [67], [84]). Survey knowledge can take a long time to develop, however, as individuals can be in a real environment for more than a year without gaining survey knowledge of it (Moeser, 1998, in [81], and if the tasks performed in the environment do not require survey knowledge it may never develop [81]. Alternatively, some survey knowledge may be available even before route knowledge, after only brief encounters with the environment [17].

Survey knowledge is thus the most complete knowledge of the environment, and this type of knowledge corresponds to Tolman's concept of a comprehensive cognitive map [76] (see Section 2.3). Route knowledge, on the other hand, corresponds to Tolman's concept of narrow strip-maps. For the type of application described in this dissertation, route knowledge is not sufficient — the user must be able to acquire survey knowledge of the environment.

2.4.3 Using Real World Research in Virtual Environments

There is some amount of uncertainty as to whether research into navigation and wayfinding in the real world can be applied to virtual environment, due in part to the issues discussed in Sections 2.1 and 2.3.3.

Satalich [67] concludes that are differences between real and virtual environments that affect performance in navigation and wayfinding tasks, while Koh *et al* [43] found that some of the same distortions occur in real and virtual environments (such underestimating ranges), and quotes conclusions by Darken and Sibert to the effect that principles developed in connection with real world navigation could be applied to VE navigation. Waller [80] quotes several studies that have found that VEs can allow users to form mental representations of large real-world spaces (e.g. Ruddle, Payne & Jones (1997), Waller, Hunt & Knapp (1998) [81], and Witmer, Bailey, Knerr, & Parsons (1996)), but also warns that

there is evidence that VEs distort perception of distances (Witmer & Kline (1998) [86]) and angles(Ellis, Smith, Grunwald & McGreevy, 1991). Vinson [78] gives several reasons for assuming that navigational behaviour is the same whether it takes place in a real or virtual environment.

2.4.4 Issues Specific to Virtual Environments

Navigation is currently the most significant usability issue in Virtual Environments [14]. As VEs become larger and more complex, it becomes easier for visitors to become disoriented [3] and become lost, and for most users becoming lost is very disturbing [78]. In addition, it is difficult for visitors to discover everything of interest in a VE without exploring it [3], and this exploration requires accurate navigation.

Navigation is often difficult for users, however [78]. They normally spend relatively short periods of time in any particular VE (generally minutes or hours, rather the days or months used by cognitive mapping researchers), and frequently lack of knowledge of their position and orientation. Not knowing the structure and layout of the particular VE as well as a lack of familiarity with VEs in general also contribute to navigation difficulties [64], and it may not always be possible to structure the VE in such a way as to help with navigation, since the spatial structure may represent data [78]. Non-intuitive navigation methods such as navigating by looking or by mouse movement [81], joystick, or keyboard [86] also contribute to difficulties in navigation.

Virtual Reality systems currently do not provide many of the cues available in the real world, some of which are described in Section 2.3.3 and Section 2.1. The lack of such cues can negatively cognitive mapping and navigation (see Sections 2.3.3 and 2.3.5). The cues that are lacking can be divided into three basic types: environmental factors, physiological factors, and system factors.

Environmental Factors. Because of computational limitations, virtual environments generally contain less visual detail than real environments, which means that there are fewer landmarks to support navigation [64, 78] and fewer depth cues (such as occlusion and texture gradients) to help users perceive distances accurately [78] (Witmer & Kline [86] found that texture did not affect distance estimates, however).

Physiological Factors. Many physiological cues are not available when interacting with a VE. Locomotive and proprioceptive cues normally provided by walking and turning one's body or head are often absent [78], and lack of scene detail and texture as well as low display resolutions can diminish motion perspective cues, which result from changes in texture gradients and give information about the rate and direction of movement [86]. The vestibular cues indicating changes in body position and orientation and the proprioceptive cues which inform one of one's body actions (such as walking speed and stride length) are also absent, and the lack of these cues can negatively affect traversed distance judgements [86]. However, some display systems such as HMDs do provide kinesthetic and vestibular cues [66]. Section 2.3.3 describes some possible effects of the lack of physiological cues on cognitive mapping.

System Factors. Some VE display systems, such as desktop and HMD systems, do not provide much (or any) peripheral vision [66, 78]. However, displays such as a CAVE do allow for peripheral vision, and some researchers have attempted to simulate peripheral vision with peripheral monitors (Slater & Usoh, 1993 in [66]). Lack of peripheral vision can cause users thus sometimes travel past their targets and also don't notice as much of the environment. The reduced field-of-view (FOV) provided by some environments also causes users to travel past their targets, as the smaller the FOV the greater the angle that users must turn their heads in order to see what they are walking past [66, 86]. The geometric FOV (the visual angle depicted in the virtual scene) can be set to greater than the physical FOV of the viewing device in order to allow users to see more of the environment, but this distorts the view and objects appear to be smaller and further away than they actually are [66]. A geometrical FOV of between 50^o and 80^o results in the greatest distance estimation accuracy [80]. A restricted FOV also negatively affects the acquisition of survey knowledge [81], although previous studies conflict as to whether or not this reduces accuracy of learning a spatial layout [64], and FOV has no effect on distance estimation [64]. However, displays such as a CAVE allows a more natural FOV [66].

Lack of pictorial cues and poor binocular disparity cues also degrade VE performance in terms of distance estimations. Witmer *et al* [87] found that while both VE and real world distances were underestimated when walking blind to a previously seen target, errors in the VE condition were twice the magnitude of those in the real world, and they suggest that the cause for this may have been lack of pictorial cues and poor binocular disparity cues.

2.4.5 Design Guidelines for Virtual Environments

In order to design VEs that are easy to navigate, a VE designer must be aware of the many design guidelines that exist, both for real and virtual environments. While it is impractical to provide an exhaustive list here, those recommendations particularly relevant to Virtual Environment design are presented (more design guidelines can be found in [4], [23], [33], [40], [48], and [49], amongst others).

Satalich mentions several design recommendations made by Gärling *et al* (1986). Navigational cues such as room numbers and building names should be added to the structure, and the structure should be built to incorporate differentiation (the sections of a building should be visually distinguishable), complexity (right-angled intersections, few segmented passages), and visual access. Visual access allows areas of the structure to be seen from other areas, such as via windows, doorways, or open courtyards [67]. This can be a problem for VE designers, however, as many VR engines rely on there being only a small number of portals in order to reduce the number of polygons visible at any given moment, thus increasing rendering speed and frame-rate.

Golbeck recommends placing many landmarks in an environment. Landmarks facilitate encoding and retrieval of information about spatial locations (such as "near", "between") and can be used for maintaining orientation, direction of movement, and for marking spatial positions. Golbeck also recommends dividing the space into subspaces using paths and barriers, as this "chunking" of the environment reduced informational processing demands [28]. This coincides with the goals of the VE designer, since many VR engines give better performance in terms of frame rate if only small parts of the environment are visible at any one time.

Chen & Stanney [14] examined several findings relating to whether or not a regular street grid layout assists with wayfinding. These findings tended to conflict on this issue, and Chen & Stanney came to the conclusion that environmental layout may have more effect in VEs than in real environments as less spatial information is available in VEs.

Vinson [78] presents thirteen guidelines for VE designers (more detail on these guidelines can be found in [78]) :

1. The VE must contain several landmarks.

2. The VE should include landmarks from all five categories (paths, such as streets and railroads; edges, such as fences and rivers; districts, or neighborhoods; nodes, such as courtyards and large intersections; and traditional landmarks, such as statues and fountains).

3. Landmarks should be made distinctive by means of having significant height, a complex shape, a bright exterior, visible signage, or unique exterior colours or textures. Landmarks can also be made distinctive by surrounding them with landscaping or leaving them free standing in an otherwise crowded area.

4. Concrete, familiar objects should be used for landmarks, rather than abstract objects.

5. Landmarks should be visible at all navigable scales. If a visitor can zoom in or out to see the VE at different scales, at least some landmarks should be visible at all scales.

6. Landmarks must be easy to distinguish from nearby objects and other landmarks.

7. The sides of a landmark must differ from each other, to help visitors determine their orientation in respect to the landmark.

8. Landmarks may be clustered to add distinctiveness, to make sure that they are visible from all angles, and to appear different from different angles.

9. Landmarks should have a common characteristic which distinguishes them from data objects, so that they cannot be mistaken for data.

10. Landmarks should be placed on major paths and at path junctions.

11. Paths and edges should be arranged to form a grid, to reduce cognitive map distortions (see Section 2.3.5).

12. Landmark's main axes should be aligned with the axes of the path grid (again, to reduce cognitive map distortions).

13. Each landmark's main axes should be aligned with those of the other landmarks.

Ruddle *et al* [63] also recommend the use of familiar objects as landmarks rather than colored patterns (or no landmarks) in order to improve route-finding accuracy.

Koh *et al* report a finding by Aginsky, Harris, Rensink & Beusmans (1996) to the effect that environmental information in the vicinity of choice points (turns, etc.) is more likely to be retained than information at other points [43]. VE designers who want visitors to

remember particular items or environmental features should thus place them near these choice points.

Similarly, Ruddle *et al* recommend placing locations which VE designers want people to visit near landmarks or other major topological features, or displaying these locations on a map of the VE, as visitors tend to explore VEs using landmarks as anchors [65].

2.4.6 Maps and Other Aids

Many different types of navigational aids have been suggested, along with various types of maps. Not all of these have been found to be effective, however, and some have actually been shown to have an adverse effect on navigation, wayfinding, and cognitive mapping.

Kaplan [39] showed that information that was irrelevant to the environment being navigated was simply ignored and did not impact on cognitive mapping, while information that was relevant but misleading had a damaging effect. In addition, contour maps of a wooded natural environment generated more confidence in wayfinding ability than photorealistic maps, although this did not affect wayfinding performance.

Holahan [33] reports on a study by D. J. Bartram in 1980, in which participants were asked to solve various route-planning tasks involving bus routes. Participants were given one of four different aids: an alphabetical list of bus stops; a list of bus stops in sequential order; a conventional road map showing the bus stops; and a simplified schematic map showing the bus stops. The two map formats proved to be more efficient to use than the list formats, and the relative efficiency differences became more distinct as the complexity of the route-planning tasks increased.

Satalich [67] and Chen & Stanney [14] report on a study on by Streeter, Vitello and Wonsiewicz in 1985 which compared navigational aids for participants driving a car through a urban environment. Participants were given either a recorded narrative of the route, a customized route map of the area, both a recorded narrative and a customized route map, or a standard road map. Those using the standard road map drove further, took more turns than necessary, and took more time to complete the task than the other three groups. Participants using the customized map drove further than those using the recorded narrative and made more errors than any other group.

Satalich [67] and Chen & Stanney [14] also report on a study conducted by Butler, Acquino, Hissong and Scott (1993) on the effect of various aids on wayfinding. The study found that the use of You-Are-Here (YAH) maps significantly slowed wayfinding as opposed to no navigational aids, and that providing signs in the environment speeded wayfinding. Signs may be detrimental to the cognitive mapping process, however — Moeser (1988, in [67]) reports that nurses in a hospital under study did not have survey knowledge of the building even after three years of working there, and comments that several nurses reported that they use the signs in the building to find their way around and thus did not have to learn how to get to certain areas.

In a study conducted by Satalich [67], participants guided through a VE by a path superimposed on the floor performed better at a spatial orientation task if they were provided with a map of the environment. In other conditions, however, map use was detrimental to the learning of the environment.

Wickens [84] discusses the relative uses of fixed and rotating maps. Fixed maps are displayed with north at the top of the map (also known as north-up maps), and the user

must mentally rotate the map so that "up" corresponds to the direction of travel. Rotating maps, however, automatically change orientation depending on the orientation of the user, so that the direction of travel is always "up" on the map, which greatly reduces the cognitive load on the user. While it may appear that rotating maps are always the best solution, studies have shown that the relative benefits of each type of map may be task dependent. North-up maps are better suited to tasks which require locating landmarks on the map itself, as well as tasks that depend on learning the relative locations of environmental features. In addition, a fixed map supports planning and communication with other users who may be at a different location in a different orientation. One solution which may combine the most useful features of each type of map is to provide users with a north-up map, but to clearly indicate the forward field of view of the user. This solution has been experimentally verified, and provides performance as good as (and sometimes better than) a rotating map. This type of map is also recommended by Sayers *et al*, and has been experimentally verified [68].

Darken & Cevik [20] also performed a study into the uses of north-up (or fixed) and forward-up (or rotating) maps. Performance on search tasks was measured in terms of search time and number of errors, and results indicated that forward-up maps are more suitable for targeted search tasks (when the participant has seen the target before), while performance on naive search tasks (when the location of the object is completely unknown) is better when participants are provided with a north-up map.

Chase [13] provides several guidelines when providing You-Are-Here maps, based on a study by Levine in 1983. Firstly, the map should be aligned with the environment — if the map has north at the top, then the visitor should be facing north when facing the map. Secondly, prominent landmarks should be visually indicated on the map. Thirdly, the map itself should be placed close to a landmark to help visitors locate themselves on the map. Finally, the location of the map itself should be indicated on the map. Levine studied YAH maps in the New York city area, and found that many violate even the most basic orientation principles.

Supporting this orientation alignment principle, Koh *et al* [43] describe a study by May, Peruch & Savoyant (1995) in which it was shown that map misalignment (where "up" on the map did not correspond to "forward" in the environment) reduced both speed and accuracy of navigation in a Virtual Environment.

Chen & Stanney [14] present a taxonomy of navigation tools with five different categories:

1. tools that display the user's current position;

2. tools that display the user's current orientation;

3. tools that log the user's movements;

4. tools that demonstrate the surrounding environment;

5. guided navigation systems.

Tools in categories 1 and 2 aid in spatial orientation tasks, particularly when landmarks are not present. Tools in category 3 may be more useful for navigation analysis than for navigational aiding, although they can also be used to display landmarks that have already

been visited. Tools in category 4 assist with spatial orientation and also collect spatial information from the surrounding environment (such as radar maps or virtual binoculars). With category 5 tools, visitors do not need to make their own wayfinding plans, but these tools should only be used for applications that do not require acquisition of spatial knowledge [14].

Ruddle *et al* [64] found that providing a compass to visitors to a seascape VE was not sufficient to provide global orientation information, but surmised that the compass may have been used more effectively if participants had been trained in using a compass. However, as with the contour maps provided by Kaplan [39], having a compass gave participants more reassurance while navigating.

In a study by Darken & Sibert (1996, in [65]), no differences in search time were found between four tool conditions: no aids; a forward-up dynamic YAH map that did not show the search targets; a grid and five tall landmarks visible from any position in the VE; and a combined condition with the map, grid and landmarks.

In a later paper, Ruddle *et al* [65] discuss the relative advantages and disadvantages of local maps and global maps. The advantage of a global map is that the visitor can see their global position at all times, which is efficient for random searches, and may be efficient for targeted searches if visitors can remember the approximate position of the target. However, a global map cannot provide enough detail to show the position and features of small targets. With local maps it is more difficult for visitors to travel past the target without noticing it, but using a local map to learn the environment may require as much integration and unaided navigation. Using both types of maps would have the advantages of both, but may increase the attentional demands on the visitor.

In this study, Ruddle *et al* [65] investigated the effectiveness of local maps, global maps, and other aids in desktop VEs. Five display conditions were used: no aids; a global map; a local map; a global map and a local map; and numerical coordinates and a compass. All maps were dynamic, north-up YAH maps which indicated a participant's position, view direction, and movement direction. In terms of distance traveled, participants in the combined local and global map condition and those in the local map condition performed significantly better than participants in the other three conditions during a search task. During an informed search task (revisiting objects already found), participants in the combined local and global map condition traveled significantly less distance than other participants, and those in either the local or global map groups traveled less distance than those in the no aid or coordinate-compass groups. Orientation and direction estimates were more accurate for the combined local and global map group and the global map group than for the other three conditions. Overall, then, the combined local and global map was the most effective display condition. On the basis of the findings from this study, Ruddle *et al* make the following suggestions:

- coordinates and compasses, or similar tools such as virtual suns, are not suitable navigation aids unless visitors are trained to use them effectively;

- if there is space for only one map, the optimum scale for the map depends on the users knowledge of the VE — local maps may be more useful for those new to the VE, while global maps are more appropriate for those familiar with the VE.

Darken & Sibert [21] measured performance on wayfinding tasks in VEs which either

provided no assistance to participants, or assisted them by either overlaying a radial grid on the environment, providing them with a map, or providing both a map and a grid. Navigational performance was worse when no assistance was provided, and was best when using map principles. The grid condition, however, provided better directional information.

Witmer, Sadowski & Finkelstein [88] compared the effect of various navigation aids on the acquisition of configuration (or survey) knowledge. Participants were assigned to one of eight conditions:

- the standard environment with the normal first-person perspective and no orientation cues;

- the standard environment with the normal first-person perspective as well as the option to use an aerial perspective but no orientation cues;

- the standard environment with the normal first-person perspective with orientation cues (such as virtual flagpoles and arrows indicating current heading);

- the standard environment with the normal first-person perspective and the option to use an aerial perspective, with orientation cues (such as virtual flagpoles and arrows indicating current heading);

- the enhanced environment (additional objects and sounds) with the normal first-person perspective and no orientation cues;

- the enhanced environment with the normal first-person perspective as well as the option to use an aerial perspective but no orientation cues;

- the enhanced environment with the normal first-person perspective with orientation cues (such as virtual flagpoles and arrows indicating current heading);

- the enhanced environment with the normal first-person perspective and the option to use an aerial perspective, with orientation cues (such as virtual flagpoles and arrows indicating current heading).

Orientation cues had no significant effects on the acquisition of configurational knowledge, and neither the VE enhancements. Participants who used the aerial perspective view performed significantly better on configuration knowledge tests.

2.4.7   Measuring Navigation and Wayfinding Performance

There are three main methods by which navigation and wayfinding performance is assessed: direction judgement, distance judgement, and accuracy of wayfinding.

Direction Judgement

Direction judgments usually involve asking participants to point in the direction of a certain landmark or environmental feature. Pointing to the same landmark from several different locations allows for inference about the location of the landmark through triangulation [55].

Distance Judgement

Distance judgements are generally more popular than direction judgements. Participants can be asked to estimate distances in absolute or relative terms, or in terms of travel time [33, 55]. The method of direct magnitude estimation may also be used, whereby a standard distance is given for the distance between a pair of objects, and participants are required to give the distances between other pairs of objects in terms of this distance [10].

Accuracy of Wayfinding

It is also possible to measure wayfinding and navigation performance by setting them several wayfinding tasks and analyzing their performance [88]. As discussed in Section 2.3.6, Koh *et al* [43] give several measures used to analyze wayfinding performance: the number of wrong turns made, the route traversal time, the number of misidentifications of the destination, and the distance traveled. Rovine [61] rated wayfinding performance on three criteria:

1. whether or not the most direct route taken;

2. the distance traveled to reach the destination;

3. the number of turns taken to reach the building.

Waller *et al* measured wayfinding ability by means of a blinded walk, and then counted the number of wrong turns taken, and measured the total time taken to cover the route.

2.5   Spatial Abilities

2.5.1   Definitions and Basic Concepts

Bishop (1980, in [83]) describes two types of spatial skills:

- spatial visualization (such as rotating objects in space);

- spatial orientation (such as recognition of relationships between objects and one's own position, such as is needed for navigation).

Bishop then goes on to distinguish between low level spatial abilities and high level spatial abilities. Someone with low level spatial abilities can visualize two-dimensional shapes, but cannot perform mental transformations of them, while someone with high level spatial abilities can visualize and manipulate objects in three-dimensional configurations. Weatherford [82] also describes two types of spatial tasks: perspective taking, where the subject must mentally determine transformations if self-environment relationships that would occur as a result of an imagined movement of the self), and mental rotation, which is similar but based on the imagined movement of the environment. All four of these spatial tasks are needed for cognitive mapping and navigation, remembering and planning routes, and drawing sketch maps and making distance and direction estimations. In one study, Georger (1998, in [43]) showed that scores from the Guildford-Zimmerman Spatial Orientation Aptitude Survey (GZ) correlated with spatial tasks such as direction estimation and unplanned route selection. In a study by Waller *et al* [81], participants were shown maps purposing to

be of a VE that they had previously explored and were asked whether each map truly was of the VE or not. GZ scores did correlate with performance on this task, although they did not correlate with scores derived from a blind-folded walk in the real environment on which the VE was based.

2.5.2 Measuring Spatial Abilities

Everyday spatial activities have been shown to correlate with a typical objective test of three dimensional space visualization. In particular, understanding of maths and science, understanding of graphs and charts, skill at arranging objects, and skill at drawing, and experience with using hand tools were most strongly correlated to spatial visualization scores. Based on these results, Lunneborg & Lunneborg have developed a self-report questionnaire which measures everyday spatial activities (the Everyday Spatial Activities Test, or ESAT), and have found that it correlates with seven pre-college variables: high school English GPA (ENGPA) and mathematics GPA (MTGPA), Vocabulary and Quantitative Composite Scores (VC and QC), Spatial Ability score (SA), Mechanical Reasoning score (MR), and all-college first-year predicted GPA (PGPA). The ESAT consists of 20 items representing 4 different factors — use of hand tools, maths/science courses, arranging objects, and mechanical drawing. Each item asks the participant to rate how good they are at certain everyday activities, on a scale from 1 to 5 (1 being very bad, and 5 being very good). A total spatial abilities score is obtained by simply adding together the ratings for each of the 20 questions.

2.5.3 Factors affecting Spatial Abilities

Gender is a major factor affecting spatial abilities, both perceived and real. In a study by Lunneborg & Lunneborg [46] college men consistently overestimated their everyday spatial abilities as compared to college women. These higher self-ratings may be related to greater experience in maths and science, as practice and ability in everyday spatial abilities were equally strongly correlated for both men and women. In addition, objective spatial ability test performance was more strongly related to mathematics preparation for women than for men. Satalich [67] found gender differences in spatial abilities as measured by the Guildford-Zimmerman Spatial Orientation Aptitude Survey.

Mathematics abilities have been shown to correlate with spatial abilities (Smith, 1964 in [83]), and evidence suggests that this relationship holds for both low level and high level spatial abilities [83]. Both maths and science teachers tend to be more successful on the Differential Aptitude Test than those in the humanities or social sciences (Martin, 1967-68 in [83]).

Age also correlates with spatial abilities. Children not yet in the fifth grade find the perspective taking and mental rotation tasks described above difficult to perform, and adults aged 20 perform better on these tasks than adults aged 60 to 70 [82].

As spatial abilities are important to cognitive mapping, and these factors (gender, mathematics ability, and age) affect spatial abilities, this research studies the effect of these factors directly on cognitive mapping (see Chapter 5).

Chapter 3

The Spatial Learning Method

3.1 The Concept of Spatial Learning

The idea behind the spatial learning method flows logically from the cognitive mapping theory presented in Section 2.3. A person can explore a physical environment and learn where certain landmarks are in relation to each other, the approximate distances between them, and the route that needs to be followed to travel between them, all without even being aware of it. So, by designing an environment where the landmarks are physical representations of data items and these landmarks are placed in accordance to some set of relationships between them, someone exploring this environment should be learning not only the spatial relationships between the data points but also the relationships that the spatial ones represent. In other words, by setting up data relationships in a spatial form, and using this spatial form as the basis for the layout of an environment, someone wandering around in this environment is learning both the spatial form and the original set of data relationships without even being aware of it.

As a very simple example, consider a timeline of the history of life on earth: 2,500 million years ago, single-cell animal life began to appear in the seas; 690 million years ago, sponges and jellyfish began to appear; 470 million years ago the first fish appeared, and 35 million years later sharks appeared; 370 million year ago amphibians began to appear, and 200 million years ago the age of the dinosaurs began. The first mammals and birds appeared 195 million years ago, and 65 million years ago the dinosaurs became extinct; 15 million years after that rabbits and rodents appeared, and from then it only took 10 million years for early primates to appear. These large amounts of time are difficult to comprehend, and a diagram can help to show the time differences by presenting them visually. Such a diagram is shown in Figure 7. This diagram can form the basis for a very simple Virtual Environment, where the types of life that appeared at each stage could be represented by a virtual animal, and the visitor could subconsciously learn the length of time between the different types of life appearing, the order in which the different animals appeared, and which types of life began appearing at approximately the same times.

As another example, a family tree could be used as the basis for the layout of a VE. Parents could be linked by a East-West passage, with a North-South passage linking them to their children. The distance between family members would indicate how closely they were related, and the exact nature of the relationship would be indicated by the route taken
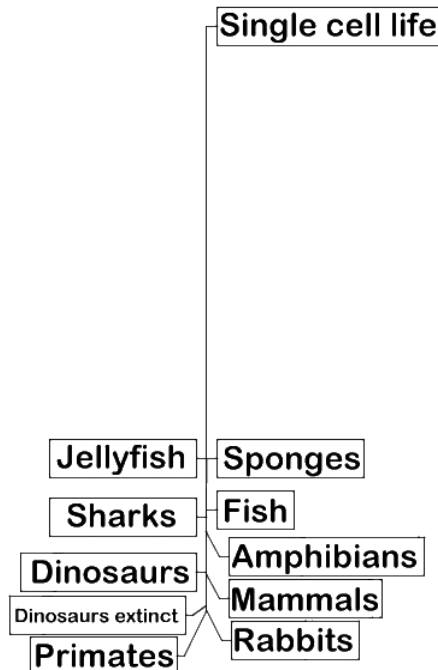
Figure 7: A diagram showing the chronological relationships between certain events in prehistory.

to travel between them.

3.1.1   Spatial Learning in Virtual Environment

Naturally, these environments would need to be created as Virtual Environments rather than as physical environments. This raises the issue of whether the cognitive mapping theory described in Section 2.3, on which the spatial learning concept is based, applies to virtual as well as real environments. Virtual environments differ from real environments in that they lack many of the physical cues normally available (see Sections 2.1, 2.3.3 and 2.3.5). VEs also suffer from other limitations, as described in Section 2.4.4.

There are few studies that focus directly on the issue of cognitive mapping in virtual environments — Chen & Stanney [14] present a theoretical model of wayfinding in Virtual Environments based on real world cognitive mapping theory, but do not seem to have any experimental validation of it as yet. However, there are many studies that look at wayfinding and navigational behaviour in Virtual Environments, summarized in Section 2.4.3. While the conclusions of these studies do conflict, it appears as though it is reasonable to assume that users do form cognitive maps of virtual environments. This assumption will, however, be tested in this research.

Both the references quoted above (relating to Virtual Environments) and those quoted in Section 2.3.5 (relating to real environments) warn that cognitive maps do contain distortions of spatial properties of the environment. This is of particular concern when the spatial properties are actually data properties and must be remembered accurately. However,

understanding the types of distortions and the reasons why they occur (Section 2.3.5) as well as careful design of environments according to the guidelines provided in Section 2.4.5 should help to eliminate most of the distortions. In addition, when the causes of the distortions are understood, virtual environments can be built not only to overcome these distortions, but to actually make use of them. For example, roads with lots of angles are generally perceived as being longer than straight roads of equal length (Sadalla & Magel, 1980, in [4], and [9]). Two objects can thus be placed close together, but joined via a winding route to give the impression of a greater distance.

If these potential difficulties are overcome, and the spatial learning method can be shown to be an effective way of teaching data sets and data relationships, then there are a wide range of situations in which it can be used, for example teaching aids at schools, as a traveling show (visiting schools, shopping centers, and such places), at exhibitions and fairs. The concept of Virtual Reality attracts many people who would not attend a presentation or look at a poster, and this attraction can be used to teach both adults and children in an enjoyable setting that nonetheless has lasting results.

3.2   Related Methods

The spatial learning method is different to conventional visualization in that users of scientific visualization are usually experts in their field [32]. Visualization aims to provide insight into complex phenomena, and to allow users to make inferences from the data provided. It is often used for transitory, changing data, and the user stands apart from the data. The spatial learning method, however, can be used to display many types of data for which visualization is not normally required (for example, the order of the planets in terms of distance from the sun), can be used by children and those with no knowledge of the subject being presented, and is primarily concerned with helping users understand and remember relationships between data items. Users are immersed in the environment created by the data, and are physically surrounded by objects as large as they are. The spatial learning method can be used for dynamic, changing data, but is useful mainly for stable data.

The spatial learning method is thus more similar to other spatial memory techniques such as the method of loci. This is the oldest known method of using spatial locations to remember data, and was originally used by students of rhetoric in Ancient Rome when memorizing speeches. To use it one must first memorize the appearance of a physical location (for example, the sequence of rooms in a building). When a list of words, for example, needs to be memorized, the learner visualizes an object representing that word in one of the pre-memorized locations. To recall the list, the learner mentally "walks through" the memorized locations, noticing the objects placed there during the memorization phase [6, 83]. Of course, using a familiar location as the starting point avoids the extra step of memorizing the physical location [83].

Another memorization technique based on spatial properties is the spatial arrangement mnemonic. Studies show that lists of words are more easily remembered when they are arranged in a distinctive spatial pattern than when they are presented as a list, and that when students used a link mnemonic to memorize these lists they recalled significantly more words than those who memorized them using conventional techniques [6]. For example, the
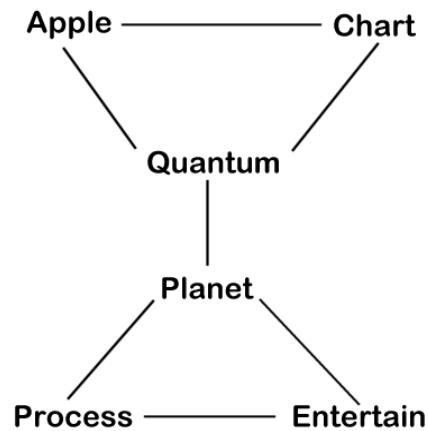
Figure 8: These six words are easier to remember when arranged in a distinctive pattern than when presented as a list.

list of words "Apple, Chart, Quantum, Planet, Process, Entertain" is much easier to remember when arranged in a distinctive pattern as in Figure 8. This is similar to the concept of diagramming, where related concepts are explicitly connected and the graphic representations rely on spatial metaphor to convey information. Such representations provide the necessary guidance for using imagery to remember abstract relationships [83].

Array memorization and manipulation can also be aided by spatial thinking. An array of items inherently contains spatial information about the items, and learning an array can be thought of as memorizing a tiny environment [83]. This concept is very similar to that of the spatial learning method, although that includes elements of the method of loci in that when recalling the locations of objects, one can picture oneself actually walking through the space.

In a study on abstract search spaces, Stubkjær [74] brings together several concepts which, combined, are similar to those behind the spatial learning method. He comments that an analogy can be made between concept maps and road maps, "with concepts corresponding to cities", and describes a wayfinding metaphor for navigating in abstract search spaces where, for example, file structures could be presented to the user as places in a landscape. He also describes a rooms metaphor where rooms might represent parts of a database. However, rather than interpreting these metaphors literally, as the spatial learning method does, Stubkjær creates another level of abstraction by using these representations as an intermediate step in the process of converting the abstract spaces to graphs and diagrams.

3.3   Characteristics of Suitable Data Sets

Any data set that can be represented as a diagram can be converted into a spatial layout, and thus into a Virtual Environment. However, data sets with certain characteristics are more likely to be successful:

- The data set must not have too many items, but also not too few. The more items there are the more complex the environment becomes, and the more time is needed

not just to travel through it, but also to traverse it often enough to build a cognitive map.

- The data items should be all be of the same type. Data sets representing an "is a" relationship (such as taxonomies) can be represented as diagrams (see the example in Figure 9(a)), and of course can then be converted into a spatial layout for a VE, but the true nature of a taxonomy is that of enclosure rather than that of a hierarchy (as in the alternative diagram of Figure 9(b)). This nature is difficult to convey when using a linear spatial form.

- The relationships within the data set should not be directly spatial. Spatial relationships can of course be represented in a Virtual Environment, but learning such an environment is straightforward cognitive mapping and as such does not really use the spatial learning method.

- It is preferable for the data relationships to form a graph rather than a tree. This allows for a more complex environment, as a graph has more cross-connections than a tree, and allows for bi-directional movement (a tree generally implies a single direction of movement down the tree, and not up-and-down movement as the user explores the environment).

- Timelines are in general not well suited to the spatial learning method, as they tend to be very linear and lead to an environment consisting of only one path. They also do not allow for up-and-down movement along the timeline.

- The data set should be static. Active data sets are possible, but the geography of the VE would have to be changed on the fly, which current technology does not yet support to any great extent. In addition, it is difficult to build a cognitive map of a constantly changing environment.

- While not essential, the items in the data set should preferably have a natural 3D graphical representation. This allows the user to immediately identify the objects in the environment with the items in the data set. If a 3D representation is not possible, a 2D representation can be placed on the walls of the environment.

- The spatial method works best when the length of the edges in the graph representing the data set are significant. This uses the true power of cognitive mapping, where not only the relative positions of the objects in the environment are important, but also the distances between them. By representing such distances in a VE, and forcing the user to walk along them, the user gets a more visceral feeling for the distances than when looking at a diagram.

The data set used in this research displays many of these characteristics, and is described in Section 4.3.

3.4 Potential Applications

There are many types of data sets that exhibit the characteristics detailed above, and thus many potential applications for the spatial learning method. Some examples are:
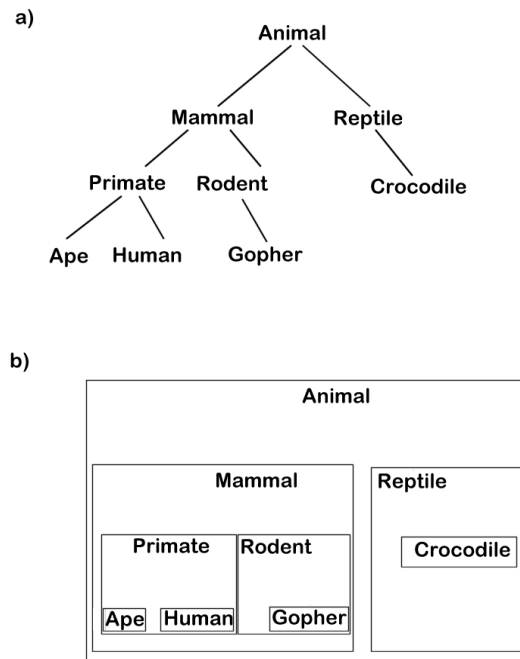
Figure 9: Two diagrams representing an "is a" relationship. Figure (a) shows the relationship drawn in a hierarchical form, while Figure (b) shows the relationship's true nature as an enclosing relationship.

- *Similarity Graphs.* Consider a group of composers. Composer A's music is similar to that of composer B because they both write music for string instruments. Composer A's music is also similar to that of composer C, because they create martial music, although composer C writes for the piano. Composer D also writes for the piano, but creates the same type of music as composer B. Then A is similar to B and C, but not to D. B is similar to A and D, while C is similar to A and D but not similar to B. D is similar to B and C. This is complicated to explain in words, but can be easily shown using a similarity graph (as in Figure 10, which can then easily be converted into a VE, where the composers are represented by objects within a room, and the passages between the rooms represent the similarity linkages on the diagram.

- *Processes.* Processes can also be learned using the spatial learning method, although they are not entirely suited to the method as they imply a certain time direction and thus an exploration direction. Actions can be represented as rooms, and choice points can be represented by a room with more than one exit (each exit corresponding to one of the choices). The length of the path between rooms could correspond to the amount of time that passes before the next action must be made or the next choice must be taken. Pipelining (for example, a CPU pipeline) can also be represented in this way, with each clock cycle being represented by a room within which the various activities are represented, and a passage which leads to the next room (i.e. the start of the next clock cycle). Delays in the pipeline could be represented by longer passages.

- *Trust Networks.* Trust networks are similar in form to similarity graphs, but the links
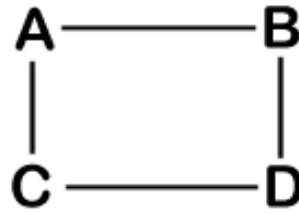
Figure 10: A similarity graph as described in Section 3.4. A link between two items implies that they are similar in some way.

between data items represent trust rather than similarity. Trust relationships are of particular importance in network security, such as in GNY analysis [30].

- *Routing Graphs.* Network routing tables can also be represented spatially. The visitor to the VE takes on the role of a packet traveling through the network under various network conditions, and experiences how the route may change depending on load and node conditions. This type of application may be better suited to a passive tour scenario, where the visitor is guided through the environment on a set path at a fixed speed, rather than a free exploration scenario.

- *Influence Graphs.* Influence graphs are again similar to trust networks and similarity graphs, except that the "Influences" relation is transitive. Such data sets can be converted to a VE in much the same way as similarity graphs.

- *Hierarchies.* Hierarchies such as family trees can also be easily converted to a Virtual Environment. Family members are represented by rooms (possibly with a photo on the walls indicating the particular family member represented by that room), North-South passages represent "child of" relationships, and East-West passages represent marriages.

- *Active Visualization.* A different type of application is that of active visualization, for example of network bandwidth. The nodes of the network could be represented by rooms in the VE, while the passages of the VE represent the bandwidth between the various nodes. The length of the passages would indicate the amount of bandwidth currently available between any given pair of nodes. The VE could dynamically reflect the current conditions, although certain problems then arise (as described in Section 3.3 above). Alternatively, the VE could represent a summary of the day's conditions, suitable for presentation to supervisors or others who need an overall picture of the conditions rather than a moment-by-moment account.

3.5   Non-linear Spatiality

3.5.1   Potential Applications

The examples given so far have all involved only one type of relationship between data items. However, intricate data sets may well have two or more sets of relationships that

need to be described. For example, an environment representing the solar system should represent distance from the sun, the size of the planets, and their relative gravities. It is very difficult to draw a single diagram showing all three relationships simultaneously, as planets which would be close together in terms of one relationship need to be far apart in terms of another. Virtual reality provides a solution to this problem, so that environments based on these relationships can be built even though they cannot be diagrammed. This is made possible through the use of items like teleports, which can transport visitors to another location in the VE instantaneously. A study by Ruddle *et al* found that the use of teleports reduced the accuracy of learning the spatial environment [62], but these problems would only apply when visitors are expected to infer the structure of the environment as it would have been without any teleports. With the spatial learning method, the teleports are used to create a new environment, with new spatial properties defined by the teleports themselves. The structure to be learned by participants is this new environment.

3.5.2   Examples

Even data sets with just one relationship can be difficult to draw diagrammatically. As an example, consider a group of four painters. Painter A is similar to painter B, because they both use oils. Painter B is similar to painter C, because they both paint the same type of subjects, and painter C is similar to D, because they both use similar subjects and use the same types of colours. Painter A is similar to painter D, because they both use oils. The following relationships then apply:

- A is similar to B and D, but not C;

- B is similar to A and C and D;

- C is similar to B and D, but not A.

A similarity graph can be drawn, as in Figure 11. The -//- symbol indicates that the link is actually shorter than it is drawn, and in a VE this would be implemented by means of a teleport. This graph fulfills all the similarity conditions that were presented:

- A is similar to B and D, but not C: only one step is needed to get to B and D from A, but 2 are needed to get to C;

- B is similar to A and C, and also D: only one step is needed to get to each of A, C, and D from B;

- C is similar to B and D, but not A: only one step is needed to get to B and D from C, but 2 are required to get to A.

As a more complicated example, with multiple relationships, consider a solar system with five planets: Alpha, Beta, Gamma, Delta, and Epsilon. That is the order in which they occur in terms of distance from their sun, with the planet nearest the sun at the top of the list (i.e. Alpha). They are all an equal distance away from each other (2 light-minutes). Alpha has 2 moons, as does Gamma; Delta has 3 moons, and Epsilon and Beta both have 5 moons. Alpha and Delta both have a gravity of 1g, Epsilon has a gravity of 2g, Gamma has a gravity of 4g, and Beta is a very dense planet with a gravity of 6g.
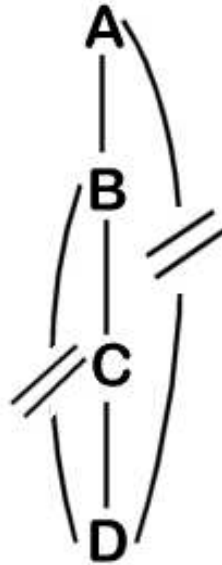
Figure 11: A similarity graph which is difficult to draw diagrammatically. The lines crossed by two diagonals lines need to be shorter than they appear.

We can then draw a similarity graph, as in Figure 12.   As before, the -//- symbol indicates that the link is actually shorter than it is drawn.

The power of a virtual reality can be used to create a single environment that exhibits all these incompatible spatial requirements. The environment could based on the "distance from the sun" relationship, with the solid black lines in the similarity graph showing the paths between the planet locations. To add in the "number of moons" relationship, a teleport could be placed between Alpha and Gamma, and between Beta and Epsilon. A path between Gamma and Delta could then be added. To connect Delta and Epsilon, either a longer, loopy path between Delta and Epsilon (twice the length of the path between Gamma and Delta) could be added, or an invisible teleport could be placed on the Delta-Epsilon path to move participants to another, disconnected section of path, and place another invisible teleport at the end of that path to bring them back to Epsilon. Similar techniques could then be used to add the "relative gravities" relationship.

Commingling all these paths and relationships, however, would probably be very confusing to participants, instead of allowing them to build up a clear picture of where things are and how they are related spatially. In a system like this one, therefore, we would need to make clear which path is part of which relationship. The "number of moons" paths could be decorated with moon-type pictures, for example, while the paths belonging to the "gravity" relationship could be decorated with a weight/mass motif. Alternatively, the "number of moons" paths could be surfaced with a texture resembling moon dust, complete with simulated craters. The particular method used is merely superficial; the true aim is simply to distinguish between the relationships and remind the participant which one they are currently traveling.
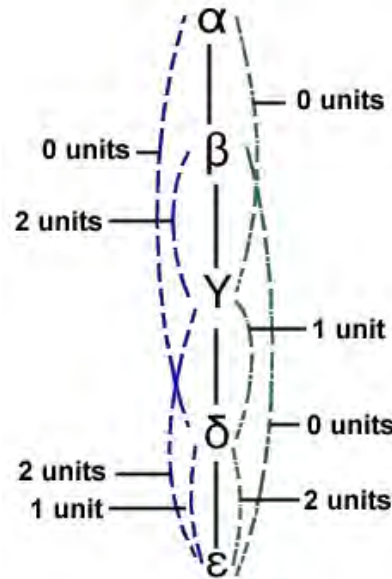
Figure 12: A similarity graph representing three relationships simultaneously. The lengths of each link are indicated in terms of generic units. The solid lines in the middle indicate the relative distances from the sun, the dashed lines on the left indicate the relative gravities, and the dashed-dotted lines on the right indicate the relative number of moons.

3.5.3   Issues with Non-linear Spatiality

There are several issues with creating environments using non-linear spatiality. Firstly, it is very difficult to create a map showing all the relationships simultaneously, and any such map is more likely to confuse visitors than help them find their way around. It might be better to create a map for each relationship, and automatically switch between them depending on which path the visitor is traveling.

A bigger problem is that of potential memory contamination. It seems likely that visitors would form a single cognitive map of the environment, containing all the relationships, rather than forming a separate cognitive map for each relationship. The cognitive mapping process may have difficulties in dealing with the incompatible spatial properties as these do not occur naturally, and the result may be confusion as to where the data items occur within which relationship.

In addition, while some VR engines do support teleports, non-linear spatiality would require undetectable teleports as well as teleports that can change their destination depending on which path the user approaches them on. This should not be difficult to program, and can certainly be implemented in the near future, but current VR engines do not support it. Authoring tools are also an issue, in that the planning and creation of environments exhibiting non-linear spatiality will be fundamentally different to that of normal, straight-forward environments. While it should be possible to use current authoring tools to create non-linear environments, it will be a difficult process and the development of authoring tools to support this type of design will be a great advantage.

Chapter 4

Experiment Design and Methodology

4.1 Areas of Interest

Three studies were designed to investigate the three areas of interest in this research described in Chapter 1: namely, issues related to cognitive mapping in Virtual Environments (VEs); issues related to the feasibility of the spatial learning method; and the psychological impact of VEs and spatial learning on the user.

Conducting exploratory studies into these three areas will help in forming a recommendation not only as to whether the spatial learning method is feasible, but also whether it is capable of replacing conventional learning methods for certain types of data sets. If visitors to a VE can form a cognitive map of the environment, can translate this to understand the underlying data set (at least as well as during a conventional presentation), and have a positive emotion set as a result of the experience, then the spatial learning method can be recommended as a tool for teaching data sets.

4.2 Study Design

The aims of the three studies are presented in Section 1.3, and repeated along with their findings in Chapter 5 (for Study 1), Chapter 6 (for Study 2), and Chapter 7 (for Study 3). Briefly, though, Study 1 explored issues relating to cognitive mapping in virtual environments, Study 2 investigated the efficacy of the spatial learning method and compared this to learning via a conventional presentation of the data set, and Study 3 explored the psychological effect of VEs and the spatial learning method on participants and compared this effect to that of a conventional presentation of the data set. The three studies were run on five different conditions:

1. A VR presentation of the data set, by means of an immersive Head Mounted Display (HMD) system with a map of the VE;

2. A VR presentation of the data set, by means of an immersive HMD-based system without a map;

3. A VR presentation of the data set, by means of a desktop system with a map;

4. A VR presentation of the data set, by means of a desktop system without a map;

5. A conventional lecture presentation of the data set.

Study 1 was performed on groups 1 to 4 (the VR groups), and followed the procedure described in Section 4.4.6.

Study 2 was performed on all five groups. The VR groups followed the procedure described in Section 4.4.6 and the conventional presentation group (group 5) followed that described in Section 4.5.2

Study 3 was again performed on all five groups as in Study 2. However, a particular psychological effect — that of the sense of presence — is specific to VR and thus was only analyzed for the VR groups.

The studies all used the same data set (described in Section 4.3), and where relevant, the same VE based on this data set (described in Section 4.4). The VR portions of all three studies were run concurrently on a single set of participants (see Section 4.4.6 for VR participant details).

4.3   Data Set

4.3.1   Description

The data set chosen for this set of studies was drawn from Greek Mythology, and forms a subset of the family tree of the Greek gods. It contains 24 Greek gods, and 32 links between them. Each god is associated with an image, representing either the god's primary attribute or some other feature of the god (such as a distinguishing event in their history, or a particular interest of theirs). For example, Zeus was represented by a lightning bolt, his attribute (see Figure 13), while Leto was represented by a palm tree, the site of the birth of her son (see Figure 14), and Hestia, the goddess of the hearth, was represented by a cooking fire (see Figure 15). The image was supplied to help participants visually distinguish between the gods, and, in cases where the names are difficult to spell and pronounce, to serve as a easier-to-remember representation of the god.

A diagram of the family tree is given in Figure 16. It is important to note that there are four types of relationships between data points:

- A horizontal link indicates that the two gods connected by the link had children together. This type of link is shown in the top left corner of Figure 17, and an example is the link between Coeus and Phoebe in the diagram of the family tree (Figure 16).

- A vertical link (sometimes angled slightly) emerging from a horizontal link indicates that the god at the end of the vertical link was a child of the two gods connected by the horizontal link. This type of link is shown in the top right corner of Figure 17, and an example is the link to Epimetheus in the diagram of the family tree (Figure 16).

Figure 13: The image associated with Zeus was that of a lightning bolt, his attribute.

Figure 14: The image associated with Leto was that of a palm tree, representing the site of the birth of her son.

Figure 15: The image associated with Hestia was that of a cooking fire, as she was goddess of the hearth.

- A vertical link emerging from a data point indicates that the god at the bottom end of the vertical link had as its only parent the god at the top end of the vertical link. This type of link is shown in the bottom left corner of Figure 17, and an example is the link between Epimetheus and Dione in the diagram of the family tree (Figure 16).

- A diagonal link coming from the horizontal link to the same data point as the horizontal link indicates that the god at the dual end of the link was both a child of and a partner to the god at the other end of the link. This type of link is shown in the bottom right corner of Figure 17, and an example is the link between Gaea and Uranus in the diagram of the family tree (Figure 16).

4.3.2   Features of Data Set

This data set was chosen for these studies because it fulfills many of the requirements given in Section 3.3 of Chapter 3 . The data points are all the same type of object, the relationships between data points are not directly spatial, it forms a static data set, and the data points can be represented graphically. The original family tree had far too many data points and relationships, but this was reduced to a manageable level. While the data set is in a sense a tree, which is not strictly desirable in data set, it also forms a graph: firstly, any path through the data set is bi-directional (visitors are not restricted from going back up the tree, or "backward in time"), and secondly, there are more cross-connections between data points than is strictly allowed in a tree (the "had children together" link could not be present in a data set which was strictly a tree).

The data set also makes use of several cognitive mapping principles. For example, the Rotation Effect (Section 2.3.5 in Chapter 2) allows us to have passages which are angled
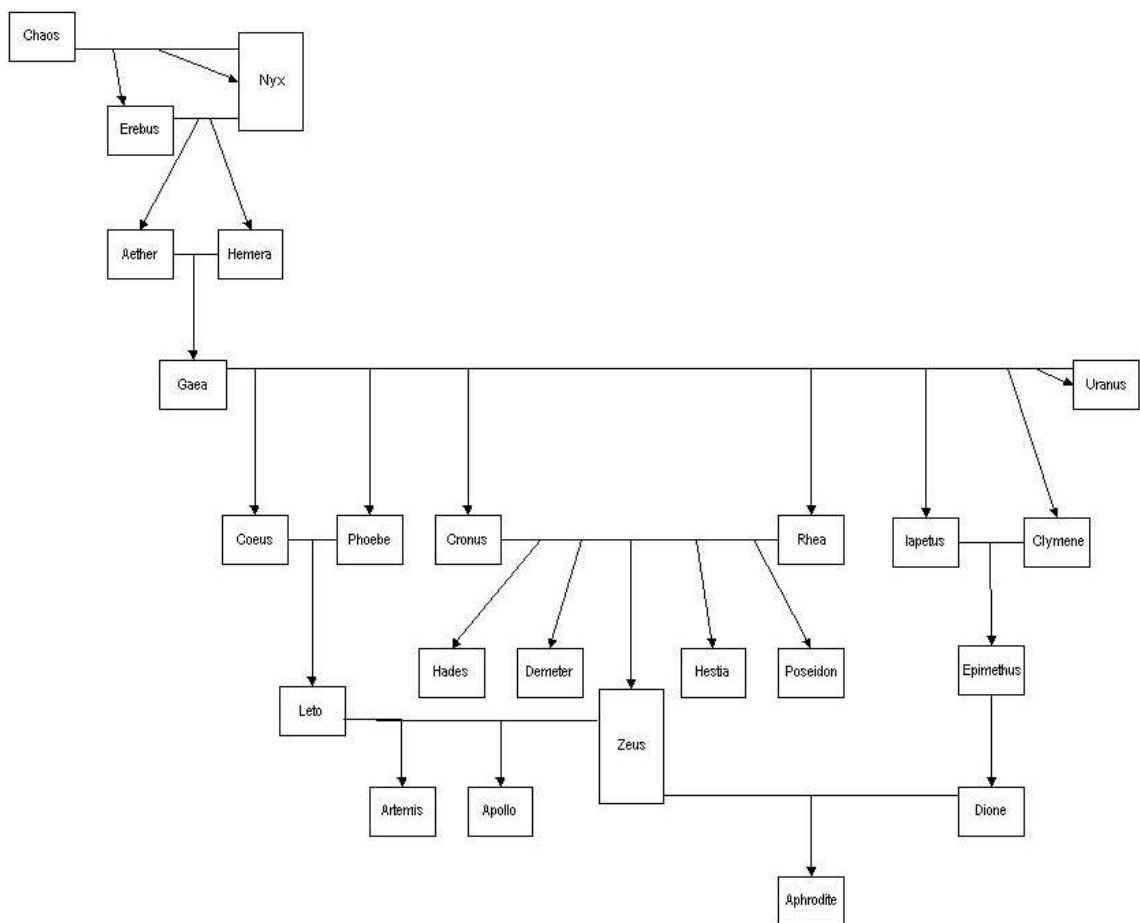
Figure 16: A subset of the family tree of the Greek gods, the data set used for this set of studies.
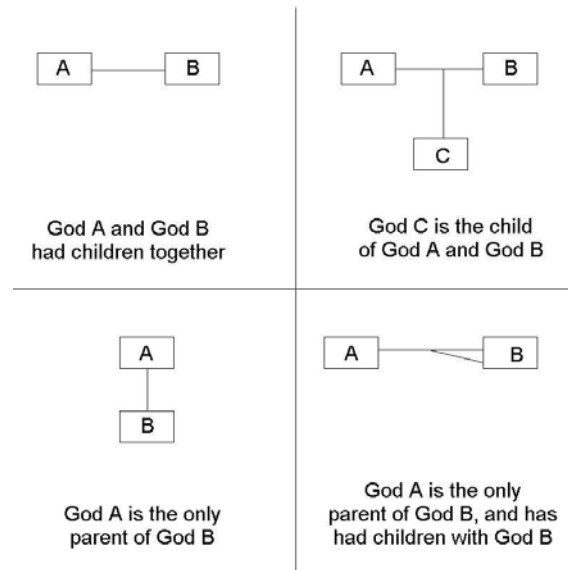
Figure 17: The four types of relationships between data points.

slightly, since visitors will mentally align them with both the cardinal directions and the non-angled passages around them.

One limitation of the data set is that it does not exhibit non-linear spatiality. Due to the difficulties described in Section 3.5.3 (in Chapter 3), including non-linear spatiality in this first study of the spatial teaching method would complicate both the study and the interpretation of the results. If this study proves to be encouraging towards the spatial teaching method, then further studies should certainly be performed to test the viability of including non-linear spatiality in the method.

4.4   Virtual Reality Presentation of Data Set

4.4.1   Environment

The layout of the virtual environment was based directly on the family tree presented in Section 4.3.1. The data points (for this data set, the Greek gods) were modelled physically as rooms, while the relationships between data points (the relationships between family members) were modelled as passages. Vertical links (i.e. parent-child links) were represented by North-South passages, while horizontal links (i.e. partnerships) were represented by East-West passages. The top-down view of the environment is thus the same as the family tree given in Figure 16.

The entire scene was set as a hedge maze in an outdoor environment (for modelling purposes, this was all contained within a larger enclosure). Textures were chosen appropriately to this setting: passage walls were textured with a plant hedge-type texture, while a texture resembling wooden planks was applied to the room walls. This helped to distinguish between the two areas. A sand-like texture was applied to floor, and a perspective-corrected sky texture was applied to the enclosing walls and ceiling. The sky texture included a
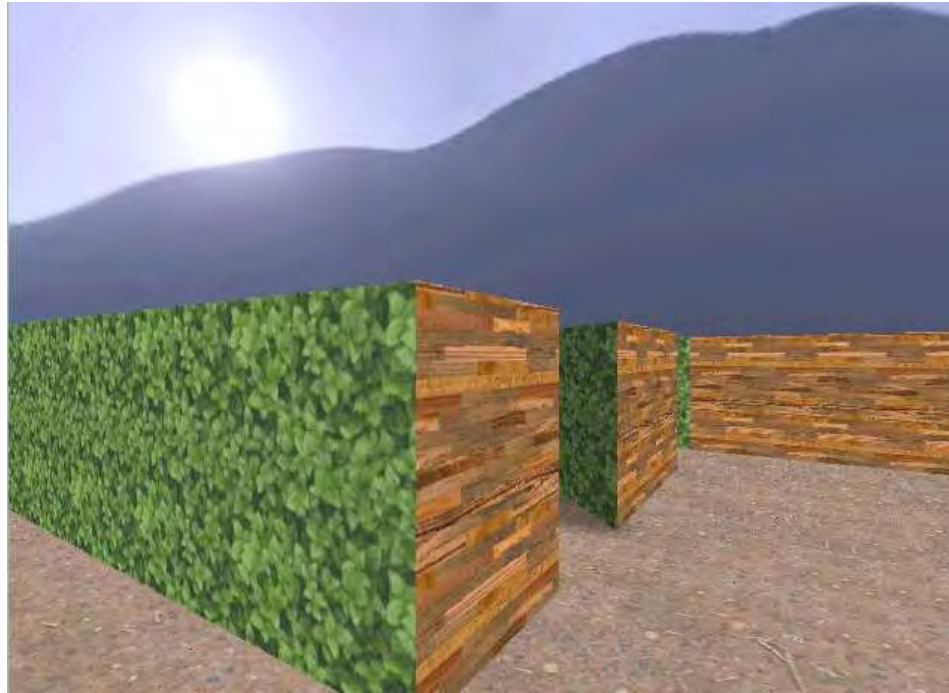
Figure 18: A screenshot of the VE, showing the different textures applied to surface features. The hedge texture was applied to passages, and the wood texture to rooms. The entire VE was enclosed in a box which was textured with a perspective-corrected sky.

virtual sun, which did not cast shadows but did serve as a reference point to help with navigation, and mountains were included on the wall-sections of the sky texture, also helping in navigation by serving to visually differentiate between different view angles. These types of navigational aids were recommended by Ruddle *et al* in [65]. Figure 18 shows a screenshot in which these features are visible.

Boxes textured with each god's name and image were placed in the appropriate rooms, as in Figures 13, 14, and 15. These "posters" were consistently placed at eye-level on east- or west-facing walls (rather than on the wall facing the passage, which varied between all four cardinal directions) in order to minimize confusion about the orientation of the participants.

Numbered tokens were placed in twelve of the twenty-four rooms, in accordance with the diagram in Figure 19. The positions of the tokens were selected to maximize the participant's exploration of the VE, and also to maximize the number of tokens that they would come across before seeing the tokens earlier in the numeric sequence (for example, before seeing token 8, we would like them to have already seen as many of tokens 9, 10, 11 and 12 as possible. See Sections 4.4.4 and 4.7.5 for more details on why this was necessary). On activation, a token would rise up into the air on a pole, high enough that it could be seen from a distance, thus becoming an additional navigational reference. Figure 20 shows an unactivated token in the foreground, with an activated token in the background.

Appropriate background noises (such as wind and birdsong) were played over earphones, and virtual footstep noises were generated in synch with the participant's walking motion within the VE. In addition, a sound was played when tokens were activated so as to give
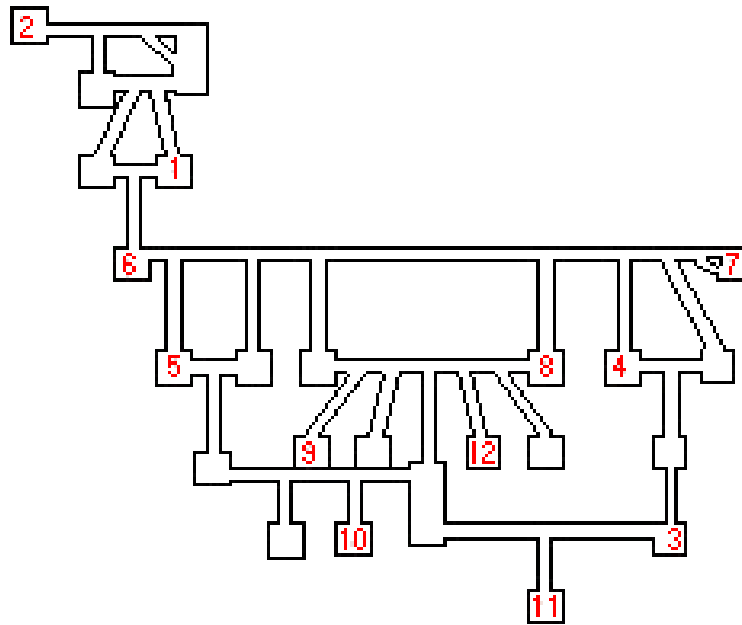
Figure 19: A top-down map of the VE, showing the locations of the numbered tokens.

some auditory feedback.

4.4.2 Movement and Interaction

Controlling movement within the VE differed for the HMD and desktop-based conditions. For the desktop system, movement was controlled via the mouse and keyboard; in the HMD condition, movement was controlled by motion trackers and a button device.

Navigation in the desktop condition was accomplished through the mouse and keyboard. The mouse was used to control the view direction (see Figure 21), while the W,A,S, and D keys were used to control the actual movement. These four keys are commonly used for this type of navigation, as they form a "⊥" shape on which the fingers can rest comfortably. For all four keys, movement occurs while the key is held down, and stops when the key is released. The W key is used to control forward movement, where forward is defined as "in the view direction"; the D key is used to control backwards movement, (i.e. in the opposite direction to the W key); and the A and S keys are used to control sideways motion, in a "side-stepping" fashion — the A key allows the participant to move perpendicular to the view direction, leftwards, while the S key allows the participant to move perpendicularly rightwards. In addition, the F key was used to activate tokens, and the M key was used to toggle the map display on and off (in the map conditions only).

Navigation in the HMD condition was controlled through the use of head-trackers and a custom-made handheld button device. The head-tracker was attached to the HMD itself, and controlled the view direction of the participant. The button device was a simple cylindrical object with 4 push-buttons built into it. The buttons were arranged ergonomically on the cylinder in order to allow easy access to each button (see Figure 22). Buttons marked 1 and 2 controlled forwards and backwards movement respectively (placed to be pressed

Figure 20: A screenshot of the VE, showing a room containing token number 6, with activated token number 5 in the background. Also visible is the poster showing the god's name (Gaea) and image (the Earth).

by thumb and forefinger), while button 3 controlled the activation of tokens. Button 4 toggled the map display on and off in the map condition, and was inactive in the non-map condition. As with the keyboard method of movement control, movement occurred while the button was held down, and stopped when the key is released. The control method used was therefore similar to the "Real Turn" mode used by Chance *et al* [12], where rotations are accompanied by visual, vestibular, and proprioceptive information, while translations are accompanied only by visual information.

4.4.3   Maps

A map of the environment was supplied to half of the participants in the VR conditions. The map was derived from the top-down view of the environment, and could be toggled on and off as desired. For the desktop conditions, the map was displayed on the top-left corner of the screen. When preparing for the HMD conditions, however, it was found that because of spherical distortion on the corners of the HMD display it was not possible to see the entire map when it was situated in the corner. It was thus moved horizontally to the center of the screen, while staying at the same vertical position.

Both the position and orientation of the participant displayed on the map, in accordance with the map principles presented in Section 2.4.6. The orientation was indicated by an
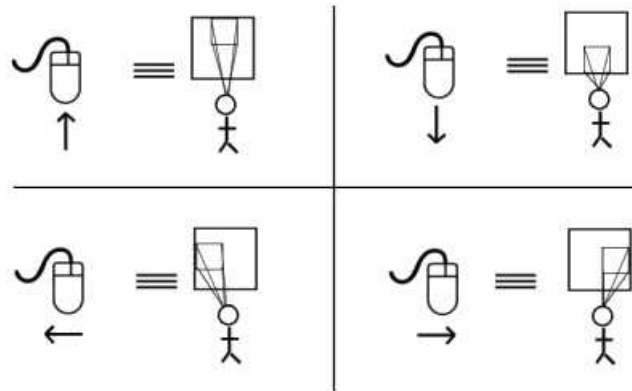
Figure 21: Controlling the view direction with the mouse: moving the mouse left/right controls the yaw rotation; moving the mouse up/down controls the pitch.

arrow which pointed in one of four directions: north, south, east, or west. If the participant's orientation was between 135° and 225°, then the arrow would point east; if it was between 225° and 315°, it would point south; and so on. The position was indicated by the position of the arrow on the map. Both position and orientation were updated in realtime.

The map did not indicate which gods were associated with which room, nor did they show the location of unactivated tokens. However, once a token was activated, its number would appear in the relevant position on the map, serving as another navigational aid. Figure 23 shows a screenshot in which the tokens 1 and 2 have been activated and thus appear on the map.

4.4.4 Task

Participants would have to spend a reasonable amount of time in the Virtual Environment in order to build up a cognitive map of the virtual space. In addition, they would have to explore the environment fully in order to cement their ideas of where things are in relation
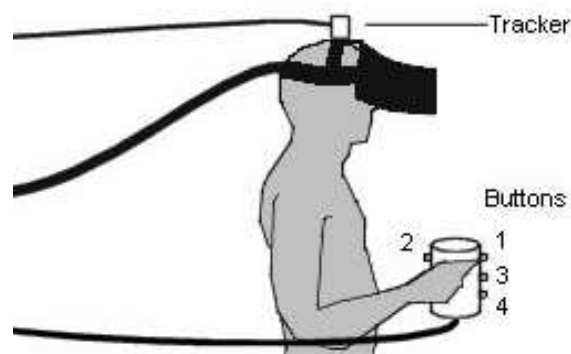
Figure 22: Controlling movement with the head-tracker and button device: the tracker controls pitch and yaw, while buttons 1 and 2 control forwards and backwards movement respectively. Button 3 activates tokens, while button 4 activates the map.

Figure 23: A screenshot of the desktop VE, showing the map supplied to some participants. Tokens 1 and 2 have been activated and appear on the map. Also note the arrow indicating the participant's position and orientation.

to each other. To facilitate these two goals, they were given a task to perform in the VE. Twelve numbered tokens were placed in various locations throughout the environment, and the participants had to find and activate these tokens in numerical order. A token could only be activated if all the tokens before it had already been activated, and the participant had to be within close range to the token. In the map conditions, the locations of activated tokens appeared on the map, but the location of tokens as yet not activated remained unknown, as in the non-map conditions.

The task thus required participants to take note of locations of objects in the VE, and to return to those locations when required. This facilitated the creation of the wayfinding measure (see Section 4.7.5), which is another method of accessing the participant's internal cognitive map of the environment.

4.4.5   Equipment

The desktop conditions were run on two desktop-based system, each with a 700Mhz AMD Athlon with a GeForce 2MX graphics card with 32Mb of onboard RAM, and a 17" monitor. The virtual environment was displayed at a resolution of 1024 x 768. The frame rate varied slightly depending on the amount of the environment visible at the time, but was in general

above 20fps. Stereo headphones were used to block out external noises, and also to allow the virtual noises to be heard.

The HMD conditions were run on a 700Mhz AMD Athlon system with a Matrox G400 graphics card to allow for dual-head display. The HMD itself was a V6 from Virtual Research, which has dual 1.3" diagonal Active Matrix Liquid Crystal Displays, and supports a resolution of 640 x 480 with 307,200 color elements and a diagonal physical FOV of 60° . The interpupillary distance was adjusted for each user. The trackers were Flock of Birds magnetic trackers used in position mode, where they have range of 3.05m and a static accuracy of 1.8mm RMS and a static resolution of 0.5mm at 30.5cm. Again, the frame rate varied slightly, but was generally about 12 fps. Participants using the HMD were seated on a swivel chair, so as to reduce the risk of tripping over cables while still allowing easy 360° rotation.

The graphics engine used was the opensource engine Genesis3D, with considerable custom-added functionality, including that for stereo.

4.4.6   Procedure

Participants were paid volunteers and were recruited by means of posters placed around campus. Participants signed up for one of the available experiment slots, which, unknown to them, were assigned to one of four conditions:

- Desktop system with map;

- Desktop system without map;

- HMD system with map;

- HMD system without map.

Participants were not aware of conditions other than their own. The desktop conditions were run first, with the HMD conditions following a few weeks later. Occasionally participants failed to turn up for a session, and slot/condition assignments were juggled so that there remained an equal distribution between those provided with a map and those without for each display condition. In total, 33 participants took part in the VR portion of the research.

Within the desktop conditions, two of the twenty subjects did not complete the session due to simulator sickness, leaving 18 participants (9 per map/no map condition) , while within the HMD conditions five of the thirteen subjects dropped out of the study due to simulator sickness, leaving only eight subjects (one of whom felt the effects of simulator sickness, but not strongly enough to withdraw from the study). This means that there were only 4 subjects per map condition within the HMD condition. In other words, participant distribution between the four groups was:

1. Desktop system with map — 9 participants;

2. Desktop system with no map — 9 participants;

3. HMD system with map — 4 participants;

4. HMD system with no map — 4 participants.

Low sample sizes are, in general, a cause for concern when using standard statistical analyses. For this reason, this study makes use of *nonparametric* statistics. More details are given in Section 4.8, but nonparametric methods make fewer and less stringent assumptions about the population, thus allowing smaller sample sizes [27].

Upon arriving for their session, participants were given a very brief introduction to the study, and were shown the equipment they would be using. They were given an opportunity to practice moving around in a training VE, during which they were instructed on how to use the controls relevant to their equipment (mouse/keyboard or tracker/button device), how to activate tokens, and those in map conditions were shown how to use and interpret the map. The task was explained to the participants, and they were encouraged to practice activating the two tokens within the training environment.

When it appeared that the participant was comfortable with the controls, the training environment was shut down and they were given instructions for the actual study. They were informed that although finding the tokens was important, their main task was to explore the environment and remember the layout and that they would be asked to draw a map of the environment afterwards. They were also told that they were to remember the locations of tokens that they saw, and that if they knew where the next token was they were to go directly to it. They were reminded that they could use the mountains and sun, as well as the activated tokens, to help orient themselves (as recommended by Ruddle *et al* [65]). They were told that the environment was based on a family tree of the Greek gods, and the gods corresponded to rooms and that the passages between rooms corresponded to relationships between the gods. The four types of relationships were explained to them (see Section 4.3.1 and Figure 17). They were also told that a poster with the god's name and a picture representing the god was placed in each of the rooms. After the first few cases of simulator sickness, participants were also warned that they might feel dizzy and nauseous, and to stop immediately if they began to feel ill. Once they were satisfied with the instructions, they were placed in the study environment. From this point onwards, they were not allowed to communicate with anyone until they finished the task. An observer remained in the room, out of sight, in case any problems arose.

Most participants finished the task within 30 minutes. A few gave up on the task when it seemed to them that they would not be able to find all the tokens, and some participants were stopped by the observer for similar reasons. A total of six participants thus did not finish the task. However, the minimum amount of time spent in the VE was over 13 minutes (a participant who *did* finish the task), and even those who did not finish the task stayed in the VE for between 22 and 45 minutes.

After finishing in the VE, participants were given a set of questionnaires to complete (see Sections 4.7 and 4.6.1 in particular), and were asked to draw a map of the VE. After finishing these tasks, they were paid and were allowed to leave.

4.5 Conventional Presentation of Data Set

4.5.1 Materials

The conventional teaching method chosen for this study was a lecture-type scenario. The lecture was directly based on the data set described in Section 4.3.1. At the beginning of the lecture, the entire family tree was displayed, and the four types of relationships were
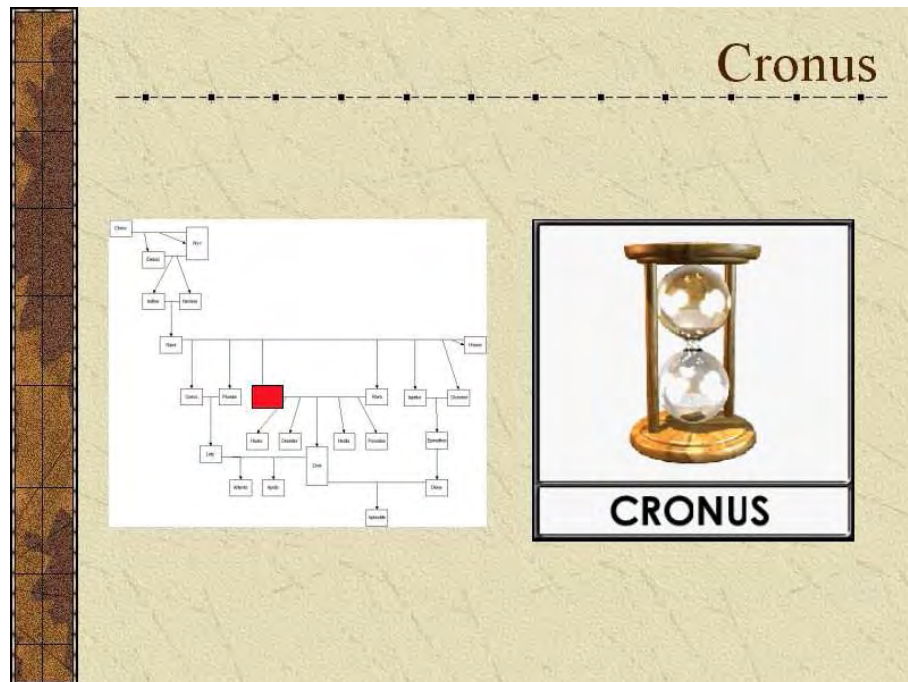
Figure 24: One of the slides from the lecture given to the non-VR participants. On the left is the family tree with the current god's position highlighted; on the right is the poster from the god's room in the VE.

explained (see Figure 17). Subsequently, each of the gods was presented per slide with a small image of the family tree with their position highlighted, together with their "poster" from the virtual environment (a sample slide can be seen in Figure 24). While showing the slide, the lecturer told the participants who the parents of the god under discussion were, who his/her partner was, and who their children were. This was done for each of the 24 gods, and at the end of the lecture the entire family tree was displayed and the relationships repeated.

4.5.2 Procedure

Participants were again paid volunteers, recruited from staff and postgraduate students working in the department (since this part of the study took place during university holidays, the recruitment procedure could be the same as for the VR conditions. However, the data set was chosen to be equally unfamiliar to all participants). There were 7 participants in this portion of the study.

Before beginning the lecture, participants were told that they would be given a set of questionnaires to fill in after the lecture, and also that they would be given a test on the subject material contained in the lecture.

The lecture took about 15 minutes, and no questions were asked during the lecture. Pencil-and-paper questionnaires where then handed out, and participants were not allowed to communicate with each other during the filling in of questionnaires and the test on the data set. As each participant finished, they were paid and allowed to leave.

4.6 Data Collection

4.6.1 Questionnaire Data

Questionnaire-style measures were administered via computer, rather than the more traditional pencil and paper approach. This allowed us to ensure that all questions were answered (and, where necessary, that only one answer is selected). It also simplified the collection of data, as responses were already captured in a digital format.

Each measure was introduced with a set of instructions relating to that particular measure before the individual items were presented. In the case of Likert-type scales, participants were presented with a statement or question, with a slider bar underneath. The points on the slider bar were numbered, and left and right anchors were named. The participant had to move the slider to the appropriate position, and click on the button labeled "next question" to continue.

Multiple choice questions were presented with the question above a column of up to 5 possible answers (with radio buttons next to the text); participants could select one and only one answer. Questions with yes/no answers were also presented in this way.

For short text answers, the participant was provided with a one-line text entry box in which to type the answer; again, they were required to type something in the box before they could continue to the next question. For text answer questions with multiple answers, multiple text entry boxes were provided, but participants needed to enter text into only one of the boxes.

Screenshots of each of these types of questions are provided in Appendix A.

4.6.2 VR Logging Data

In order to have a record of activity in the VE for later analysis, certain key data was written to a log file for each participant. Every 0.01 seconds a record containing the participant's x, y, and z positions were written to a file, along with their yaw and pitch rotations, which token they were currently looking for, and whether the map provided was currently visible or not (in the "no map" conditions, this was a fixed value of "not visible"). Enough data was recorded for the entire session to be recreated within the VE, as if it had been videotaped and was being played back.

4.6.3 Biographical Data

Biographical data about the participant was collected digitally via the questionnaire administration program described above. The participant was prevented from beginning the post-experimental questionnaires until all the biographical data had been entered.

The data collected was:

- Age (thought to influence cognitive mapping [33] and spatial abilities [82]);

- Gender (thought to influence cognitive mapping [1, 81] and spatial abilities [18]);

- Whether the participant had ever done any orienteering (which could influence cognitive mapping abilities);

- Whether the participant had ever taken a technical drawing course (thought to affect spatial abilities [46]);

- Primary means of transportation of the participant (walking, public transport, driving as driver, driving as passenger — thought to affect cognitive mapping abilities — see Section 2.3.3);

- How many hours per week the participant normally spent playing first-person computer games (a measure of previous VR experience, since most modern 3D games use the same technology as VR). Previous VR experience is thought to affect navigation ability [64];

- The highest level maths course taken by the participant (Final year school, 1st year undergraduate, 2nd year undergraduate, 3rd year undergraduate, one year postgraduate, more than one year postgraduate) (thought to influence spatial performance [18, 46, 83]).

This data was collected for all participants, although the non-VR participants completed a pencil-and-paper version rather than the digital one.

4.7   Measures

4.7.1   Psychological Factors

Emotional Factors — Interest,Enjoyment, Surprise, Distress, Shyness and Fear

The emotional factors were measured by the use of the Differential Emotions Scale (DES) developed by Izard [36]. The DES asks participants to rate their emotions during the session on a scale from 1 to 5. For the purposes of these studies, the questions were posed in the form: "To what extent did you feel ... during the session?", and anchor points were labelled "Not at all" and "Very much". The DES was presented via the questionnaire administration program described in Section 4.6.1 above, and a sample screenshot of the layout of the questions is shown in Figure 45.

The original DES measures 10 different emotions, not all of which were appropriate for this study. Questions measuring Interest, Enjoyment, Surprise, Distress, Shyness and Fear were included in the study, while those measuring Anger, Contempt and Guilt were excluded on the basis that these emotions were extremely unlikely to be evoked by the learning experience and would confuse the participant, possibly polluting the measurement of the other emotions.

Each emotion was measured using three related questions. The ratings for each of the questions was then summed to form a rating for that emotion. For example, the "Interest" emotion was measured using the following three questions:

- To what extent did you feel that you were concentrating during the session?

- To what extent did you feel attentive during the session?

- To what extent did you feel alert during the session?

On a scale from 1 to 7:

How natural did your interactions with the environment seem?
(1 being extremely artificial, 7 being completely natural)

How much did you focus on using the display and control devices instead of the
virtual experience and experimental tasks?
(1 being not at all, 7 being very much)

How compelling was your sense of moving around inside the virtual environment?
(1 being not compelling, 7 being very compelling)

Table 1: A few sample questions from the PQ, as described in Section 4.7.1. The full
questionnaire is given in Appendix C.

If a participant scored 3, 2, and 3 on these questions, then their "Interest" score was 8.

In order to avoid presenting similar questions directly after each other, the question
order was determined randomly. However, each participant received the questions in the
same order. Appendix B contains the DES questionnaire as used in these studies.

This measure was applied to subjects in both the VR and non-VR conditions, although
the non-VR participants completed a pencil-and-paper version rather than the digital one.

Cognitive Factors — Presence

The sense of presence was measured using the Presence Questionnaire (PQ) developed by
Witmer and Singer [89], as discussed in Section 2.2.2 . It consists of 32 likert-style ques-
tions representing 6 different factors that affect the sense of presence: involvement/control,
naturalness, interface quality, auditory, haptics, and resolution. A sample of the questions
asked is given in Table 1, while Appendix C contains the complete PQ as used in these
studies.

The PQ was presented via the questionnaire administration program described above
in Section 4.6.1, and was used in the VR conditions only.

4.7.2 Spatial Abilities

The participants' spatial abilities were measured using the Everyday Spatial Abilities Test
(ESAT), a self-report spatial abilities questionnaire developed by Lunneborg and Lunneborg
[46]. The ESAT consists of 20 items representing 4 different factors — use of hand tools,
maths/science courses, arranging objects, and mechanical drawing. Each item asks the
participant to rate how good they are at certain everyday activities, on a scale from 1 to
5 (1 being very bad, and 5 being very good). A total spatial abilities score is obtained by
simply adding together the ratings for each of the 20 questions. A sample of the questions
asked is given in Table 2, and the complete set of questions is given in Appendix D.

Additional factors which are thought to influence spatial abilities were also measured
(such as primary transportation method) as described in Section 4.6.3.

The ESAT was chosen above other spatial abilities measures partly because it is easy to

administer and fairly quick for participants to complete, but also because it correlates well with objective spatial abilities tests such as the spatial abilities and mechanical reasoning scores of the Washington Pre-College Test Battery [46].

ESAT data was collected via the questionnaire administration program described above in Section 4.6.1 for the participants in the VR conditions only.

4.7.3   Map Use

A measure of how much the participants who were provided with a map actually used that map to explore the Virtual Environment would be a more accurate measure of map use than a simple "map/no map" measure. Since the map was toggleable, the reasonable assumption is made that while the map was on (and thus visible to the participant), the map was being used by the participant. However, since each participant was in the VE for a different amount of time, a straight measure of *how long* the map was on for would not allow any comparisons to be made. Instead, the ratio of the amount of time the map was visible to the amount of time the participant spent in the VE must be calculated, and this ratio can then be used as a measure of map use.

A simple VR data log analysis program was written to count the number of log records written to file. It also counted the number of records for which the map was visible, and then computed the percentage of time that the map was on, as shown in Equation 1. Running this program on each log file provided us with our measure of map use for each participant.

$$map\ visible\ percentage = \frac{num\ records\ with\ map\ visible}{total\ number\ of\ records} \tag{1}$$

This is measure is not wholly accurate, because the recording process started when the VR application was started and was ended when the application was closed, rather than when the participant started and ended their task. Unfortunately, this means that some extraneous data was included at both the beginning and end of the log file. However, as the extraneous data lasted only a few seconds on either side, and even the shortest VR session lasted more than 13 minutes, the effect on the map use ratio was minimal.

4.7.4   Percentage Explored

It would be useful to know whether the participant was exposed to the entire VE, or whether they saw only a small section of it. This measure aims to address this issue by representing the number of rooms they explored as a percentage of the total number of rooms in the environment.

On a scale from 1 (very bad) to 5 (very good):

How good are you at sketching geometric designs?
How good are you at packing and storing boxes and luggage quickly and efficiently?
How good are you at solving practical problems using maths?
How good are you at understanding graphs and charts in textbooks and magazines?

Table 2: A few sample questions from the ESAT, as described in Section 4.7.2. The full questionnaire is given in Appendix D.
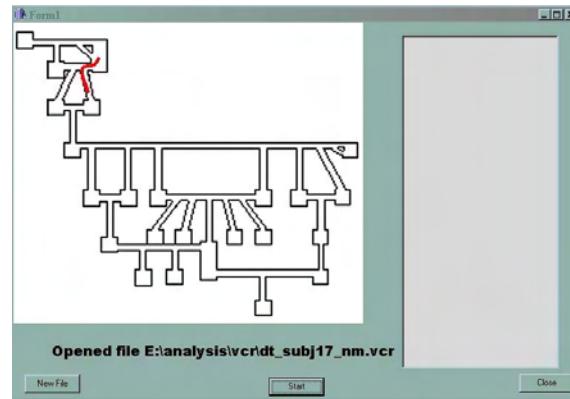
Figure 25: Analysing the participant's route through the VE

This is a conservative measure, because the assumption is made that a participant has only explored a room if they actually entered that room. In actual fact, in some cases it may have been possible for them to see everything within the room from an adjacent room or from a passage leading to the room, but as this can be a difficult judgement to make the more conservative but more consistent method of requiring the participant to have entered the room before considering them to have examined it will be used.

The data for this measure was again obtained from the log files described in Section 4.6.2. A second log file analysis program was written to read in the x and z positions from the log file and superimpose a small marker at the scaled position on a 2D map of the VE. A delay was imposed between the placing of subsequent position markers so that the route taken by the participant could be viewed as it progressed (to avoid cluttering the display, markers were removed every 1000 positions, thus allowing backtracking to be seen clearly). A screenshot of this process can be seen in Figure 25.

While viewing this animation for each participant, a written list was compiled of each room that the participant entered, in the order in which they were entered — in other words, a written route of the path that the participant took through the VE. From this it was determined how many of the 24 rooms were entered, and the percentage explored was determined.

4.7.5   Cognitive Mapping

Because cognitive mapping is a difficult construct to measure (see Section 2.3.6) two different approaches were used: a wayfinding score, and a sketch map score.

Wayfinding

The wayfinding score is based on three assumptions: firstly, that participants would remember if they had previously seen the token that they were currently looking for; secondly, that if they knew where the next token was they would go directly to it, and thirdly that they have only seen a token if they have entered the room in which the token is located.

The first two assumptions are fairly reasonable, especially as participants were instructed to keep track of what tokens they had seen and to go directly to a token if they knew where

it was. The third assumption is conservative, as for the "Percentage Explored" measure (see Section 4.7.4) — there were situations where a token could be seen from a neighbouring room or from the passage outside a room. However, it becomes difficult to establish the exact viewing angle of the participant in order to determine whether or not they have a seen a token (even with the playback allowed by the log files described in Section 4.6.2), and even if the viewing angle could be determined it would still require eye-tracking to definitely say whether or not the token was seen.

The wayfinding score is based on the number of excess rooms that the participant visited while searching for a token that they know the location of (that is, one that they have seen before).

To determine the wayfinding score for a participant, the written route through the environment obtained for the "Percentage Explored" measure (see Section 4.7.4) was used. This was supplemented by marking down on the written route each time a token was activated. The wayfinding score was then obtained as follows: For each token T:

- Determine if the room containing token T+1 was visited before the activation of T.

- If it wasn't, then begin again with token T+1. Otherwise:

    - Count the number of rooms that the participant visited when moving from the room in which T was activated until activating T+1 (including the room in which T+1 is situated).

    - Count the number of rooms on the shortest route between T and T+1.

    - Divide the number of rooms entered by the participant by the number of rooms on the shortest route. This number is the proportion of excess rooms visited while trying to find token T+1, and is a component of the final score (the two numbers cannot simply be subtracted because each sub-route will have different number of minimum rooms, and the score must reflect the number of excess rooms in proportion to the minimum in order to be able to aggregate them).

Once all the tokens have been processed, the final wayfinding score is obtained by adding up all the scores obtained in the last step above, and dividing by the total number of tokens for which this was done. This gives the average proportion of excess rooms visited during the task (we need to take the average because not all participants followed the same route through the VE, and so would not have encountered the tokens in the same order, implying that each participant would have different number of component scores. A participant who never visited any excess rooms would have a perfect wayfinding score of 1.

Sketch Maps

Participants had been told before they began the task that they would be asked to draw a map of the virtual environment when they were finished. They were given a sheet of plain A4 paper, and were instructed to draw the map to help other participants find their way around (as recommended by [34]). Participants were not instructed as to what details they should include (e.g. token placement).

The sketch maps thus produced were analyzed according to two different criteria: amount of detail, and accuracy. The detail score was obtained by averaging together four percentages: the percentage of rooms in the VE that were drawn on the map; the percentage of

passages in the VE that were drawn on the map; the percentage of god's locations indicated on the map; and the percentage of tokens indicated on the map. Rooms, passages, gods, and tokens were counted irrespective of whether or not they were in the correct position; however, maps which bore absolutely no resemblance to the VE at all were given a detail score of 0. A map which contained all possible features would have a detail score of 100.

The accuracy score was loosely based on the concept of local accuracy (as used by Rovine [61]), whereby accuracy is defined in terms of what information is included in the sketch map, rather than in terms of what existed in the environment. The accuracy score was derived from the number of errors in the sketch map. There were essentially five types of errors:

- God/token labels drawn in the incorrect position;

- Rooms/passages shown which clearly did not exist in the VE (as in Figure 26A);

- Room/passage placements which violated the child/partner relationship:

    - a child shown as a partner, or vice versa (Figure 26B);

    - a child shown as having one (or two) parents when they actually had two (or one) (Figure 26C);

    - a god being shown as a child of the wrong parent;

    - a god shown as only a child or partner when they were actually both (Figure 26D).

- Room/passage placements which violated the generational relationships (i.e. a god (or group of gods) being shown as a child rather than a grandchild (Figure 26E));

- Structural errors, such as a passage ending in a corner rather than entering a room (Figure 26F).

Distance and direction distortions were not counted as errors, and neither was missing data (as this was covered by the detail score).

After determining the number of errors for each sketch map, they were assigned to one of five categories:

- Maps in category A had 0 to 3 errors;

- Maps in category B had 4 to 7 errors;

- Maps in category C had 8 to 11 errors;

- Maps in category D had 12 to 15 errors;

- Maps in category E bore no resemblance to the VE at all.

The data was converted to this ordinal form primarily because it allows for easy expression of the fact that some maps were so incorrect that their errors could not even be counted.
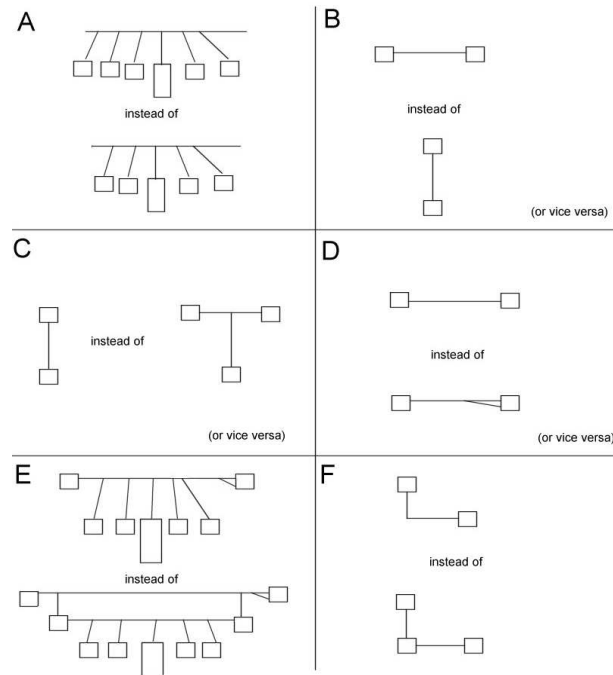
Figure 26: The types of errors made in sketch maps of the environment — A: passages/rooms which did not exist; B,C,D: violating the parent/child relationship; E: violating generational relationships; F: structural errors.

Name two gods without any children.
Was Leto of the same generation as the children of Cronus and Rhea?
Name two gods who had only one parent.
Was Iapetus Clymene's child or partner?

Table 3: A few sample questions from the test on the data set, as described in Section 4.7.6. The full test, together with the correct answers, is given in Appendix E.

4.7.6 Learning

In order to determine whether participants had learned the data points and the relationships between them, they were asked to complete a test on the data set. This test was administered via the questionnaire administration program described in Section 4.6.1 in the case of the VR participants, and via the standard pencil-and-paper method for the non-VR participants.

The test was designed to evoke the participant's knowledge of the data set, and included a range of questions (from very straightforward, requiring only a very basic understanding, to fairly tricky, requiring a true understanding of the data set). There were 14 questions in the test, some requiring single word answers, others requiring 2-part answers, and others that were multiple choice. The full set of questions (together with the correct answers) is included in Appendix E, but a sample is shown in Table 3.

The scoring of the questions took into account the fact that some of the VR participants would know the answer in spatial terms rather than in terms of the data set, and that while

this is better than not knowing the answer at all, it is not fully satisfactory. Accordingly, half marks were awarded if the answer was given in spatial terms (for example, "the room at the end of the long passage", or "box number 4"). Answers that were misspelled scored full marks, provided that they were intelligible, as were answers given in terms of the picture representing the data point rather than the name of the data point (e.g. "the lightning bolt" instead of "Zeus"), provided that the description was a reasonable and unambiguous approximation of the picture. The maximum score possible was 36.

4.8 Methods of Analysis

Details of the analyses performed (as well the results thereof) will are given in Chapters 5, 6 and 7. However, it is important to note here that nonparametric statistical methods were used to analyze the data. This was due in large part to the small sample sizes, and also because much of the data was ordinal.

Traditional (parametric) statistical procedures have a range of assumptions about the data that need to be satisfied in order for conclusions to be valid (a common assumption is that the population data fits a certain distribution). If the assumptions are disregarded or cannot be satisfied, there is far greater chance of the inferences made being incorrect. Nonparametric methods, however, make fewer and weaker assumptions about the population that the data comes from, and in some cases, no assumptions are made at all [27].

In many research situations, there is no basis for assuming any given distribution. Small sample sizes, in particular, mean that the distribution cannot be assumed, and thus any analysis done on such samples cannot be relied upon. With nonparametric methods, however, there is no minimum sample size required for validity and reliability. In addition, many procedures in parametric statistics require the data to be measured on at least an interval or ratio scale, while most nonparametric methods only require data to be ordinal [27].

There are nonparametric equivalents to many (although not all) of the traditional statistical tests. These equivalents are given below for the tests used in these studies.

4.8.1 t-test For Two Independent Samples

There are three nonparametric alternatives to the traditional t-test for independent samples: the Wald-Wolfowitz runs test, the Kolmogorov-Smirnov test, and the Mann-Whitney U test. This last is the one used in these studies, as it is essentially identical to the conventional t-test except that it is computed using rank sums rather than means. It can be interpreted in the same way as the conventional t-test, and is the most powerful nonparametric alternative [72].

4.8.2 χ^2 Contingency Tables

The traditional χ^2 contingency table analysis is actually nonparametric, as it deals with frequency data and does not assume that the population data fits any particular distribution [27]. It can thus be used in these analyses without reservation.

4.8.3 Correlations

The traditional parametric test for correlation is Pearson's product-moment correlation coefficient. There are three alternative nonparametric tests — Spearman R, Kendall *tau*, and *Gamma*. The alternative used in these studies is the Spearman R, can be interpreted in the same way as Pearson's product-moment correlation coefficient. The test assumes that the data is measured on at least an ordinal scale, i.e that the observations can be ranked into two ordered series. The Spearman R is the closest to the conventional parametric test, and has at least as much statistical power as the other two alternatives.[72].

4.8.4 1-way ANOVA

There is only one alternative to a 1-way ANOVA that applies to the data being analyzed in these studies, and this is the Kruskal-Wallis ANOVA by ranks. This test assumes that the data is continuous and was measured on at least an ordinal scale, and can be interpreted identically to a parametric 1-way ANOVA except that it is based on ranks rather than means [72].

4.8.5 2-way ANOVA

Unfortunately, there is no nonparametric alternative to a 2-way ANOVA. As a rough approximation, the following 3-step process has been used:

1. Perform an Mann-Whitney U test (nonparametric t-test) on factor 1 and the dependent variable to get the first main effect. If factor 1 has more than two levels (e.g. hmd vs desktop vs lecture rather than just male vs female), use a Kruskal-Wallis ANOVA by ranks instead.

2. Perform an Mann-Whitney U test (nonparametric t-test) on factor 2 and the dependent variable to get the second main effect. If factor 2 has more than two levels (e.g. hmd vs desktop vs lecture rather than just male vs female), use a Kruskal-Wallis ANOVA by ranks instead.

3. Combine the factors into one set of categories. For example, if the original factor 1 was "Gender" with levels "male" and "female", and the original factor 2 was "Presentation method" with levels "hmd", "desktop" and "lecture", the new set of categories would have 6 levels:

   - male-hmd;
   - male-desktop;
   - male-presentation;
   - female-hmd;
   - female-desktop;
   - female-presentation.

   Then perform a Kruskal-Wallis ANOVA by ranks on these combined categories to get the interaction effect.

It is important to note that this process does not give quite the same result for the interactions as a true 2-way ANOVA. However, it does give some idea of the interactions that are occurring, and the small cell sizes in these studies (typically 9 or 4) preclude the use of a parametric 2-way ANOVA (since any results obtained based on such small samples would be invalid and unreliable).

Chapter 5

Study 1 — Cognitive Mapping

5.1 Aims

This study investigated the issues surrounding the process of forming a cognitive map of a virtual environment (VE). The primary issue that needed to addressed was whether a visitor to a virtual environment could form a reasonably accurate and detailed cognitive map of that environment within the limited amount of time available.

The first section of this study was an experimental study, with two hypotheses:

1. The VR display system and input device used would affect cognitive mapping;

2. Providing participants with a dynamic You-Are-Here map showing their position and orientation would affect cognitive mapping.

This part of the study aimed to determine whether immersive equipment (such as head-trackers and Head Mounted Displays) was necessary to form a cognitive map of the VE, or whether a keyboard-and-monitor-based desktop system was sufficient. Immersive systems may give better cognitive mapping results as they provide more ideothetic information, such as vestibular and proprioceptive cues, but results in the literature are mixed (see Section 2.3.3). As such systems are more expensive to install and maintain than desktop systems, if these ideothetic cues are found to be not essential to the formation of a cognitive map, cheaper desktop systems could be used in place of immersive systems.

The second aim of this section of the study was to determine whether providing a map of the VE would help or hinder visitors with building an accurate cognitive map of the area. Some studies have shown that maps help with cognitive mapping, while other studies suggest that providing a map may actually hinder cognitive mapping (see the discussion in Section 2.4.6).

The possibility of an interaction between display system and map provision was also be investigated.

The second section of this study was a relational study, in which the relationships between various variables and cognitive mapping scores were investigated, namely:

- the percentage of the environment that was explored (a visitor cannot incorporate areas of the environment into their cognitive map if they have not seen them);

- whether or not the task was finished (finishing the task required participants to revisit places several times, which should improve cognitive mapping);

- previous VR/gaming experience (lack of familiarity with VEs may hinder navigation [64], while those who regularly play games requiring them to remember the location of virtual objects may form cognitive maps of higher quality, and may have developed better navigation skills);

- age (the ability to adopt new technology may decrease with age, and Holahan [33] suggests that age may affect cognitive mapping skills although Appleyard [1] found no age-based differences.);

- gender (some studies have shown gender effects in cognitive mapping [1, 39, 81]);

- primary travel mode (transportation method has been shown to affect cognitive abilities [1, 11, 16, 33, 73]; see Section 2.3.3);

- spatial abilities (several studies have shown that spatial abilities affect cognitive mapping skills [8, 18, 67, 83]; see Section 2.5);

- percentage of time that a map of the VE was visible (a refinement of the map / no map test, which looks at how long the participant actually saw the map during the VE exposure);

- Psychological factors such as emotional aspects (positive emotions may help cognitive mapping, while negative emotions may be detrimental) and presence (if a visitor feels that they are in a real place, that may help them to form a cognitive map of that place).

The design of this study and the procedure followed in implementing it is as described in Chapter 4, as are the measures used.

5.2 Results

5.2.1 Formation of a Cognitive Map

Out of 26 participants, only five (19.23%) drew sketch maps that did not resemble the VE at all. This implies that it is certainly possible to form a cognitive map of a virtual environment after being exposed to it for a relatively short period of time — exposure times varied between 13 and 60 minutes, with a mean exposure time of 30.6 minutes and standard deviation of 11.96 minutes.

However, the cognitive maps formed of the VE were of fairly low quality. Sketch map detail scores (as described in Section 4.7.5) varied between 0 and 66.4 (see Figure 27 , but when excluding the maps that bore no resemblance to the VE the minimum score was 15 (a perfect detail score was 100). With regards to accuracy (also described in Section 4.7.5), three maps had between 0 and 3 errors, seven had between 4 and 7 errors, five had between 8 and 11 errors, six had between 12 and 15 errors, and as mentioned five maps showed no similarity to the VE (these error rates are shown as a histogram in Figure 28). Wayfinding scores (described in Section 4.7.5) varied between 1.25 and 12 (see Figure 29), where a
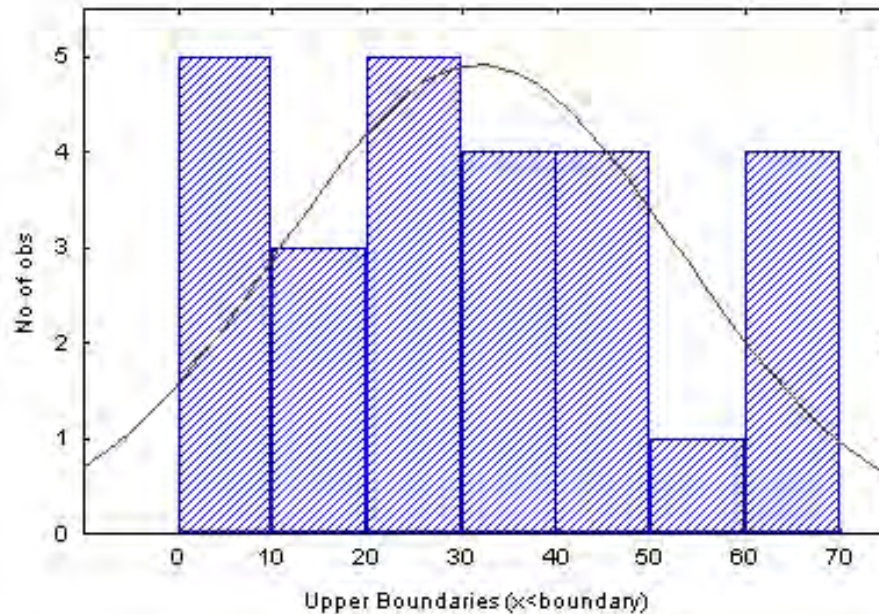
Figure 27: Histogram showing the sketch map detail scores. The solid curve indicates the expected normal.

perfect wayfinding score would be 1 (a wayfinding score could not be determined for one participant, as their route through the VE did not pass through any of the target rooms before activation of the relevant tokens).

Longer exposure to the VE may help to increase the quality of the cognitive map formed, but substantial increases in exposure time would have to be made — with exposure times varying between 13 minutes and 60 minutes in the current study, the time spent in the VE did not significantly affect any of the three cognitive mapping scores (Table 4).

5.2.2 Effect of Display System and Map

To investigate the combined effect of display type and map provision on cognitive mapping scores, a 2-way ANOVA would normally be used. However, there is no non-parametric equivalent to a 2-way ANOVA, so a replacement 3-step process has been used (see Section 4.8).

Step 1 is to determine the effect of display type on cognitive mapping, by means of a

	Valid N	Spearman R	p-level
Time & detail	26	-.172539	.399308
Time & accuracy	26	-.329383	.100352
Time & wayfinding	25	.216805	.297888

Table 4: Spearman R and p values for the relationship between time spent in the VE and the three cognitive mapping scores. No significant relationships were found.
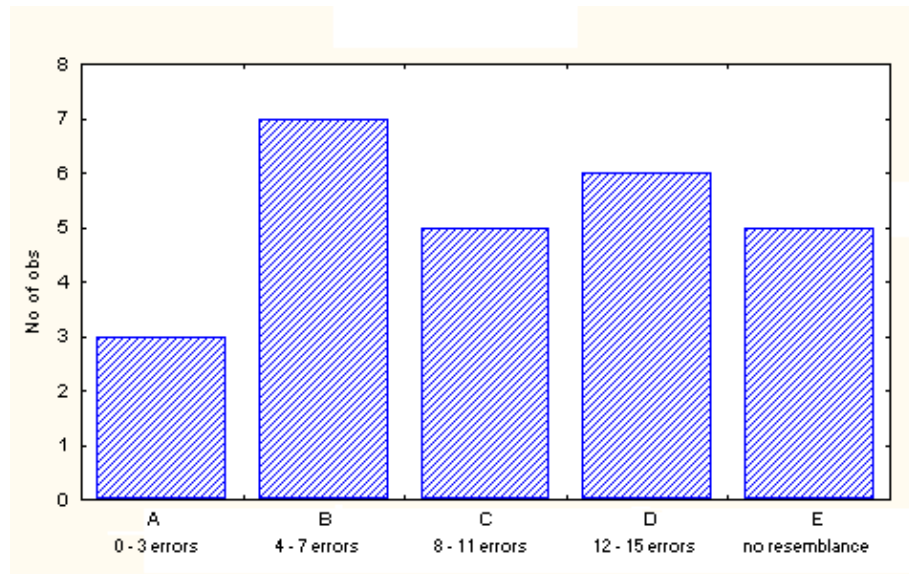
Figure 28: Histogram showing the number of sketch maps in each error category. The solid curve indicates the expected normal.
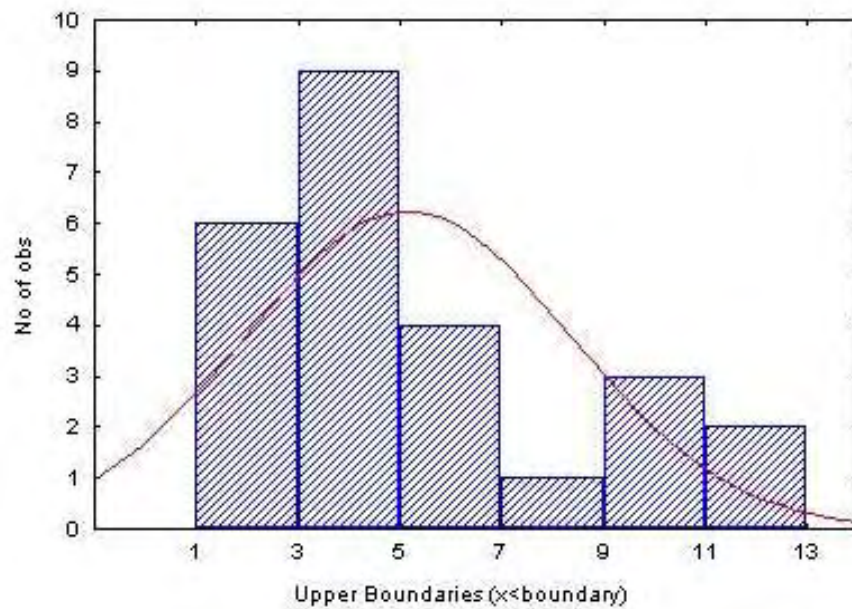
Figure 29: Histogram showing the wayfinding scores. The solid curve indicates the expected normal.

	U	Z	p-level	N (desktop)	N (HMD)
Detail	65.00000	.38889	.697361	18	8
Accuracy	51.50000	-1.13889	.254758	18	8
Wayfinding	54.50000	.51444	.606946	18	7

Table 5: Mann-Whitney U test on the effect of display type (HMD/desktop) on cognitive mapping scores. No significant differences were found.

Mann-Whitney U test. The results are presented in Table 5. For the effect on sketch map detail scores, a U value of 65.00 and a corresponding p-value of .697361 were obtained; for the effect on sketch map accuracy scores a U value of 51.50 and a corresponding p-value of .254758 were obtained; and for the effect on wayfinding a U value of 54.50 and a corresponding p-value of .606946 were obtained. Hypothesis 1 (as described in Section 5.1 is thus rejected. As no significant differences were found, display type did not affect cognitive mapping.

Step 2 is to determine the effect of having a map on cognitive mapping, again by means of a Mann-Whitney U test. The results are presented in Table 6. For the effect on sketch map detail scores, a U value of 75.50 and a corresponding p-value of .644415 were obtained; for the effect on sketch map accuracy scores a U value of 65.50 and a corresponding p-value of .329886 were obtained; and for the effect on wayfinding a U value of 54.00 and a corresponding p-value of .191757 were obtained. Hypothesis 2 (as described in Section 5.1 is thus rejected. As no significant differences were found, having a map of the VE did not affect cognitive mapping.

Step 3 is to determine whether there were any interactions between display type (HMD vs Desktop) and having a map. This was done by means of a Kruskal-Wallis ANOVA by ranks. No significant results were thus found, which means that there was no interaction between display type and having a map on cognitive mapping. These results are shown in Table 7.

5.2.3 Effect of Other Variables

A significant relationship was found between percentage explored and cognitive mapping — participants who explored more of the environment included more detail on their sketch maps (Spearman R gives R=.432960, p=.027152, α=0.05). Accuracy and wayfinding were not related to the percentage explored, however (see Table 8), and the percentage of time

	U	Z	p-level	N (map)	N (no map)
Detail	75.50000	.46154	.644415	13	13
Accuracy	65.50000	.974359	.329886	13	13
Wayfinding	54.00000	-1.30543	.191757	13	12

Table 6: Mann-Whitney U test on the effect of having a map on cognitive mapping scores. No significant differences were found.

Sketch Map Detail:
H (3, N= 26) = 3.643978, p =.3026

	Valid N	Sum of Ranks
desktop-map	9	115.0000
desktop-no-map	9	135.0000
hmd-map	4	69.5000
hmd-no-map	4	31.5000

Sketch Map Accuracy:
H (3, N= 26) = 4.34649,9 p =.2264

	Valid N	Sum of Ranks
desktop-map	9	112.0000
desktop-no-map	9	110.5000
hmd-map	4	82.5000
hmd-no-map	4	46.0000

Wayfinding:
H (3, N= 25) = 3.795747, p =.2844

	Valid N	Sum of Ranks
desktop-map	9	93.0000
desktop-no-map	9	149.5000
hmd-map	4	52.0000
hmd-no-map	3	30.5000

Table 7: Kruskal-Wallis ANOVA by ranks on the combined display/map categories for sketch map detail, sketch map accuracy, and wayfinding scores. No significant effects were found.

that participants had the map displayed for did not influence any aspect of cognitive mapping (Table 10).

Having a map of the VE did not affect the percentage explored (Table 9), which seems counterintuitive but fits in with previous results — if having a map *had* influenced the percentage of the environment explored, then as percentage explored did affect the sketch map detail scores, having a map would indirectly affect cognitive mapping. But it has already been shown in Step 2 above that having a map did *not* affect cognitive mapping, and so the finding that having a map did not affect the percentage explored is consistent with these results.

Finishing the task did not affect cognitive mapping (Table 11), and neither gaming experience nor age were found to be related to cognitive mapping (see Tables 12 and 13). One participant refused to disclose their age, so sample sizes for the age analysis are one smaller than those for other analyses. The average age was 24.6, with the youngest participant being 18 and the oldest, 50.

A Mann-Whitney U test on gender and cognitive mapping showed a significant relationship between gender and wayfinding (U=33.0 , p=.014383, α=0.05), although not between gender and either of the sketch map variables (see Table 14). The median wayfinding score

	Valid N	Spearman R	p-level
Percent explored & detail	26	**.432960**	**.027152**
Percent explored & accuracy	26	.135158	.510338
Percent explored & wayfinding	25	.134422	.521769

Table 8: Spearman R and p values for the relationship between percentage of environment explored and the three cognitive mapping scores. A significant relationship (marked in bold, α=0.05) was found between percentage explored and cognitive mapping.

	U	Z	p-level	N (map)	N (no map)
Percent Explored	51.50000	1.692308	.090597	13	13

Table 9: Mann-Whitney U test on the effect of having a map on the percentage of the environment explored. No significant difference was found.

Spearman Rank Order Correlations
MD pairwise deleted

	Valid N	Spearman R	p-level
Percent map on & detail	26	.047311	.818478
Percent map on & accuracy	26	.081929	.690718
Percent map on & wayfinding	25	-.284678	.167816

Table 10: Spearman R and p values for the relationship between percentage of time the map was displayed and the three cognitive mapping scores. No significant relationships were found.

	U	Z	p-level	N (not finished)	N (finished)
Detail	53.50000	-.395577	.692419	6	20
Accuracy	55.00000	-.304290	.760909	6	20
Wayfinding	46.50000	-.237778	.812055	5	20

Table 11: Mann-Whitney U test on the effect of finishing the task on cognitive mapping. No significant difference was found.

	Valid N	Spearman R	p-level
Gaming experience & detail	26	.180335	.378012
Gaming experience & accuracy	26	.224989	.269141
Gaming experience & wayfinding	25	.139769	.505188

Table 12: Spearman R and p values for the relationship between gaming experience and the three cognitive mapping scores. No significant relationships were found.

	Valid N	Spearman R	p-level
Age & detail	25	-.193305	.354542
Age & accuracy	25	-.033378	.874151
Age & wayfinding	24	-.068908	.749017

Table 13: Spearman R and p values for the relationship between age and the three cognitive mapping scores. No significant relationships were found.

was higher for males than for females (Figure 30), which means that women performed better than men (the wayfinding score give the average excess percentage of rooms entered when trying to find the target tokens, and thus a higher score indicates worse performance).

Travel mode did not affect any aspect of cognitive mapping (see Table 15, and neither did orienteering experience (Table 16 or maths level (Table 17). Technical drawing experience affected both detail (U=43.50, p=.054399, which is close to significant at a 5% significance level) and wayfinding scores (U=24.00, p=.004673, α=0.05), but not accuracy (Table 18. Those participants who had technical drawing experience included more detail in their sketch maps (as can be seen in Figure 31, the median score is higher for those who had taken a technical drawing course) but performed worse at wayfinding (the median wayfinding score is higher for those with technical drawing experience, as can be seen in Figure 32). ESAT scores correlated significantly with the detail score (R=.408662, p=.038192, α=0.05), showing that participants with better spatial abilities included more detail in their sketch maps, but did not correlate with the other cognitive mapping scores (Table 19)

Fear and surprise were associated with decreased accuracy of the sketch maps (R=-.404995, p=.040131 for fear; R=-.556910, p=.003125 for surprise, α=0.05) and shyness correlated with a decreased amount of detail included in the sketch maps (R=-.493673, p=.010376, α=0.05). No other relationships between psychological factors and cognitive mapping aspects were found, including the sense of presence (see Table 20).

	U	Z	p-level	N (male)	N (female)
Detail	80.50000	.180021	.857138	14	12
Accuracy	79.00000	.257172	.797047	14	12
Wayfinding	33.00000	2.447677	**.014383**	13	12

Table 14: Mann-Whitney U test on the effect of gender on cognitive mapping. A significant difference was found (marked in bold, α=0.05), in that females performed significantly better at wayfinding than males.
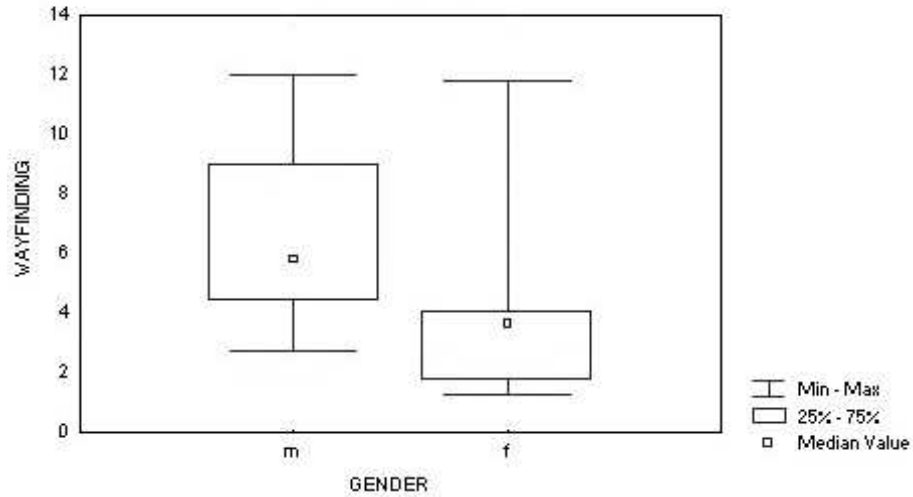
Figure 30: A box plot showing the median and quartile wayfinding scores for males and females. The median for males was higher than that for females.

Sketch Map Detail:
H (3, N= 26) = .3414184, p =.9521

	Valid N	Sum of Ranks
Walking	6	84.00000
Public transport	8	98.50000
Driving (driver)	6	81.00000
Driving (passenger)	6	87.50000

Sketch Map Accuracy:
H (3, N= 26) = 1.048387, p =.7895

	Valid N	Sum of Ranks
Walking	6	95.00000
Public transport	8	94.00000
Driving (driver)	6	84.00000
Driving (passenger)	6	79.00000

Wayfinding:
H (3, N= 25) = 1.649619, p =.6482

	Valid N	Sum of Ranks
Walking	6	86.00000
Public transport	7	87.00000
Driving (driver)	6	91.00000
Driving (passenger)	6	61.00000

Table 15: Kruskal-Wallis ANOVA by ranks on the effect of transportation method on cognitive mapping. No significant effects were found.

	U	Z	p-level	N (no orienteering)	N (orienteering)
Detail	78.50000	-.079057	.936988	16	10
Accuracy	69.50000	.553399	.579994	16	10
Wayfinding	61.00000	-.776580	.437412	15	10

Table 16: Mann-Whitney U test on the effect of previous orienteering experience on cognitive mapping. No significant differences were found.

Sketch Map Detail:
H (3, N= 26) = .2.282837, p =.5158

	Valid N	Sum of Ranks
1st year university course	3	24.5000
2nd year university course	10	152.5000
3rd year university course	10	140.0000
Postgraduate maths course	3	34.0000

Sketch Map Accuracy:
H (3, N= 26) =.6787634, p =.8782

	Valid N	Sum of Ranks
1st year university course	3	43.5000
2nd year university course	10	142.5000
3rd year university course	10	120.0000
Postgraduate maths course	3	45.0000

Wayfinding:
H (3, N= 26) = 1.004772, p =.8001

	Valid N	Sum of Ranks
1st year university course	3	43.0000
2nd year university course	10	119.5000
3rd year university course	10	143.0000
Postgraduate maths course	2	19.5000

Table 17: Kruskal-Wallis ANOVA by ranks on the effect of maths level on cognitive mapping. No significant effects were found.

	U	Z	p-level	N (no tech draw)	N (tech draw)
Detail	43.50000	-1.92372	**.054399**	16	10
Accuracy	70.50000	.50069	.616590	16	10
Wayfinding	24.00000	-2.82897	**.004673**	15	10

Table 18: Mann-Whitney U test on the effect of technical drawing experience on cognitive mapping. A significant relationship was found between technical drawing and amount of detail included in sketch map, and wayfinding score (marked in bold, α=0.05).
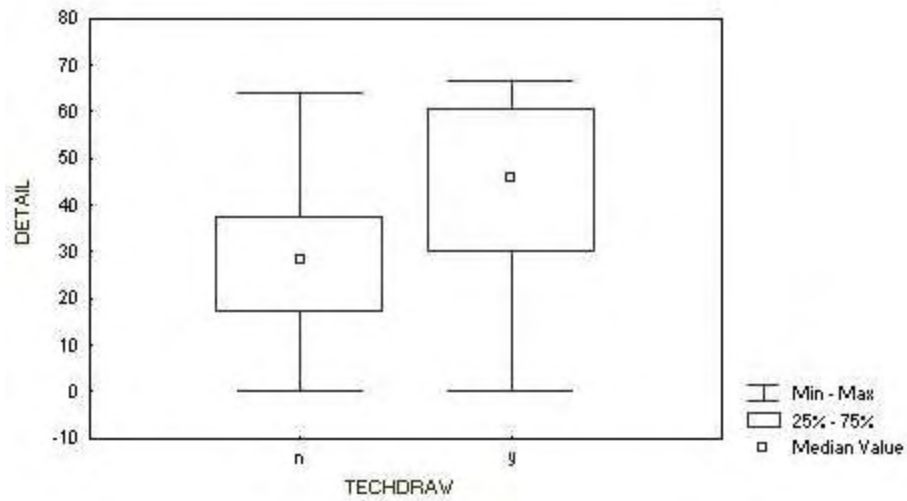
Figure 31: A box plot showing the median and quartile detail scores for technical drawing experience. The median for those with technical drawing experience was higher than that for those without, showing that those with technical drawing experience included more detail in their sketch maps.
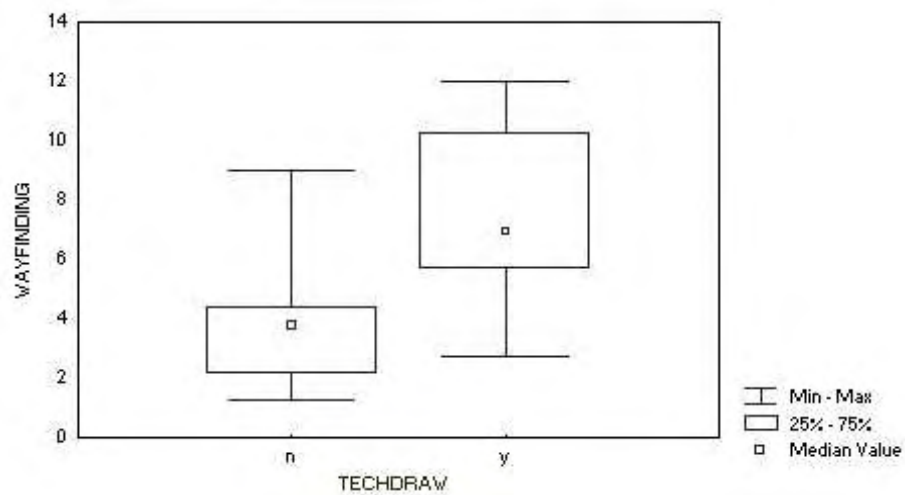
Figure 32: A box plot showing the median and quartile wayfinding scores for technical drawing experience. The median for those with technical drawing experience was higher than that for those without, showing that those *without* technical drawing experience performed better.

	Valid N	Spearman R	p-level
ESAT & Detail	26	**.408662**	**.038192**
ESAT & Accuracy	26	-.005785	.977625
ESAT & Wayfinding	25	.198497	.341511

Table 19: Spearman R and p values for the relationship between spatial abilities and cognitive mapping. A significant positive correlation was found between spatial abilities and the amount of detail included in the sketch map (marked in bold, α=0.05).

	Valid N	Spearman R	p-level
Presence & detail	26	.293957	.144949
Presence & accuracy	26	.071460	.728667
Presence & wayfinding	25	-.020019	.924330
Enjoyment & detail	26	.049232	.811237
Enjoyment & accuracy	26	-.202989	.319958
Enjoyment & wayfinding	25	-.250343	.227444
Fear & detail	26	-.252398	.213520
Fear & accuracy	26	**-.404995**	**.040131**
Fear & wayfinding	25	-.089613	.670118
Surprise & detail	26	-.003822	.985218
Surprise & accuracy	26	**-.556910**	**.003125**
Surprise & wayfinding	25	.120088	.567474
Shyness & detail	26	**-.493673**	**.010376**
Shyness & accuracy	26	-.367625	.064654
Shyness & wayfinding	25	.037036	.860484
Interest & detail	26	.160716	.432857
Interest & accuracy	26	-.148283	.469728
Interest & wayfinding	25	-.054980	.794075
Distress & detail	26	-.338201	.091048
Distress & accuracy	26	-.316838	.114783
Distress & wayfinding	25	.086649	.680456

Table 20: Spearman R and p values for the relationship between psychological factors (presence and emotions) and the three cognitive mapping scores. Significant negative correlations were found between fear and detail, surprise and detail, and shyness and accuracy (marked in bold, α=0.05).

5.3 Discussion of Results

5.3.1 Formation of a Cognitive Map

It has been established that while it is possible to form a cognitive map of virtual environment within a relatively short period of time, the maps thus formed are of mediocre quality. As the cognitive mapping literature normally reports exposure times in terms of days or months [16], substantially longer exposure times may be required to increase the quality of cognitive mapping process.

5.3.2 Effect of Display System and Map

Cognitive mapping scores were not affected by whether participants had a map or not (adding more evidence to the map controversy discussed in Section 2.4.6) and neither were they affected by whether participants used an immersive or a desktop system. This shows that ideothetic cues such as proprioception are not essential to forming a cognitive map of a virtual environment, and that visitors will form the same quality map whether they use an immersive system or a desktop system, implying that expensive immersive systems are not necessary to form a cognitive map of a VE. This finding is supported by several other studies, such as those by Koh *et al* [43], Waller [80], but is contradicted by studies by Iwata & Yoshida [35] and Chance *et al* [12]. Section 2.3.3 contains more information on these studies.

5.3.3 Effect of Other Variables

The more of the environment that a participant explored, the more detail they included on their sketch map. This is not surprising, as it follows from the fact that one cannot draw a sketch map of areas that one hasn't explored, but it is important in that it points out that if the spatial learning method is to be used successfully, visitors to VEs must be subtly encouraged to explore the entire environment. Finishing the assigned task did not affect cognitive mapping, which implies either that participants made multiple passes over the environment whether they finished the task or not, or that multiple passes do not add further knowledge of the environment. The data seems to suggest the former, although a much deeper analysis of participants' movements through the VE would be needed to confirm this.

Neither previous VR or gaming experience affected cognitive mapping, which shows firstly that there is no novelty effect (participants without previous experience did not pay more attention to their surroundings just because it was a new experience to them) and secondly that neither the quality of the cognitive maps formed of the VE nor wayfinding skills are likely to increase with increasing familiarity with VEs in general (contradicting Ruddle, Payne and Jones [64], who suggest that general lack of familiarity with VR may contribute to difficulties in navigation).

Age did not affect cognitive mapping abilities, but only adults were included in the study (ages ranged between 18 and 50). It is possible that younger children may perform differently to adults (see Section 2.3.7).

Women performed better at wayfinding than men, although no there was no gender effect on sketch map detail or accuracy. This contradicts findings by Waller, Hunt and

Knapp [81], who found that men perform better at wayfinding tasks, and Appleyard [1], who found that women made more errors in their sketch maps. Appleyard's study was conducted in 1969, however, when ability to draw a map was considered unfeminine, and he reports that there was a reluctance among women to demonstrate knowledge that they were later found to possess. Kaplan found that amongst 12 year old children, girls drew more accurate and more detailed sketch maps [39], while Holahan [33] reports that while some studies show gender differences in sketch maps, most do not.

Primary transportation method did not affect cognitive mapping scores, which agrees with findings by Carr and Schissler [11], who found no difference in cognitive mapping between drivers, passengers, and commuters. However, it is contradicted by Holahan [33] who reports that drivers produce more accurate sketch maps, and Cohen [16] who report that those who use mainly public transport tend to have the most inaccurate maps. Appleyard [1] also found that those who use public transport drew the lowest quality sketch maps, and those who drive a car produced the highest quality sketch maps.

Participants who had previously taken a technical drawing course included more detail in their sketch maps, but performed worse at wayfinding. Previous orienteering experience did not affect cognitive mapping, and neither did the mathematical abilities of the subjects (indicated by the highest level maths course they had taken). This is in contradiction with finding by Darken *et al* [18] and findings quoted by West and Morris [83], who found that maths abilities correlate positively with cognitive mapping abilities. Spatial abilities were found to affect the amount of detail included in sketch maps, but did not affect accuracy of the sketch maps or wayfinding.

The amount of presence felt by participants did not affect cognitive mapping, showing that visitors do not have to feel as though they are in a real place in order to form a cognitive map of the area. This corroborates findings by Darken *et al* who found that spatial knowledge acquisition was not affected by presence [19]. Positive emotions (such as interest and enjoyment) were not related to cognitive mapping, but fear and surprise were associated with decreased accuracy of the sketch maps and shyness with decreased amount of detail. Which is cause and which is effect cannot be determined — it is equally possible that fear and surprise were caused by an inaccurate cognitive map of the environment as that they disrupted the cognitive mapping process and thus caused less accurate cognitive maps (and similarly for shyness and detail — participants may have felt shy because they weren't absorbing much detail, or shyness may have prevented them from paying attention to their surroundings and thus prevented them from absorbing much detail). This does, however, indicate that it is possible for visitors to form an accurate cognitive map of the environment even if the data set is not one that they find interesting, and also implies that fear and surprise should be minimized in order for visitors to learn the space in full.

Chapter 6

Study 2 - Learning the Data Set

6.1 Aims

This study investigated whether, having formed a cognitive map of a virtual environment, it was possible for participants to absorb and comprehend the underlying data set not as a map, but as the actual data points and the relationships between them.

The first section of this study was an experimental study, with two hypotheses:

1. the presentation style (HMD VR system, desktop VR system, and conventional lecture) would affect the learning of the data set;

2. providing VR participants with a dynamic You-Are-Here map showing their position and orientation would affect their learning of the data set.

The first aim of the experimental part of this study was to determine whether immersive equipment (such as head-trackers and Head Mounted Displays) was necessary to learn the data set represented by the VE, or whether a keyboard-and-monitor-based desktop system was sufficient. It also compared both forms of VR presentation with a more conventional lecture-style presentation of the data. The second aim was to determine whether providing a map of the VE would help or hinder visitors with learning the underlying data set, separately to any effect it may have on the cognitive mapping process — seeing the layout of the environment (and thus the data set) may help visitors to remember the overall structure and relationships. The possibility of an interaction between presentation type and map provision was also be investigated.

In the relational section of the study, the relationships between various variables and learning scores of the VR participants were investigated, namely:

- cognitive mapping (the previous study, described in Chapter 5, showed that visitors to a virtual environment can form a cognitive map of the environment, but that the map formed is of mediocre quality — is the quality of the cognitive map related to how well visitors can learn the underlying data set?);

- psychological factors, such as emotions (enjoyment, fear, surprise, shyness, interest, distress) and the sense of presence (negative emotions may detract from learning, and positive emotions enhance learning; and the extent to which visitors to the VE feel that they are in a real place may affect the amount of data they absorb);

- the percentage of the environment that was explored (a visitor cannot learn data points if they have not traveled to them);

- the time spent in the VE (if the spatial learning method is to be useful, visitors must be able to learn the underlying data set with fairly short exposure times);

- age (to use the spatial learning method to teach the general public, an age effect would need to ruled out);

- gender (to use the spatial learning method to teach the general public, a gender effect would need to ruled out);

- spatial abilities (the spatial abilities of visitors may affect how well they can learn the underlying data set);

- percentage of time that a map of the VE was visible (a refinement of the map / no map test, which looks at how long the participant actually saw the map during the VE exposure).

The design of this study and the procedure followed in implementing it is as described in Chapter 4, as are the measures used.

6.2 Results

6.2.1 Absorbing and Comprehending the Data Set

The average learning score (as described in Section 4.7.6) for the VR participants was 12.85 with a standard deviation of 6.21. The lowest score obtained was 4, and the highest score obtained was 31, i.e. a range of 27. The maximum possible score was 36. A histogram of the VR participants' learning scores is given in Figure 33.

By comparison, the average learning score for the non-VR participants (those who attended a lecture on the data set rather than exploring it via the VE) was 23.14 with a standard deviation of 6.20. The lowest score obtained was 16, and the highest was 34, i.e. a range of 18. A histogram of the VR participants' learning scores is given in Figure 34.

6.2.2 Effect of Presentation Type and Map

To investigate the combined effect of presentation type and map provision on learning scores, a 2-way ANOVA would normally be used. However, there is no non-parametric equivalent to a 2-way ANOVA, so a replacement 3-step process has been used (see Section 4.8).

Step 1 is to determine the effect of presentation type on learning, by means of a Kruskal-Wallis ANOVA by Ranks. A significant difference in learning scores was found (H=11.95633, p=.0025, α=0.05, as shown in Table 21), and hypothesis 1 of section 6.1 is thus accepted. From this result, however, it cannot be determined which presentation group performed differently to the other two, or whether this group performed better or worse, as this requires a post-hoc test. The Sheffé test would be used when using parametric statistics, but when using non-parametric statistics, as now, a series of Mann-Whitney U tests are run between all possible combinations of groups in order to tell where the significant difference
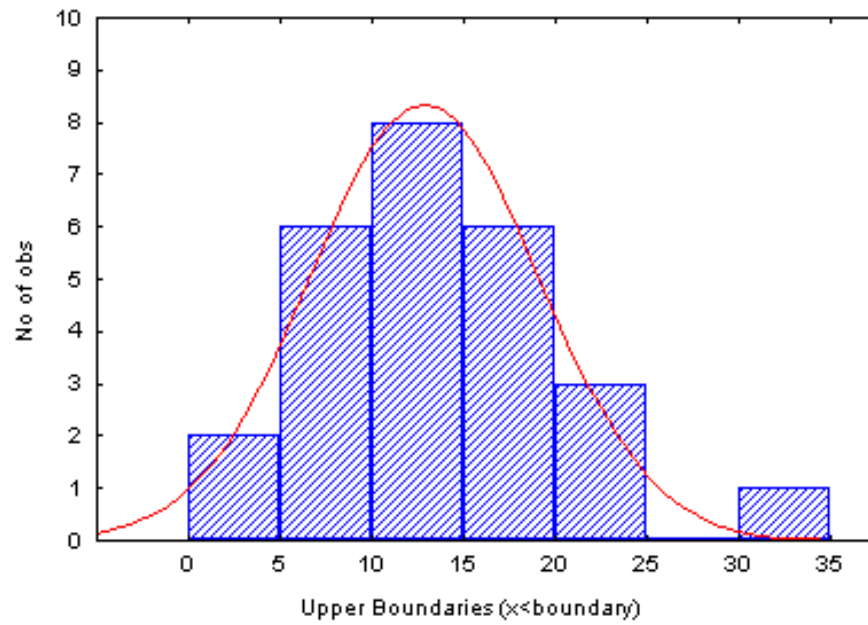
Figure 33: Histogram showing the learning scores for the VR participants. The solid curve indicates the expected normal.
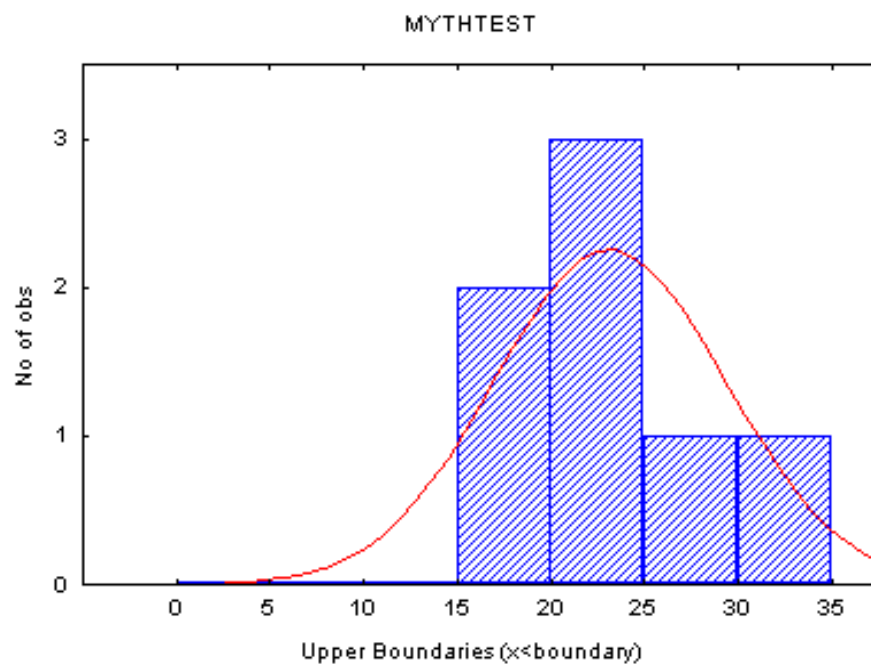
Figure 34: Histogram showing the learning scores for the non-VR participants. The solid curve indicates the expected normal.

H (2, N= 33) = 11.95633, **p =.0025**

	Valid N	Sum of Ranks
Desktop	18	287.0000
HMD	8	83.0000
Lecture	7	191.0000

Table 21: Kruskal-Wallis ANOVA by ranks testing the effect of presentation type on learning. A significant difference was found (indicated in bold, α=0.05).

	U	Z	p-level	N (lecture)	N (HMD)
Learning	0.00	3.240370	**.001195**	7	8

Table 22: Mann-Whitney U test on the difference in learning score between the lecture group and the HMD group. A significant difference was found (shown in bold, α=0.05).

lies. The direction of the difference is then determined by looking at a medians plot of the relevant groups. In this case, the difference was found to be between the HMD and lecture groups (U¡0.00, p¡.001195 α=0.05, as shown in Table 22) and between the desktop and lecture groups (U=19.00000, p=.007749, α=0.05, as shown in Table 23). From Figure 36 it is clear that the lecture group performed better than the HMD group, and similarly for the desktop group in Figure 35. The lecture group thus outperformed both the VR groups.

Step 2 in the 3-step process is to determine the effect of having a map on learning in the VR conditions, by means of a Mann-Whitney U test. The results are presented in Table 24, but no significant differences were found and hypothesis 2 of section 6.1 is thus rejected. Having map of the VE therefore did not affect learning.

Step 3 is to determine whether there were any interactions between presentation type (for HMD and Desktop) and having a map. This was done by means of a Kruskal-Wallis ANOVA by ranks. No significant results were found, as shown in Table 25.

6.2.3   Effect of Other Variables

Because all three studies were run on the same group of participants, the correlation between cognitive mapping scores and learning scores can be investigated. A Spearman R correlation between learning score and the three cognitive mapping scores (Table 26) shows that cognitive mapping and learning were somewhat related. Learning was not correlated with sketch map detail or sketch map accuracy, but was positively correlated with wayfinding scores (R=.429683, p=.032059, α=0.05). This means that those participants who performed worse at wayfinding actually learned the underlying data set better than those who had performed

	U	Z	p-level	N (lecture)	N (desktop)
Learning	19.00000	2.663001	**.007749**	7	18

Table 23: Mann-Whitney U test on the difference in learning score between the lecture group and the desktop group. A significant difference was found (shown in bold, α=0.05).
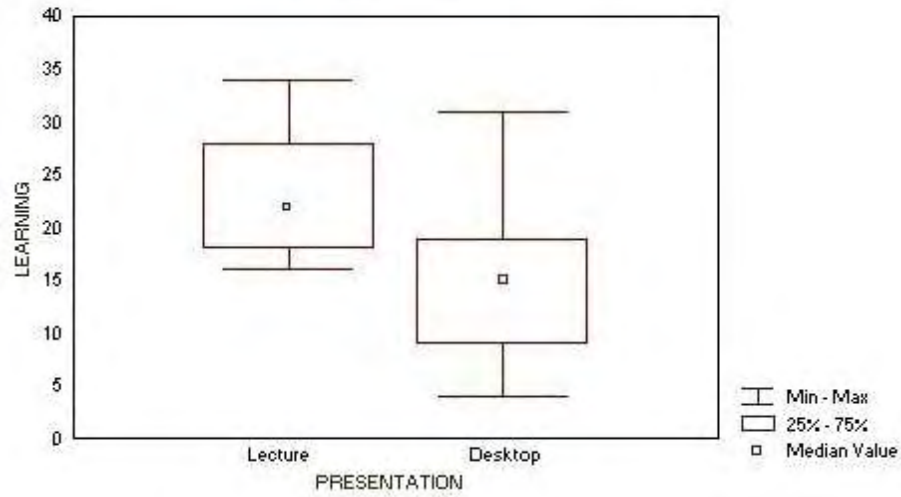
Figure 35: A box plot showing the median and quartile detail scores for the desktop and lecture groups. The median score for the lecture group is greater than that of the desktop group.
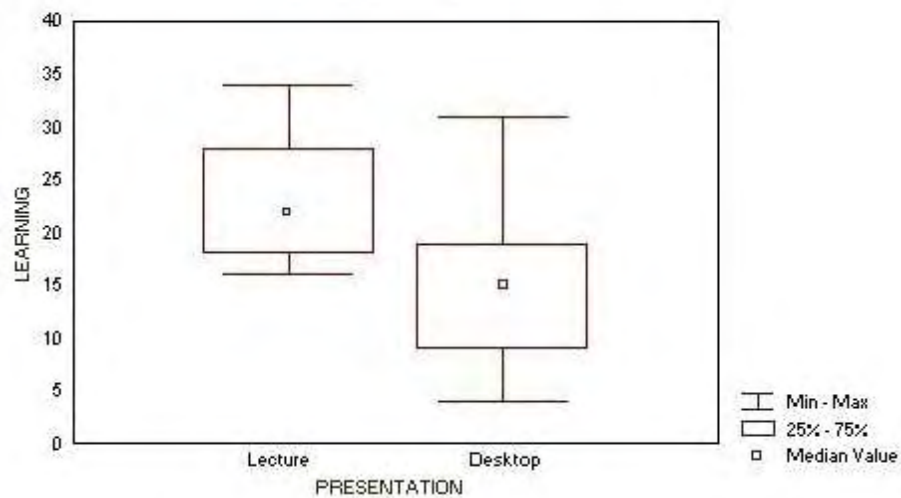
Figure 36: A box plot showing the median and quartile detail scores for the HMD and lecture groups. The median score for the lecture group is greater than that of the HMD group.

	U	Z	p-level	N (map)	N (no map)
Learning	79.50000	.256410	.797636	13	13

Table 24: Mann-Whitney U test on the effect of having a map on learning scores. No significant differences were found.

H (3, N= 26) = 4.836405, p =.1842

	Valid N	Sum of Ranks
desktop-map	9	115.5000
desktop-no-map	9	152.5000
hmd-map	4	55.0000
hmd-no-map	4	28.0000

Table 25: Kruskal-Wallis ANOVA by ranks on the combined display/map categories for learning (VR participants only). No significant effects were found.

	Valid N	Spearman R	p-level
Learning & detail	26	.250431	.217229
Learning & accuracy	26	-.112068	.585714
Learning & wayfinding	25	**.429683**	**.032059**

Table 26: Spearman R and p values for the relationship between learning and the three cognitive mapping scores. A significant positive relationship was found between learning and wayfinding (marked in bold, α=0.05), which means that better learning correlates with worse wayfinding performance.

better at wayfinding, and that learning was *negatively* correlated to wayfinding *performance*.

No emotional factors correlated to learning scores, and neither did the sense of presence (Table 27).

The percentage of the environment that was explored did not affect the learning score, nor did the percentage of time that the map was displayed. The amount of time spent in the environment did not affect the learning score either (see Table 28 for all these results).

Neither orienteering experience (Table 29), technical drawing experience (Table 30), maths skills (Table 31) nor transport method (Table 32) correlated with learning scores. ESAT scores were significantly and positively correlated with learning scores (R=.450356, p=.020959, α=0.05, as shown in Table 33). Neither age (Table 34) nor gender (Table 35) affected learning score (for both the VR and non-VR groups).

	Valid N	Spearman R	p-level
Presence & learning	26	.003095	.988028
Enjoyment & learning	26	.229372	.259674
Fear & learning	26	-.243562	.230524
Surprise & learning	26	.340068	.089165
Shyness & learning	26	-.197567	.333326
Interest & learning	26	.094514	.646050
Distress & learning	26	-.319728	.111331

Table 27: Spearman R and p values for the relationship between the psychological factors and learning. No significant relationships were found.

	Valid N	Spearman R	p-level
Percent explored & learning	26	.121120	.555594
Percent map on & learning	26	-.083000	.686873
Time & learning	26	.094305	.646784

Table 28: Spearman R and p values for the relationships between the percentage of the environment explored, the percentage of time that the map was displayed for, the amount of time spent in the VE, and learning. No significant relationships were found.

	U	Z	p-level	N (no orienteering)	N (orienteering)
Learning	51.00000	-1.52843	.126415	16	10

Table 29: Mann-Whitney U test on the effect of orienteering experience on learning scores. No significant differences were found.

	U	Z	p-level	N (no tech draw)	N (tech draw)
Learning	59.50000	-1.08044	.279952	16	10

Table 30: Mann-Whitney U test on the effect of technical drawing experience on learning scores. No significant differences were found.

H (3, N= 26) = 5.441140, p =.1422

	Valid N	Sum of Ranks
1st year university course	3	37.0000
2nd year university course	10	117.5000
3rd year university course	10	175.0000
Postgraduate maths course	3	21.5000

Table 31: Kruskal-Wallis ANOVA by ranks on the effect of maths level on learning. No significant effects were found.

H (3, N= 25) = 1.971611, p =.5783

	Valid N	Sum of Ranks
Walking	6	96.50000
Public transport	8	85.50000
Driving (driver)	6	89.00000
Driving (passenger)	6	80.00000

Table 32: Kruskal-Wallis ANOVA by ranks on the effect of transportation method on learning. No significant effects were found.

	Valid N	Spearman R	p-level
ESAT & Learning	26	**.450356**	**.020959**

Table 33: Spearman R and p values for the relationships between spatial abilities and learning. A significant relationship was found (marked in bold, α=0.05).

	Valid N	Spearman R	p-level
Age & learning	25	-.029630	.888190

Table 34: Spearman R and p values for the relationship between age and learning. No significant relationships were found.

	U	Z	p-level	N (male)	N (female)
Learning	77.00000	-.360041	.718819	14	12

Table 35: Mann-Whitney U test on the effect of gender on learning scores. No significant differences were found.

6.3 Discussion of Results

6.3.1 Absorbing and Comprehending the Data Set

This study has shown that while it is possible for visitors to a VE to absorb the data set underlying the layout of the VE, the results obtained vary greatly between individuals (the scores varied between 4 and 31, where the maximum possible score was 36). While some participants did very well on the learning test, others did very badly, and as Figure 33 shows, most VR participants scored less than 50%. By comparison, the learning scores for the non-VR participants had a range of only 18, and most participants scored more than 50% (see Figure 34).

6.3.2 Effect of Presentation Type and Map

A Kruskal-Wallis ANOVA by Ranks confirms that the lecture group performed significantly better than either of the VR groups, although within the VR groups there was no difference between those who used a desktop system and those who used an immersive system, nor was there a difference between those who had a map of the VE and those who did not. Moreover, there was no interaction between display system and having a map or not.

The lack of display system effect is encouraging, as it implies that cheaper desktop systems can be used instead of expensive and difficult to install immersive systems, and it is also encouraging that some participants did manage to learn the data set through VE exploration. The fact that the VR participants performed significantly worse than the lecture group is discouraging — for the spatial learning method to be useful in teaching the general public, it must give satisfactory results for the majority of people and must preferably give results at least as good as conventional teaching methods, such as lectures.

6.3.3 Effect of Other Variables

Participants' learning scores did not correlate with their sketch map scores, but did correlate positively with their wayfinding scores. This is perplexing, because higher wayfinding scores mean worse wayfinding performance, and thus better wayfinding performance implies worse learning performance. A possible explanation for this is that participants who performed well at wayfinding put more effort into finding their way, and were paying more attention to the structure of the VE and the features that helped them to find their way around (such as the tokens that the task focused on, which were more visible than the data posters

— especially once the tokens were activated (see Section 4.4.1)), thus leaving fewer mental resources available for them to remember the data points and their relationships.

No emotions were related to the learning scores, which means that positive emotions did not contribute to learning, but neither did negative emotions detract from it. Conversely, it also implies that learning well did not evoke positive emotions, and that learning badly did not evoke negative emotions.

Presence scores did not correlate to learning scores either, which means that feeling that one is in a real place does not help one to absorb the data underlying the layout of the place. This is a fortunate result in that it shows that sophisticated and expensive immersive systems, thought to induce presence [31, 85], are not necessary to obtain results from the spatial learning method.

Surprisingly, the amount of the environment that was explored did not affect the learning score, even though it did affect the amount of detail drawn on the sketch maps (see Section 5.3.3). This is possibly because participants remembered only a relatively small subsection of the data set even when they had explored the entire environment.

The percentage of time that the map was displayed for (in the case of those participants who were provided with a map) did not affect the learning score, confirming the finding that having a map did help with learning (Section 6.3.2).

The amount of time spent in the environment did not affect learning either, which implies that visitors do not have to spend long periods of time in the VE as shorter periods produce the same results.

Orienteering experience, technical drawing experience, maths skills and transportation method all had no effect on learning, although participants with better spatial abilities did learn the underlying data set better.

Finally, neither gender nor age affected learning, implying that the spatial learning method works equally well for all (note, however, that the youngest participant was 18 years old, and that results may well differ for children — refer to Section 2.3.7).

Chapter 7

Study 3 - Psychological Impact

7.1   Aims

This study investigated the impact of the use of a VR system (whether desktop-based or immersive) on a visitor's emotions (and whether these effects were due simply to the novelty of the situation or whether they would also be felt by those with more experience in VR) — not just in terms of their effect on learning, but also because of the effect of these emotions on the participants. It also compared the emotions evoked by the VR system used in the spatial learning method with the emotions evoked by a conventional presentation of the data set, and investigated another aspect of the visitor's psychological experience, namely that of *presence.*

The first section of this study was an experimental study, with four hypotheses:

1. presentation style (HMD VR system, desktop VR system, and conventional lecture) would affect the emotions of the participants;

2. providing VR participants with a dynamic You-Are-Here map showing their position and orientation would affect their emotions;

3. within the VR conditions, the display type would affect the sense of presence of the participants;

4. providing VR participants with a dynamic You-Are-Here map showing their position and orientation would affect their sense of presence.

The possibility of an interaction between presentation type and map provision was also be investigated.

The second section of the study was a relational study in which the relationships between various variables and the emotions and sense of presence of the VR participants were investigated, namely:

- cognitive mapping (positive emotions may help with cognitive mapping, as may the sense that one is in a real place);

- learning (positive emotions may enhance learning, as may the sense that one is in a real place);

- time spent in the VE (those who spent more time in the VE may have felt more positive emotions, or may have become despondent and felt more negative emotions; also, the sense of presence may be affected by the amount of time spent in the VE);

- the percentage of the environment explored (those who explored only a small part of the VE may have become frustrated and felt more negative emotions);

- whether or not they finished the task (those who did not manage to finish the task may have felt negative emotions);

- previous VR/gaming experience (those who have previous experience will not display any novelty effect in their emotions or sense of presence);

- gender (there may be a gender effect on emotions and presence)

- age (there may be an age effect on emotions and presence).

Finally, the relationship between the emotions felt by participants and their sense of presence was investigated.

The design of this study and the procedure followed in implementing it is as described in Chapter 4, as are the measures used.

7.2 Results

The emotions measured were enjoyment, fear, surprise, shyness, interest, and distress. Participants were asked to rate the intensity to which they felt each emotion, and scores for each emotion could vary from 0 to 15. For the VR participants, the average score for fear was 4.6538, for surprise 9.6923, for shyness 6.1538, for interest 12.9231, for enjoyment 10.69231 and for distress 6.5769. For the participants who attended the lecture instead of exploring the VE, the average score for fear was 3.57143, for surprise 9.00000, for shyness 4.85714, for interest 11.57143, for enjoyment 8.571429 and for distress 5.42857.

Presence scores could range from 0 (they didn't feel present in the VE at all) to 224 (they felt very strongly present in the VE). The average presence score was 112.7308.

7.2.1 Effect of Presentation Type and Map

To investigate the combined effect of presentation type and map provision on emotions, a 2-way ANOVA would normally be used. However, there is no non-parametric equivalent to a 2-way ANOVA, so a replacement 3-step process has been used (see Section 4.8).

Step 1 is to determine the effect of presentation type on emotions, by means of a Kruskal-Wallis ANOVA by Ranks. The complete results are shown in Table 36, but no significant differences were found between the desktop, HMD and lecture groups. Hypothesis 1 of section 7.1 is thus rejected.

Enjoyment:

H (2, N= 33) = 2.761139, p =.2515

	Valid N	Sum of Ranks
Desktop	18	331.0000
HMD	8	148.5000
Lecture	7	81.5000

Fear:

H (2, N= 33) = 1.990149, p =.3697

	Valid N	Sum of Ranks
Desktop	18	309.0000
HMD	8	157.0000
Lecture	7	95.0000

Surprise:

H (2, N= 33) = .4431117, p =.8013

	Valid N	Sum of Ranks
Desktop	18	323.5000
HMD	8	130.0000
Lecture	7	107.5000

Shyness:

H (2, N= 33) = 2.961973, p =.2274

	Valid N	Sum of Ranks
Desktop	18	298.0000
HMD	8	171.5000
Lecture	7	91.5000

Interest:

H (2, N= 33) = 4.534734, p =.1036

	Valid N	Sum of Ranks
Desktop	18	333.0000
HMD	8	156.0000
Lecture	7	72.0000

Distress:

H (2, N= 33) = 1.305822, p =.5205

	Valid N	Sum of Ranks
Desktop	18	302.0000
HMD	8	159.0000
Lecture	7	100.0000

Table 36: Kruskal-Wallis ANOVA by ranks on the effect of presentation type on the emotions of participants. No significant effects were found.

	U	Z	p-level	N (map)	N (no map)
Enjoyment	80.00000	-.23077	.817496	13	13
Fear	66.00000	.94872	.342771	13	13
Surprise	75.50000	-.46154	.644415	13	13
Shyness	71.00000	-.69231	.488749	13	13
Interest	72.50000	-.61538	.538305	13	13
Distress	62.50000	-1.12821	.259242	13	13

Table 37: Mann-Whitney U test on the effect of having a map on the emotions of the participants.

Step 2 in the 3-step process is to determine the effect of having a map on emotions (for the VR participants), by means of a Mann-Whitney U test. Again, no significant differences were found (Table 37), and hypothesis 2 of section 7.1 is thus rejected.

Step 3 is to determine whether there were any interactions between presentation type (for HMD and Desktop) and having a map. This was done by means of a Kruskal-Wallis ANOVA by ranks (see the results in Table 38). Two significant results were found, for fear (H=8.660244, p=.0342, α=0.05) and for distress (H=8.649237, p=.0344, α=0.05). However, these figures merely indicate that at least one of the four groups was significantly different to the other groups, and do not tell us which group(s) or in which direction the difference lies. For this, the Sheffé test would be used when using parametric statistics, but when using non-parametric statistics, as now, this option is not available. Instead, a series of Mann-Whitney U tests are run between all possible combinations of groups in order to tell where the significant difference lies. The direction of the difference is then determined by looking at a medians plot of the relevant groups. In this case, the difference was found to be between the desktop-map and desktop-no-map groups (U=16.50, p=.034077, α=0.05), and between the desktop-no-map and hmd-no-map groups (U=5.00, p=.044871, α=0.05) for fear (see Table 39), and between desktop-map and hmd-no-map (U=4.50, p=.037250, α=0.05), desktop-no-map and hmd-no-map (U=3.00, p=.020644, α=0.05), and hmd-map and hmd-no-map (U=.00, p=.020928, α=0.05) for distress (see Table 40). Figure 37 shows that the desktop-map group felt *more* fear than the desktop-no-map group, while Figure 38 shows that the desktop-no-map group felt *less* fear than the hmd-no-map group. For distress, Figure 39 shows that the desktop-map group felt *less* distress than the hmd-no-map group, Figure 40 shows that the desktop-no-map group felt *less* distress than the hmd-no-map group, and Figure 41 shows that the hmd-map group felt *less* distress than the hmd-no-map group. These results are summarized in Table 41.

Enjoyment:

H (3, N= 26) = 2.655541, p =.4478

	Valid N	Sum of Ranks
desktop-map	9	103.0000
desktop-no-map	9	138.0000
hmd-map	4	68.0000
hmd-no-map	4	42.0000

Fear:

H (3, N= 26) = 8.660244, **p =.0342**

	Valid N	Sum of Ranks
desktop-map	9	150.0000
desktop-no-map	9	81.5000
hmd-map	4	44.0000
hmd-no-map	4	75.5000

Surprise:

H (3, N= 26) = .5086662, p =.9170

	Valid N	Sum of Ranks
desktop-map	9	115.0000
desktop-no-map	9	134.5000
hmd-map	4	51.5000
hmd-no-map	4	50.0000

Shyness:

H (3, N= 26) = 6.869970, p =.0762

	Valid N	Sum of Ranks
desktop-map	9	121.0000
desktop-no-map	9	100.5000
hmd-map	4	41.0000
hmd-no-map	4	88.5000

Interest:

H (3, N= 26) = .4382002, p =.9322

	Valid N	Sum of Ranks
desktop-map	9	112.5000
desktop-no-map	9	127.5000
hmd-map	4	51.0000
hmd-no-map	4	60.0000

Distress:

H (3, N= 26) = 8.649237, **p =.0344**

	Valid N	Sum of Ranks
desktop-map	9	123.0000
desktop-no-map	9	107.0000
hmd-map	4	30.5000
hmd-no-map	4	90.5000

Table 38: Kruskal-Wallis ANOVA by ranks on the combined display/map categories for emotions felt by participants. Significant effects are marked in bold (α=0.05).

	U	Z	p-level	N (desktop-map)	N (desktop-no-map)
Fear	6.50000	2.119252	**.034077**	9	9

	U	Z	p-level	N (desktop-no-map)	N (hmd-no-map)
Fear	5.000000	-2.00594	**.044871**	9	4

Table 39: Mann-Whitney U tests to pinpoint the difference in fear between the four VR groups. Only the two significant results are shown (indicated in bold, α=0.05).
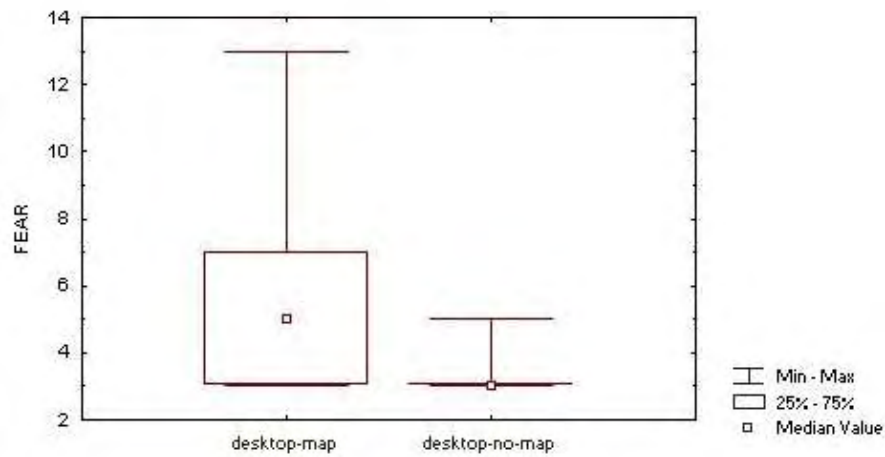
Figure 37: A box plot showing the median and quartile fear scores for the desktop-map and desktop-no-map groups. Those with a map felt more fear.

	U	Z	p-level	N (desktop-map)	N (hmd-no-map)
Distress	4.500000	-2.08310	**.037250**	9	4

	U	Z	p-level	N (desktop-no-map)	N (hmd-no-map)
Distress	3.000000	-2.31455	**.020644**	9	4

	U	Z	p-level	N (hmd-map)	N (hmd-no-map)
Distress	0.00	-2.30940	**.020928**	4	4

Table 40: Mann-Whitney U tests to pinpoint the difference in distress between the four VR groups. Only the three significant results are shown (indicated in bold, α=0.05).
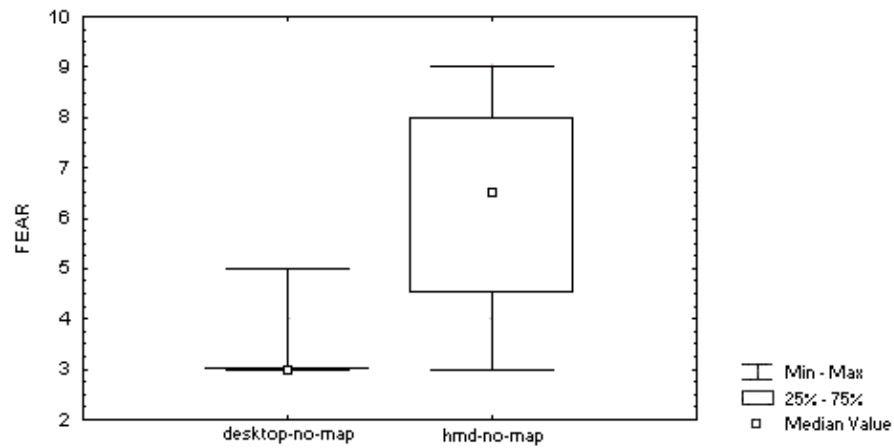
Figure 38: A box plot showing the median and quartile fear scores for the desktop-no-map and hmd-no-map groups. Those with using the HMD felt more fear.
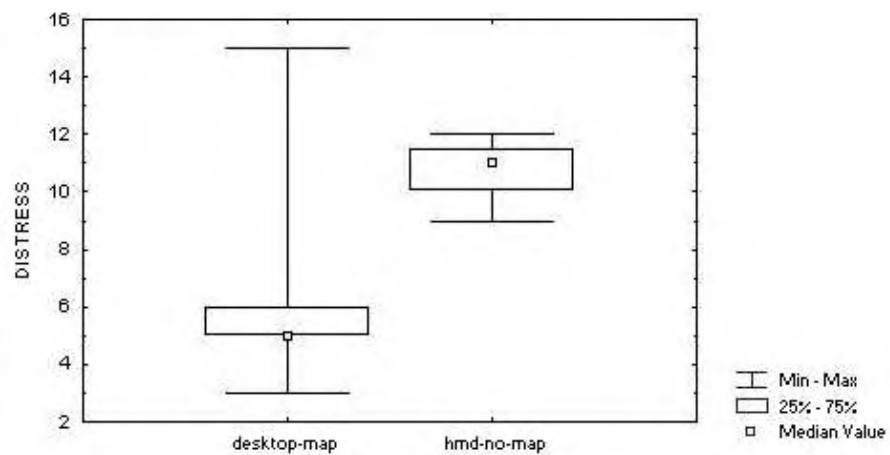
Figure 39: A box plot showing the median and quartile distress scores for the desktop-map and hmd-no-map groups. The hmd-no-map group were significantly more distressed.
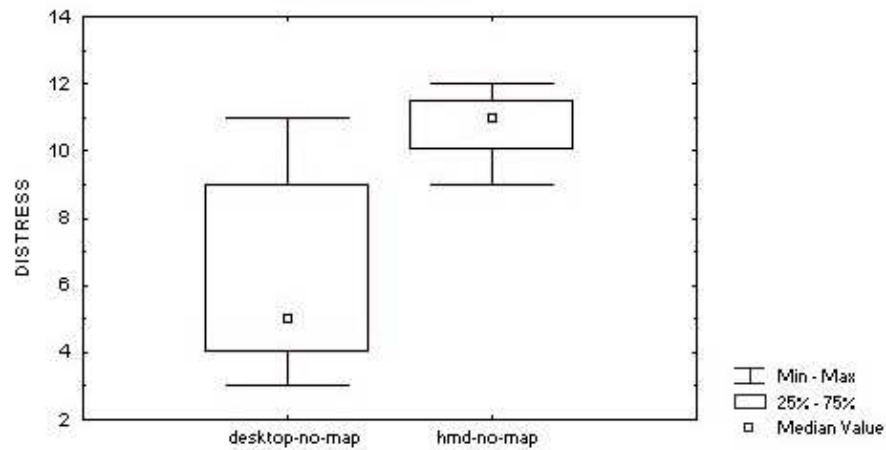
Figure 40: A box plot showing the median and quartile distress scores for the desktop-no-map and hmd-no-map groups. Those using the HMD felt significantly more distressed.
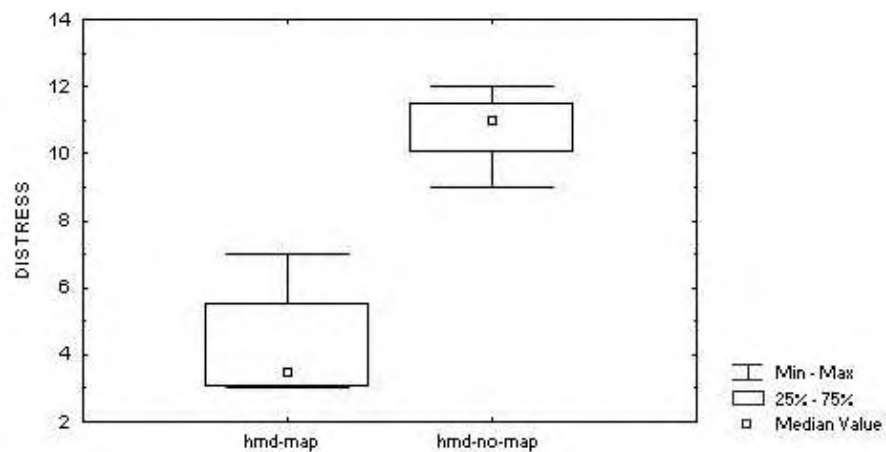
Figure 41: A box plot showing the median and quartile distress scores for the hmd-map and hmd-no-map groups. Those with a map felt significantly less distressed.

	Fear	Distress
desktop-map vs desktop-no-map	>	
desktop-map vs hmd-map		
desktop-map vs hmd-no-map		<
desktop-no-map vs hmd-map		
desktop-no-map vs hmd-no-map	<	<
hmd-map vs hmd-no-map		<

Table 41: A summary of the interaction between map and display type for those emotions where a significant difference was found. A < indicates that participants in the first condition felt significantly less of the relevant emotion than those in the second condition, while a > indicates that participants in the first condition felt significantly more of the relevant emotion than those in the second condition.

	U	Z	p-level	N (desktop)	N (HMD)
Presence	52.00000	1.111111	.266529	18	8

Table 42: Mann-Whitney U test on the difference in presence between the desktop and HMD groups. No significant difference was found.

To determine the effect of display system and map on the sense of presence, the same 3-step process is followed.

Step 1 is to determine the effect of display system on presence, by means of a Mann-Whitney U test. The complete results are shown in Table 42, but no significant results were found, and hypothesis 3 of section 7.1 is thus rejected.

Step 2 is to determine the effect of having a map on presence, also by means of a Mann-Whitney U test. The results are presented in Table 43, but no significant differences were found. Hypothesis 4 of section 7.1 is thus rejected.

Step 3 is to determine whether there were any interactions between display type and having a map. This was done by means of a Kruskal-Wallis ANOVA by ranks. Again, no significant results were found, as shown in Table 44.

7.2.2 Effect of Other Variables

Cognitive mapping scores were found to be related to some emotional aspects. Fear and surprise were negatively correlated with sketch maps accuracy (R=-.404995, p=.040131 for fear; R=-.556910, p=.003125 for surprise) and shyness negatively correlated with sketch map detail (R=-.493673, p=.010376). No other relationships between psychological factors

	U	Z	p-level	N (map)	N (no map)
Presence	79.50000	.256410	.797636	13	13

Table 43: Mann-Whitney U test on the difference in presence between the map and no-map groups. No significant difference was found.

H (3, N= 26) = 5.104796, p =.1643

	Valid N	Sum of Ranks
desktop-map	9	110.5000
desktop-no-map	9	152.5000
hmd-map	4	60.0000
hmd-no-map	4	28.0000

Table 44: Kruskal-Wallis ANOVA by ranks on the combined display/map categories for presence. No significant effects were found.

Figure 42: A box plot showing the median and quartile shyness scores for those who did and did not finish the task. Those who finished felt less shyness.

and cognitive mapping aspects were found, including the sense of presence (see Table 45).

Learning scores did not correlate with any emotional factors, and neither did they correlate with the sense of presence (Table 46).

The amount of time spent in the environment correlated positively with shyness (R=.437397, p=.025448, α=0.05), although not with any other emotions or the sense of presence (see Table 47).

The percentage of the environment explored was negatively correlated with distress (R=-.427834, p=.029235, α=0.05), but did not correlate with any other emotions or the sense of presence (Table48).

Finishing the task was not related to presence, but was found to be related to shyness (U=24.00000, p=.028467, α=0.05) and distress (U=21.50000, p=.019134, α=0.05). These results are shown in Table 49, but are not sufficient to indicate the direction of the relationship. Figure 42 shows that finishing the task was resulted in less shyness, while Figure 43 shows that not finishing the task resulted in more distress.

The only emotion that was found to correlate with previous VR/gaming experience was

	Valid N	Spearman R	p-level
Presence & detail	26	.293957	.144949
Presence & accuracy	26	.071460	.728667
Presence & wayfinding	25	-.020019	.924330
Enjoyment & detail	26	.049232	.811237
Enjoyment & accuracy	26	-.202989	.319958
Enjoyment & wayfinding	25	-.250343	.227444
Fear & detail	26	-.252398	.213520
Fear & accuracy	26	**-.404995**	**.040131**
Fear & wayfinding	25	-.089613	.670118
Surprise & detail	26	-.003822	.985218
Surprise & accuracy	26	**-.556910**	**.003125**
Surprise & wayfinding	25	.120088	.567474
Shyness & detail	26	**-.493673**	**.010376**
Shyness & accuracy	26	-.367625	.064654
Shyness & wayfinding	25	.037036	.860484
Interest & detail	26	.160716	.432857
Interest & accuracy	26	-.148283	.469728
Interest & wayfinding	25	-.054980	.794075
Distress & detail	26	-.338201	.091048
Distress & accuracy	26	-.316838	.114783
Distress & wayfinding	25	.086649	.680456

Table 45: Spearman R and p values for the relationship between psychological factors (presence and emotions) and the three cognitive mapping scores. Significant negative correlations were found between fear and detail, surprise and detail, and shyness and accuracy (marked in bold, α=0.05).

	Valid N	Spearman R	p-level
Presence & learning	26	.003095	.988028
Enjoyment & learning	26	.229372	.259674
Fear & learning	26	-.243562	.230524
Surprise & learning	26	.340068	.089165
Shyness & learning	26	-.197567	.333326
Interest & learning	26	.094514	.646050
Distress & learning	26	-.319728	.111331

Table 46: Spearman R and p values for the relationship between the psychological factors and learning. No significant relationships were found.

	Valid N	Spearman R	p-level
Presence & time	26	-.273257	.176795
Enjoyment & time	26	.165979	.417737
Fear & time	26	.01830	.929291
Surprise & time	26	.353920	.076094
Shyness & time	26	**.437397**	**.025448**
Interest & time	26	.166519	.416205
Distress & time	26	.303858	.131268

Table 47: Spearman R and p values for the relationship between the psychological factors and time spent in the VE. A significant relationship was found between exposure time and shyness (marked in bold, α=0.05).

	Valid N	Spearman R	p-level
Presence & percent explored	26	.141990	.488988
Enjoyment & percent explored	26	-.162965	.426363
Fear & percent explored	26	-.142717	.486741
Surprise & percent explored	26	-.290767	.149567
Shyness & percent explored	26	-.311902	.120863
Interest & percent explored	26	-.261892	.196214
Distress & percent explored	26	**-.427834**	**.029235**

Table 48: Spearman R and p values for the relationship between the psychological factors and percentage of the environment explored. A significant relationship was found between percent explored and distress (marked in bold, α=0.05).

	U	Z	p-level	N (not finished)	N (finished)
Presence	49.50000	-.639010	.522821	6	20
Enjoyment	35.00000	1.521452	.128156	6	20
Fear	58.50000	-.091287	.927265	6	20
Surprise	30.50000	1.795313	.072613	6	20
Shyness	24.00000	2.190890	**.028467**	6	20
Interest	59.00000	.060858	.951473	6	20
Distress	21.50000	2.343035	**.019134**	6	20

Table 49: Mann-Whitney U test on the difference in psychological effects between those participants who finished the task and those who didn't. Significant differences were found for shyness and distress (marked in bold, α=0.05).
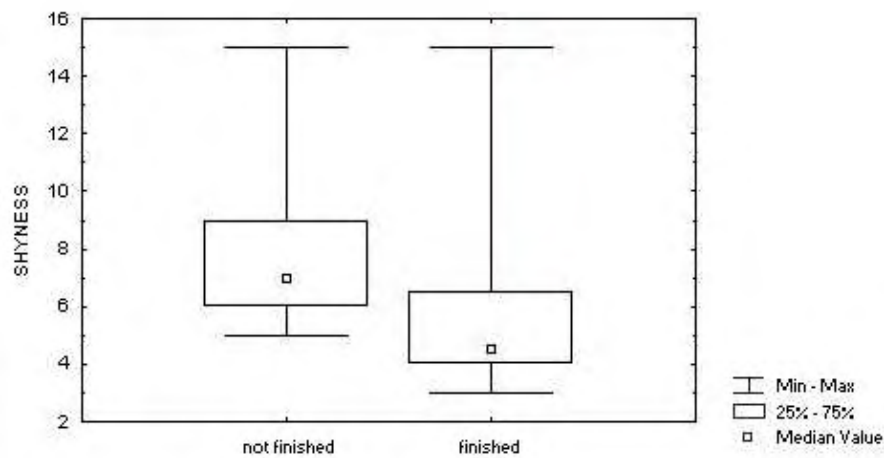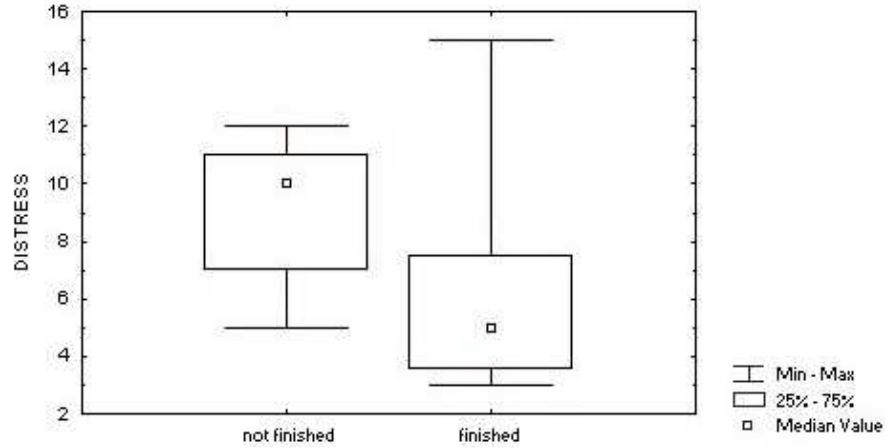
Figure 43: A box plot showing the median and quartile distress scores for those who did and did not finish the task. Those who finished felt less distressed.

	Valid N	Spearman R	p-level
Presence & gaming experience	26	**.507289**	**.008165**
Enjoyment & gaming experience	26	.184515	.3668683
Fear & gaming experience	26	**-.487790**	**.011476**
Surprise & gaming experience	26	.098686	.631492
Shyness & gaming experience	26	-.111185	.588693
Interest & gaming experience	26	-.189683	.353359
Distress & gaming experience	26	-.366018	.065922

Table 50: Spearman R and p values for the relationship between the psychological factors and previous gaming/VR experience. A significant relationship was found between gaming experience and fear (marked in bold, α=0.05).

fear (R=-.487790, p=.011476, α=0.05, as shown in Table 50). This was a negative correlation, showing that novices felt more fear than those with previous experience. Presence correlated positively with gaming experience, showing that presence is most likely not a novelty effect (also shown in Table 50).

No relationships were found between emotions and age or presence and age (Table 51), nor were any relationships found between gender and emotions or presence (Table 52.

Finally, the sense of presence itself correlated with enjoyment, although not with any of the other emotions (Table 53).

	Valid N	Spearman R	p-level
Presence & age	25	-.055482	.792236
Enjoyment & age	25	-.044284	.833521
Fear & age	25	.105146	.616929
Surprise & age	25	.163493	.434869
Shyness & age	25	-.022316	.915679
Interest & age	25	.173400	.407150
Distress & age	25	.296080	.150700

Table 51: Spearman R and p values for the relationship between the psychological factors and age. No significant relationships were found.

	U	Z	p-level	N (male)	N (female)
Presence	74.50000	.488627	.625109	14	12
Enjoyment	68.50000	-.797234	.425321	14	12
Fear	83.00000	-.051434	.958980	14	12
Surprise	71.50000	-.642931	.520274	14	12
Shyness	78.50000	-.282889	.777263	14	12
Interest	80.50000	.180021	.857138	14	12
Distress	74.00000	-.514345	.607015	14	12

Table 52: Mann-Whitney U test on the difference in psychological effects between genders. No significant differences were found.

	Valid N	Spearman R	p-level
Enjoyment & presence	26	**.418296**	**.033452**
Fear & presence	26	-.218857	.282750
Surprise & presence	26	.310324	.122854
Shyness & presence	26	-.185904	.363208
Interest & presence	26	.174994	.392530
Distress & presence	26	-.291661	.148262

Table 53: Spearman R and p values for the relationship between the sense of presence and emotions. A significant relationship was found between presence and enjoyment (marked in bold, α=0.05).

7.3 Discussion of Results

7.3.1 Effect of Presentation Type and Map

The presentation type did not affect participants' emotions. This means that the VR experiences were not more interesting or enjoyable than the lecture experience, which is surprising. Fortunately, though, the VR experiences did not induce any more negative emotions than the lecture, either, which means that VR participants were at least not more afraid and distressed than those who attended the lecture.

Having a map did not affect emotions, either, showing those who did not have a map did not feel more distressed than those who did.

There was, however, an interaction between VR display type and having a map. Those who used the desktop system without a map felt more fear than those using the desktop system with a map, and out of those who did not have a map the participants using the HMD felt more fear than those using the desktop system. It is therefore possible to surmise that not having a map left participants feeling lost and helpless, and that this effect was more pronounced with those using the HMD as contact with the real world is lessened while wearing earphones and an HMD. In addition, those participants who used the desktop system with a map felt far less distress than those who used the desktop system without a map (the two extremes of the confidence spectrum), those who used the desktop without a map felt more distressed than those who used the HMD without a map (again, the HMD isolates the user from the real world, magnifying feelings of being lost and alone), and within those who used the HMD those with a map felt less distressed than those who did not have a map (the map providing a reassuring anchor).

The display type (i.e. HMD or desktop) did not affect the sense of presence, however, and neither did having a map. In addition, there were no interactions between display type and having a map for presence — in other words, there was no difference in presence between the combined conditions of desktop-map, desktop-no-map, hmd-map, and hmd-no-map.

7.3.2 Effect of Other Variables

Fear and surprise were negatively correlated with accuracy of the sketch maps, while shyness was negatively correlated with the amount of detail. While these negative correlation cannot justify reducing the amount of shyness, fear and surprise in order to increase cognitive mapping, the spatial learning method can be more easily considered to be successful if users do not experience negative emotions during the session.

Presence was not related to any of the cognitive mapping scores or learning scores, which implies that expensive, immersive, presence-inducing systems are not necessary for the success of the spatial learning method. This corroborates findings by Darken *et al* who found that spatial knowledge acquisition was not affected by presence [19].

Learning scores were not associated with any emotions, which is fortunate in that if participants do experience negative emotions, it will not affect the absorption and comprehension of the data set underlying the layout of the VE.

The amount of shyness felt by participants was found to be correlated with the amount of time they spent in the VE. It is not possible to tell, however, whether participants felt shy because they were taking a longer time to finish the task, or whether they took longer to finish the task because they were shy and hesitant.

Increased shyness was also found to be associated with not finishing the task. Again, it is not statistically possible to tell whether participants felt shy because they could not finish the task, or whether increasing shyness prevented them from finishing the task.

Distress was associated with the amount of the environment explored. Again, it is impossible to determine statistically whether distress caused less of the environment to be explored, or whether participants felt distressed because they were stuck in a small part of the VE. However, the latter is more likely than the former, because otherwise one would expect a correlation with time (i.e. if participants were distressed because they couldn't find the rest of the tokens used in the task because they only discovered a small portion of the environment, they would surely give up and ask to leave (as some participants did, but not in statistically significant quantities).

Participants with previous gaming experience felt less fear and more presence than novices to 3D virtual environments, but did not feel less enjoyment or surprise, showing that the emotional reactions to the VR experience are stable and unlikely to be merely a novelty effect.

Presence itself correlated with enjoyment, although not with any of the other emotions. A possible explanation for this is that participants who were enjoying themselves let go of their physical surroundings more easily and were more willing to accept the reality of the virtual place.

Chapter 8

Conclusion

The aim of this dissertation was to introduce the spatial learning method as an alternative to conventional ways of presenting data sets, and to present a set of exploratory studies designed to investigate whether or not this new teaching method is feasible. The ultimate aim was to either recommend the spatial learning method and provide recommendations as to what equipment would be needed and what other factors need to be considered, or to eliminate the method as a feasible learning option.

8.1   Conclusions

Study 1 showed in Chapter 5 that while it is possible to form a cognitive map of a virtual environment within a relatively short period of time, the cognitive maps thus formed are not of very good quality in terms of wayfinding ability and sketch map accuracy and detail. If the theory behind the spatial learning method is correct, then participants should not have been able to learn the data set — and indeed, Study 2 (Chapter 6) showed that while some participants did very well on the test on the data set, most scored less than 50%. However, Study 2 also showed that the ability to learn the data set underlying the structure of the VE was not significantly correlated to sketch map scores. This may point towards a flaw in the use of sketch maps — it is possible that participants knew more detail about the environment than they thought was necessary to include on their sketch map (such as the names and locations of the gods).

With regards to learning ability, Study 2 showed that the learning of the underlying data set varied greatly between participants, with some remembering almost all of the data points and the relationships between them while others could barely answer the most rudimentary questions about the data set. Discouragingly, participants who attended the lecture presentation performed significantly better on the learning test than the VR participants, regardless of whether the VR participants used the HMD (with all the ideothetic cues available) or the desktop system, or whether they had a map to show them the structure of the VE or not. This means that on average the spatial learning method is not even at least as effective as conventional methods. However, the spatial learning method was not totally ineffective, as some VR participants did do very well on the mythology test.

The findings in Study 3 (Chapter 7) were surprising in that they showed that the VR experience was no more enjoyable or interesting than the lecture, and neither were

the negative emotions felt any less strongly. This seems counter-intuitive, but demand characteristics may have played a role — the lecture participants were paid volunteers and thus may have overrated their enjoyment and interest [56]. Alternatively, the problem may lie in the fact that emotions were rated using a self-report scale. A common problem with self-report scales is that there is no clear baseline against which to rate highly subjective phenomena (such as emotions), so, for example, the lecture group may have rated their lecture enjoyment as lower if they had previously experienced the VR presentation and could compare the two experiences. This theory is supported by anecdotal evidence, as many of the VR participants spontaneously remarked that they had fun during the session.

An interesting result was that those participants who performed better at wayfinding actually performed worse in the learning test. It is possible that these participants were so focused on the task of finding the tokens that they stored only the information needed to return to any specific token, and did not pay enough attention to their surroundings and especially not to the posters in the rooms which indicated which room represented which god. They were therefore able to draw a reasonable sketch map of the area, but without reference to the gods' locations, and without being able to answer the questions about the gods and their relationships.

Another interesting and useful result was that using the immersive, head-tracked VR system did not lead to better cognitive mapping (despite the extra ideothetic cues), did not lead to better learning scores, and in certain cases (for example, when no map was available) caused more fear and distress than the desktop system. In addition, more participants felt dizzy and nauseous while using the HMD than the desktop, even though there were fewer participants using the HMD, showing that the immersive system induces simulator sickness more readily than the desktop system (this data was analyzed using a χ^2 contingency table, and was found to be statistically significant — χ^2=5.610, p=.0179). Using the HMD therefore is not only not beneficial, it may actually be detrimental. This is a fortunate finding in that immersive systems are more expensive than desktop systems, are less portable than desktop systems, and are more difficult to set up for each individual user. If the spatial learning method is to be used at exhibitions and in other temporary displays, a desktop system would be far easier to use.

Another encouraging result was that the amount of time spent in the VE did not affect cognitive mapping or learning. This result implies that exposure times can be kept within the lower range of those in this study without affecting the amount of learning.

Study 1 and Study 2 showed that having a map did not improve either cognitive mapping or learning, which adds more evidence to the map/no map controversy discussed in Section 2.4.6. This finding also helps to improve the practicability of non-linear spatiality, as one of the difficulties with implementing and testing the use of non-linear spatiality in the spatial learning method (Section 3.5.3) is that these environments are topologically impossible in the real world — which means that it is very difficult to draw a map of the environment to provide to participants.

Study 3 showed that the sense of presence was not related to cognitive mapping or learning. This is again fortunate in that the systems which are thought result in increased presence are generally the immersive systems, which are more expensive than simple desktop systems (although in this study, the type of VR system used did not affect presence at all).

8.2 Recommendations

With the current findings it is difficult to recommend the use of the spatial learning method as a teaching tool. However, some of the results are encouraging, and it may be possible to improve the method at least to the point where it could be used a teaching aid to supplement conventional methods rather than replacing them.

To do this, the learning rates would need to be improved, and while there will always be some range in the learning results, the extreme range found in this study would need to be reduced. Increasing the amount of time spent in the VE would probably not improve learning scores, as time spent and learning were found to be not associated. Equally, finding a method to encourage participants to explore more of the environment would not be effective as the percentage of the environment explored did not relate to learning, either.

Wayfinding scores were inversely related to learning scores, and a possible explanation for this may point to a way to improve wayfinding — it is possible that participants focused on the items related to the task, and used *these* items as landmarks and to orient themselves, thus largely ignoring the items representing the data set. This explanation is supported by real-time observation of the participants, which shows that the participants start off examining the data set items, but rapidly deteriorate into scanning rooms purely for the task items and rarely look at the data set items at all. By integrating the task items and the data items, the data itself becomes the focus of the task, and as this is what participants seem to be paying the most attention to, recall of the data set items should be improved, thus improving the overall learning scores. It is important to bear in mind, however, that while this hypothesis was suggested by the data gathered in this study, there is insufficient information available to claim that this *is* the explanation for the inverse correlation between wayfinding performance and learning, and further studies would have be conducted to test whether integrating the task and the data items would actually improve learning.

Another possible contributing factor to the low learning scores is suggested by anecdotal evidence to the effect that the names of the data items were difficult to enunciate and thus to remember. Using a data set with more familiar items may help participants to remember the items and their names more easily, and thus also help to improve learning [2].

To make the spatial learning experience a positive one, certain points must be borne in mind:

1. if a desktop system is to be used, visitors should be provided with a map of the VE (when using a desktop system, participants felt more afraid if they were not provided with a map of the VE than if they were given a map);

2. if a map cannot be provided, a desktop system is still recommended over a HMD system (out of participants did not have a map of the VE, those using HMDs were more afraid than those using desktop systems);

3. a desktop system with a map causes less distress than a HMD system without a map;

4. HMD-based systems induce simulator sickness far more often than desktop-based systems;

5. visitors with poor spatial abilities may not learn as well using the spatial learning method as others;

6. Novices to the VR experience feel more fear than those who regularly play VR-style games, and so they may need more reassurance.

The equipment recommended for the spatial learning method is a desktop system rather than an immersive one, due to the points mentioned above and the conclusions drawn above. While providing a map is not necessary for learning purposes, it is recommended to make the experience more positive.

8.3   Future Work

While this initial investigation into the spatial learning method produced some discouraging results by showing that teaching via a lecture produced better results than the spatial learning method, there is enough evidence to suggest that some participants did learn the data relationships from the VE. There are thus two major directions which further research should take. The first is to repeat the three studies reported here, but with the improvements noted above (such as avoiding self-report scales for emotions, and integrating the task with the data set). Such replications should also focus on the findings that contradicted established results (such as those relating to the effect of providing a map of the environment on cognitive mapping), if possible finding the measures used in previous studies and incorporating them into the studies under consideration. This would help to determine whether the established findings are upheld when their own measures are used, and would also establish the relationship (if any) between those measures and the ones used in these studies.

The second direction that future research into the spatial learning method should take is to conduct studies focusing on the issues not investigated in the three studies presented in this dissertation. For example, different data sets could be compared as to "learnability", such as the different types of data sets described in Section 3.4. Further research should also be conducted into the concept of non-linear spatiality, which was not investigated at all in these studies. Some interesting areas for further research are designing non-linear environments, programming VR engines to deal with the necessary features (such as teleports which are transparent and change destination depending on the path taken to reach them), creating editing tools to help in the planning and creation of these environments, and exploring techniques to create intuitive maps of these types of environments.

Bibliography

[1] D. Appleyard. Styles and methods of structuring a city. *Environment and Behavior*, 2:100 – 116, 1970.

[2] A. Baddeley, editor. *Human Memory: Theory and Practice, Revised Edition*. Allyn & Bacon, 1997.

[3] A. Van Ballegooij and A. Eliëns. Navigation by query in virtual worlds. In *Web 3D 2000*, pages 77–83, 2000.

[4] R. B. Bechtel. *Environment and Behaviour: An Introduction*. SAGE Publications, 1997.

[5] R. Beck. *Spatial Meaning and the Properties of the Environment*, pages 134–141. In Proshansky et al. [58], 1970.

[6] F. S. Belleza. The spatial arrangement mnemonic. *Journal of Experimental Psychology: Applied*, 5(1):54 – 75, 1999.

[7] M. Bonnes and G. Secchiaroli. *Environmental Psychology: A Psycho-social Introduction*. SAGE Publications, 1995.

[8] D. Bowman, E.T. Davis, L.F. Hodges, and A.N. Badre. Maintaining spatial orientation during travel in an immersive virtual environment. *Presence: Teleoperators and Virtual Environments*, 8(6):618–631, December 1999.

[9] R. Briggs. *Methodologies for the measurement of cognitive distance*, chapter 30, pages 325–334. In Moore and Golledge [53], 1976.

[10] M. Cadwallader. *Cognitive Distance in Intraurban Space*, chapter 29, pages 316–324. In Moore and Golledge [53], 1976.

[11] S. Carr and D. Schissler. The city as a trip: perceptual selection and memory in the view from the road. *Environment and Behaviour*, 1:7–36, 1969.

[12] S. Chance, F. Gaunet, A. Beall, and J. Loomis. Locomotion mode affects the updating of objects encountered during travel: The contribution of vestibular and proprioceptive inputs to path integration. *Presence: Teleoperators and Virtual Environments*, 7(2):168–178, April 1998.

[13] W. G. Chase. *Visual Information Processing*, volume 2, chapter 28. Wiley, 1986.

[14] J.L. Chen and K.M. Stanney. A theoretical model of wayfinding in virtual environments: Proposed strategies for navigational aiding. *Presence: Teleoperators and Virtual Environments*, 8(6), December 1999.

[15] R. Cohen, editor. *The Development of Spatial Cognition*. Lawrence Erlbaum Associates, 1985.

[16] S. Cohen and R. Cohen. *The Role of Activity in Spatial Cognition*, chapter 8, pages 199–221. In Cohen [15], 1985.

[17] H. Colle and G. Reid. The room effect: Metric spatial knowledge of local and separated regions. *Presence: Teleoperators and Virtual Environments*, 7(2):116–128, April 1998.

[18] R. Darken, T. Allard, and L. Achille. Spatial orientation and wayfinding in large-scale virtual spaces: An introduction. *Presence: Teleoperators and Virtual Environments*, 7(2):101–107, April 1998.

[19] R. Darken, D. Bernatovich, J. Lawson, and B. Peterson. Quantitative measures of presence in virtual environments: The roles of attention and spatial comprehension. *CyberPsychology and Behavior*, 2(4):337–347, August 1999.

[20] R. Darken and H. Cevik. Map usage in virtual environments: Orientation issues. In *Proceedings of IEEE Virtual Reality 99*, pages 133–140, 1999.

[21] R. Darken and J. Sibert. Navigating large virtual spaces. *International Journal of Human Computer Interaction*, 8(1):49–71, January - March 1996.

[22] R.P. Darken, T. Allard, and L. Achille. Spatial orientation and wayfinding in large-scale virtual environments - guest editor's introduction. *Presence: Teleoperators and Virtual Environments*, 8(6):iii–vi, December 1999.

[23] A. Devlin. *The Small Town Cognitive Map: Adjusting to a New Environment*, chapter 4, pages 58–66. In Moore and Golledge [53], 1976.

[24] R. Downs. *Cognitive Mapping and Information Processing: A Commentary*, chapter 5, pages 67–70. In Moore and Golledge [53], 1976.

[25] T. FurnessIII, W. Winn, and R. Yu. The impact of three dimensional immersive virtual environments on modern pedagogy: Global change, VR and learning. Workshops on The Impact of Three Dimensional Immersive Virtual Environments on Modern Pedagogy in Seattle, Washington, and Loughborough, England, May/June 1997.

[26] T. Gärling, A. Böök, and E. Lindberg. *Adults' Memory Representations of the Spatial Properties of the their Everyday Physical Environment*, chapter 6, pages 141–184. In Cohen [15], 1985.

[27] J. D. Gibbons. *Nonparametric methods for Quantitative Analysis*. Holt, Rinehart and Winston, 1976.

[28] S. Golbeck. *Spatial Cognition as a Function of Environmental Characteristics*, chapter 9, pages 225–275. In Cohen [15], 1985.

[29] R. Golledge. *Methods and Methodological Issues in Environmental Cognition Research*, chapter 28, pages 300–313. In Moore and Golledge [53], 1976.

[30] Li Gong, Roger Needham, and Raphael Yahalom. Reasoning about belief in cryptographic protocols. In Deborah Cooper and Teresa Lunt, editors, *Proceedings 1990 IEEE Symposium on Research in Security and Privacy*, pages 234–248. IEEE Computer Society, 1990.

[31] C. Hendrix and W. Barfield. Presence in virtual environments as a function of visual and auditory cues. In *Proceedings of the Virtual Reality Annual International Symposium '95*, pages 74–82, 1995.

[32] I. Herman, G. Melancon, and M.S. Marshall. Graph visualisation and navigation in information visualisation: A survey. *IEE Transactions on Visualisation and Computer Graphics*, 6(1):24–43, 2000.

[33] C. J. Holahan. *Environmental Psychology*. Random House, 1982.

[34] C.J. Holahan and M. Bonnes Dobrowolny. Cognitive and behavioral correlates of the spatial environment. *Environment and Behavior*, 10(3):317–334, September 1978.

[35] H. Iwata and Y. Yoshida. Path reproduction tests using a torus treadmill. *Presence: Teleoperators and Virtual Environments*, 8(6):587–597, December 1999.

[36] C. E. Izard. *Psychology of Emotions*. Plenum Press, 1991.

[37] R. Jackson, W. Taylor, and W. Winn. Peer Collaboration and Virtual Environments: A Preliminary Investigation of Multi-Participant Virtual Reality Applied in Science Education. http://www.hitl.washington.edu/publications/r-98-10, October 1998.

[38] C. Johns, D. Nuñez, M. Daya, D. Sellars, J. Casanueva, and E. Blake. *The Interaction Between Individuals' Immersive Tendencies and the Sensation of Presence in a Virtual Environment*, chapter 7. In Mulder and van Liere [54], 2000. Proceedings of the Eurographics Workshop in Amsterdam, The Netherlands, June 1-2, 2000.

[39] R. Kaplan. *Way-Finding in the Natural Environment*, chapter 3, pages 46–57. In Moore and Golledge [53], 1976.

[40] S. Kaplan. *Adaptation, Structure, and Knowledge*, chapter 2, pages 32–45. In Moore and Golledge [53], 1976.

[41] R. M. Kitchin. Cognitive maps: What are they and why study them? *Journal of Environmental Psychology*, 14:1 – 19, 1994.

[42] F. Klett and D. Alpaugh. *Environmental Learning and Large-Scale Environments*, chapter 10, pages 121–130. In Moore and Golledge [53], 1976.

[43] G. Koh, T.E. von Wiegand, R. Garnett, N. Durlach, and B. Shinn-Cunningham. Use of virtual environments for acquiring configurational knowledge about specific real-world space I: Preliminary experiment. *Presence: Teleoperators and Virtual Environments*, 8(6):632–656, December 1999.

[44] J. Lackner and P. DiZio. Spatial orientation as a component of presence: Insights gained from nonterrestrial environments. *Presence: Teleoperators and Virtual Environments*, 7(2):108–115, April 1998.

[45] J. Loomis, R. Golledge, and R. Klatzky. Navigation system for the blind: Auditory display modes and guidance. *Presence: Teleoperators and Virtual Environments*, 7(2):193–203, April 1998.

[46] P. Lunneborg and C. Lunneborg. Everyday spatial activities test for studying differential spatial experience and vocational behavior. *Journal of Vocational Behavior*, 28:135–141, 1986.

[47] K. Lynch. *The Image of the City*. MIT Press, 1960.

[48] K. Lynch and L. Redwin. *A Theory of Urban Form*, pages 84–101. In Proshansky et al. [58], 1970.

[49] K. Lynch and M. Rivkin. *A Walk Around the Block*, pages 631–642. In Proshansky et al. [58], 1970.

[50] T. McNamara, J. Hardy, and S. Hirtle. Subjective hierarchies in spatial memory. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 15(2):211–227, 1989.

[51] G. Moore. *Theory and Research on the Development of Environmental Knowing*, chapter 12, pages 138–164. In Moore and Golledge [53], 1976.

[52] G. Moore and R. Golledge. *Environmental Knowing: Concepts and Theories*, chapter 1, pages 3–24. In [53], 1976.

[53] G. Moore and R. Golledge, editors. *Environmental Knowing: Theories, Research and Methods*. Dowden, Hutchinson and Ross, Inc., 1976.

[54] J.D. Mulder and R. van Liere, editors. *Virtual Environments 2000*. Lecture Notes in Computer Science. SpringerWienNewYork, 2000. Proceedings of the Eurographics Workshop in Amsterdam, The Netherlands, June 1-2, 2000.

[55] N. Newcombe. *Methods for the Study of Spatial Cognition*, chapter 11, pages 277–299. In Cohen [15], 1985.

[56] M. Orne. On the social psychology of the psychological experiment: with particular reference to demand characteristics and their implications. *American Psychologist*, 17:776–783, 1962.

[57] K. Osberg, W. Winn, H. Rose, A. Hollander, H. Hoffman, and P. Char. The effect of having grade seven students construct virtual environments on their comprehension of science. In *Proceedings of the Annual Meeting of the American Educational Research Association*, 1997.

[58] H.M. Proshansky, W. Ittelson, and L. Rivlin, editors. *Environmental Psychology: Man and his physical setting*. Holt, Rinehort and Winston, Inc., 1970.

[59] J. Prothero. The role of rest frames in vection, presence and motion sickness. Dissertation: University of Washington, 1998.

[60] J. Prothero, H. Hoffman, D. Parker, T. Furness III, and M. Wells. Foreground and background manipulations affect presence. In *Human Factors and Ergnomics Society*, pages 1410–1414, 1995.

[61] M. Rovine and G. Weisman. Sketch-map variables as predictors of way-finding performance. *Journal of Environmental Psychology*, 9:217–232, 1989.

[62] R. Ruddle, A. Howes, S. Payne, and D. Jones. The effects of hyperlinks on navigation in virtual environments. *International Journal of Human Computer Studies*, 53(4):551–581, October 2000.

[63] R. Ruddle, S. Payne, and D. Jones. Navigating buildings in desktop virtual environments: Experimental investigations using extended navigational experience. *Journal of Experimental Psychology: Applied*, 3(2):143–159, June 1997.

[64] R. Ruddle, S. Payne, and D. Jones. Navigating large-scale "Desktop" virtual buildings: Effects of orientation aids and familiarity. *Presence: Teleoperators and Virtual Environments*, 7(2):179–192, April 1998.

[65] R. Ruddle, S. Payne, and D. Jones. The effects of maps on navigation and search strategies in very-large-scale virtual environments. *Journal of Experimental Psychology: Applied*, 5(1):54 – 75, 1999.

[66] R. Ruddle, S. Payne, and D. Jones. Navigating large-scale virtual environments: What differences occur between helmet-mounted and desk-top displays? *Presence: Teleoperators and Virtual Environments*, 8(2):157 – 168, 1999.

[67] G. A. Satalich. *Navigation and Wayfinding in Virtual Reality - Finding Proper Tools and Cues to Enhance Navigation Awareness*. PhD thesis, University of Washington, 1986.

[68] H. Sayers, S. Wilson, W. Myles, and M. McNeill. Navigation in non-immersive virtual environments. In *Proceedings of the 7th UK VR-SIG Conference*, pages 43–51, 2000.

[69] A. Siegel and J. Cousins. *The Symbolising and Symbolised Child in the Enterprise of Cognitive Mapping*, chapter 14, pages 347–365. In Cohen [15], 1985.

[70] M. Slater, V. Linakis, M. Usoh, and R. Kooper. Immersion, presence, and performance in virtual environments: An experiment with tri-dimensional chess. In Mark Green, editor, *ACM Virtual Reality Software and Technology (VRST)*, pages 163–172, July 1996.

[71] M. Slater, M. Usoh, and Y. Chrysanthou. The influence of dynamic shadows on presence in immersive virtual environments. In M. Goebel, editor, *Virtual Environments '95*, Springer Computer Science, pages 8–21. Springer, 1995.

[72] StatSoft. Statistica help file, 2001.

[73] D. Stea. *Program Notes on a Spatial Fugue*, chapter 9, pages 106–120. In Moore and Golledge [53], 1976.

[74] E. Stubkjær. Roads to cities, to geographic information systems, and to knowledge. In *Proceedings of the International Conference of Computer Based Learning in Science*, pages 521–529, 1993.

[75] D. Tate, L. Sibert, F. Williams, T. King, and D. Hewitt. Virtual environment fire-fighting/ship familiarization feasability tests aboard the ex-uss shadwell. November 1995.

[76] E. C. Tolman. Cognitive maps in rats and men. *Psychological Review*, 55(4):189 – 208, July 1948.

[77] B. Tversky. Distortions in memory for maps. *Cognitive Psychology*, 13:407–433, 1981.

[78] N.G. Vinson. Design guidelines for landmarks to support navigation in virtual environments. In *CHI '99*, pages 278–285, May 1999.

[79] S. Volbracht and G. Domik. *Developing Effective Navigation Techniques in Virtual Environments*, chapter 6. In Mulder and van Liere [54], 2000. Proceedings of the Eurographics Workshop in Amsterdam, The Netherlands, June 1-2, 2000.

[80] D. Waller. Factors affecting the perception of inter-object distances in virtual environments. *Presence: Teleoperators and Virtual Environments*, 8(6):657–670, December 1999.

[81] D. Waller, E. Hunt, and D. Knapp. The transfer of spatial knowledge in virtual environment training. *Presence: Teleoperators and Virtual Environments*, 7(2):129–143, April 1998.

[82] D. Weatherford. *Representing and Manipulating Spatial Information from different Environments: Models to Neighbourhoods*, chapter 3, pages 41–70. In Cohen [15], 1985.

[83] R. West and C. Morris. *Spatial Cognition on Nonspatial tasks: Finding Spatial Knowledge when you;re not looking for it*, chapter 2, pages 13–39. In Cohen [15], 1985.

[84] C. D. Wickens. *Engineering Psychology and Human Performance*. Harper Collins Publishers Inc., 1992.

[85] W. Winn. A conceptual basis for educational applications of virtual reality. HITL Technical Report R-93-9. Human Interface Technology Laboratory, University of Washington, 1993.

[86] B. Witmer and P. Kline. Judging perceived and traversed distance in virutal environments. *Presence: Teleoperators and Virtual Environments*, 7(2):144–167, April 1998.

[87] B. Witmer and W. Sadowski. Nonvisually guided locomotion to a previously viewed target in real and virtual environments. *Human Factors*, 40(3):478–488, September 1998.

[88] B. Witmer, W. Sadowski, and N. Finkelstein. Training dismounted soldiers in virtual environments: Enhancing configuration learning. United States Army Research Institute for the Behavioral and Social Sciences Technical Report 1103, July 2000.

[89] B. Witmer and M. Singer. Measuring presence in virtual environments: A presence questionnaire. *Presence: Teleoperators and Virtual Environments*, 7(3):225–240, June 1998.

[90] D. Wood and R. Beck. *Talking with Environmental-A, an experimental mapping language*, chapter 32, pages 351–361. In Moore and Golledge [53], 1976.

Appendix A

Questionnaire Data Collection Screenshots

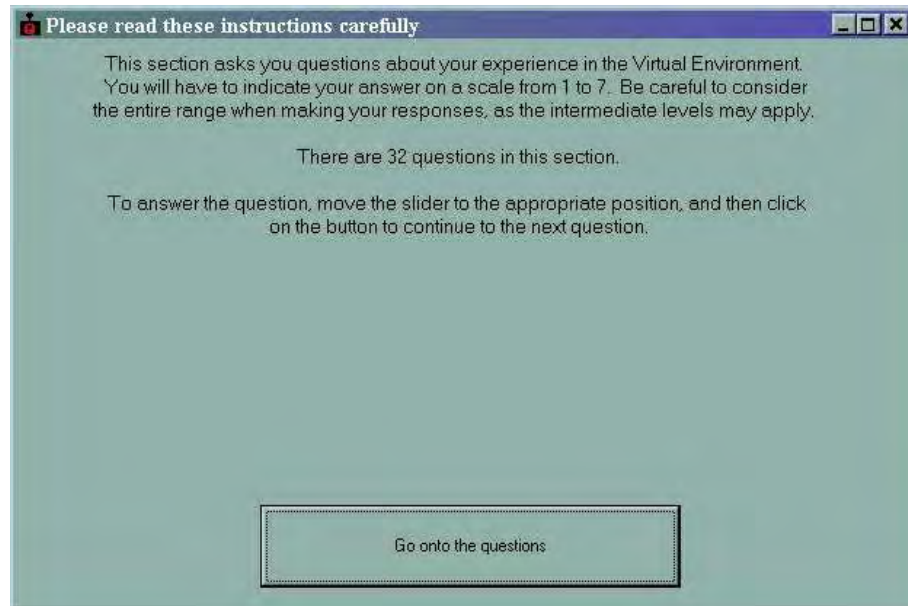*This appendix contains screenshots of the digitally presented questionnaires, described in Section 4.6.1*

Figure 44: A screenshot of the questionnaire administration program. This shows an example of the instructions given to participants before each section.
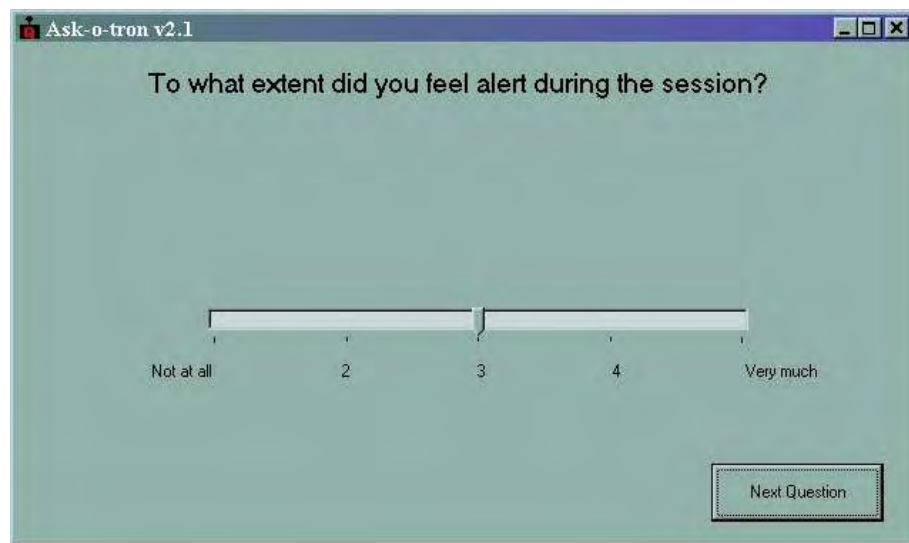
Figure 45: A screenshot of the questionnaire administration program. This shows a Likert-type question, where the user moves the slider bar to the appropriate position and then clicks on the button to continue to the next question. The user cannot continue to the next question without answering the current one.
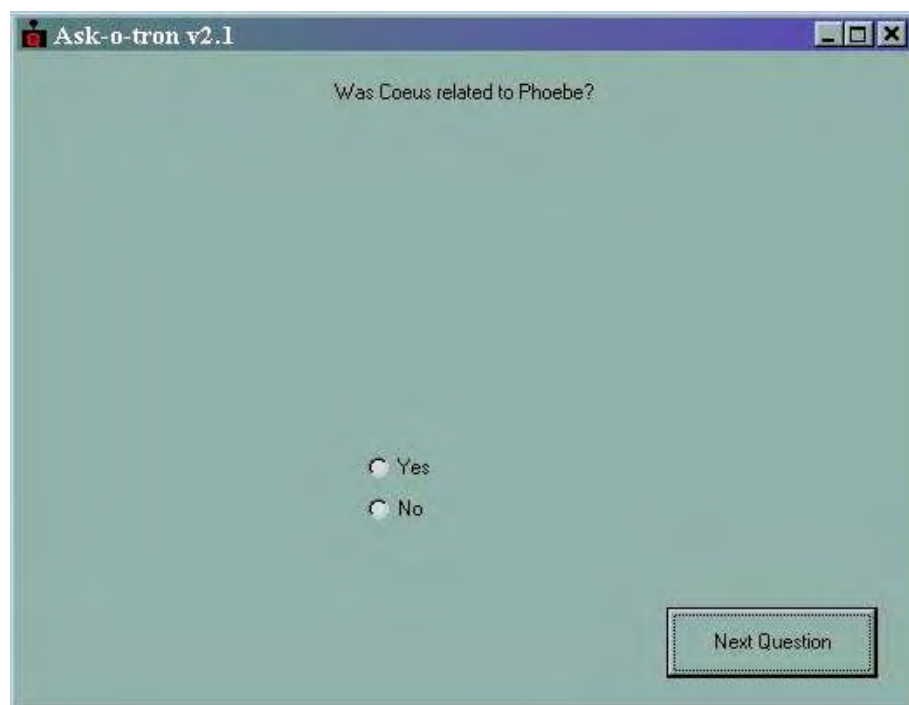

Figure 46: Another screenshot of the questionnaire administration program. This shows a multiple choice question, where the user selects a radio button and then clicks on the button to continue to the next question. The user must select a radio button before they are allowed to continue to the next question.
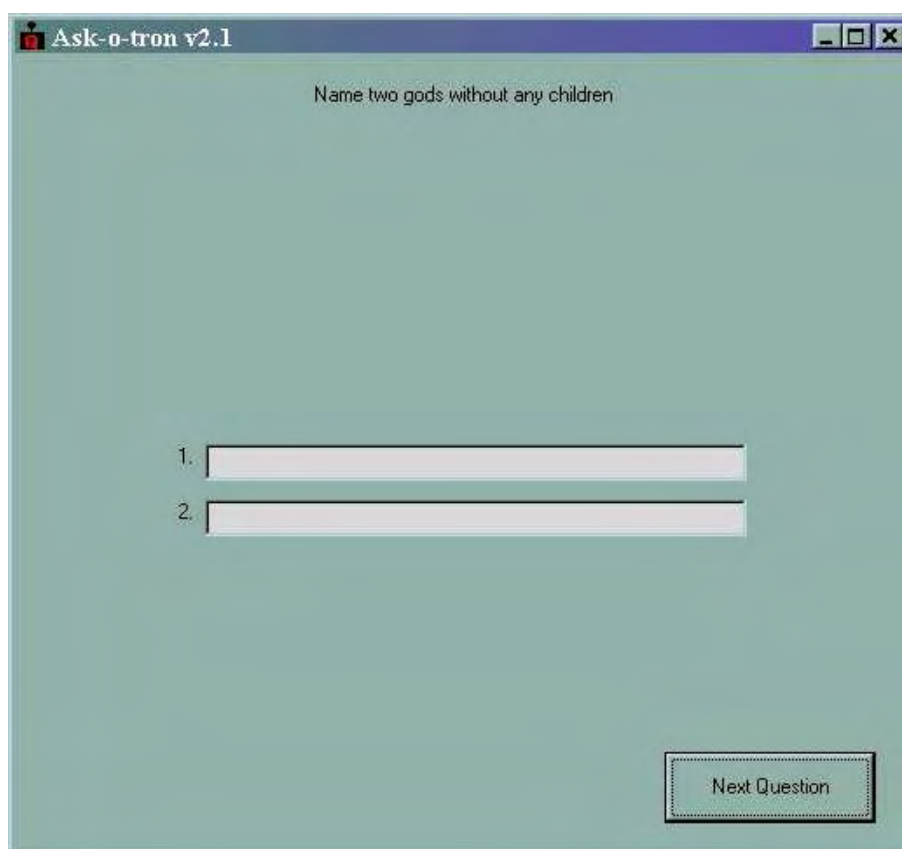
Figure 47: Another screenshot of the questionnaire administration program. This shows a multiple text entry question, where the user types their answers in the corresponding boxes. Text must be entered in at least the first box before the user can continue to the next question.

Appendix B

Differential Emotions Scale (DES)

This appendix contains the Differential Emotions Scale. The instructions at the top are those as given to the participants. VR participants were told to move the slider to the appropriate position, and click on the button to continue to the next question, while non-VR participants were told to place a tick in the box containing the appropriate number for each question.

B.1 About Your Feelings

This set of questions asks about your feelings during the session.

There are 18 questions in this section. Each question will ask to you rate how intensely you felt a particular emotion during the session on a scale from 1 (Not at all) to 5 (Very much).

1. To what extent did you feel happy during the session?

2. To what extent did you feel scared during the session?

3. To what extent did you feel amazed during the session?

4. To what extent did you feel fearful during the session?

5. To what extent did you feel astonished during the session?

6. To what extent did you feel sheepish during the session?

7. To what extent did you feel afraid during the session?

8. To what extent did you feel alert during the session?

9. To what extent did you feel surprised during the session?

10. To what extent did you feel downhearted during the session?

11. To what extent did you feel attentive during the session?

12. To what extent did you feel shy during the session?

13. To what extent did you feel delighted during the session?

14. To what extent did you feel joyful during the session?

15. To what extent did you feel bashful during the session?

16. To what extent did you feel sad during the session?

17. To what extent did you feel discouraged during the session?

18. To what extent did you feel that you were concentrating during the session?

Appendix C

Presence Questionnaire (PQ)

This appendix presents the Presence Questionnaire given to the VR participants. Anchor points are given below each question.

C.1   Your Virtual Reality Experience

This section asks you questions about your experience in the Virtual Environment. You will have to indicate your answer on a scale from 1 to 7. Be careful to consider the entire range when making your responses, as the intermediate levels may apply.

There are 32 questions in this section.

To answer the question, move the slider to the appropriate position, and then click on the button to continue to the next question.

1. How much were you able to control events?
   Not at all, Completely

2. How responsive was the environment to actions that you initiated (or performed)?
   Not responsive, Completely responsive

3. How natural did your interactions with the environment seem?
   Extremely artificial, Completely natural

4. How much did the visual aspects of the environment involve you?
   Not at all, Completely

5. How much did the auditory aspects of the environment involve you?
   Not at all, Completely

6. How natural was the mechanism which controlled movement through the environment?
   Extremely artificial, Completely natural

7. How compelling was your sense of objects moving through space?
   Not at all, Very compelling

8. How much did your experiences in the virtual environment seem consistent with your real world experiences?
Not consistent,Very consistent

9. Were you able to anticipate what would happen next in response to the actions that you performed?
Not at all, Completely

10. How completely were you able to actively survey or search the environment using vision?
Not at all, Completely

11. How well could you identify sounds?
Not at all, Completely

12. How well could you localize sounds?
Not at all, Completely

13. How well could you actively survey or search the virtual environment using touch?
Not at all, Completely

14. How compelling was your sense of moving around inside the virtual environment?
Not compelling, Very compelling

15. How closely were you able to examine objects?
Not at all, Very closely

16. How well could you examine objects from multiple viewpoints?
Not at all, Extensively

17. How well could you move or manipulate objects in the virtual environment?
Not at all, Extensively

18. How involved were you in the virtual environment experience?
Not involved, Completely engrossed

19. How much delay did you experience between your actions and expected outcomes?
No delays, Long delays

20. How quickly did you adjust to the virtual environment experience?
Never did, Less than a minute

21. How proficient in moving and interacting with the virtual environment did you feel at the end of the experience?
Not proficient, Very proficient

22. How much did the visual display quality interfere or distract you from performing assigned tasks or required activities?
Not at all, Prevented task performance

23. How much did the control devices interfere with the performance of assigned tasks or with other activities?
Not at all, Interfered greatly

24. How well could you concentrate on the assigned tasks or required activities rather than on the mechanisms used to perform those tasks or activities?
Not at all, Completely

25. How completely were your senses engaged in this experience?
Not engaged, Completely engaged

26. To what extent did events occurring outside the virtual environment distract from your experience in the virtual environment?
Not at all, Very much

27. Overall, how much did you focus on using the display and control devices instead of the virtual experience and experimental tasks?
Not at all, Very much

28. Were you involved in the experimental task to the extent that you lost track of time?
Not at all, Completely

29. How easy was it to identify objects through physical interaction; like touching an object, walking over a surface, or bumping into a wall or object?
Impossible, Very easy

30. Were there moments during the virtual environment experience when you felt completely focused on the task or environment?
None, Frequently

31. How easily did you adjust to the control devices used to interact with the virtual environment?
Difficult, Easily

32. Was the information provided through different senses in the virtual environment (e.g., vision, hearing, touch) consistent?
Not consistent, Very consistent

Appendix D

Everyday Spatial Abilities Test (ESAT)

This appendix presents the Everyday Spatial Abilities Test given to the VR participants. Left anchor points were labelled "very bad", right anchor points were labelled "very good", with three intermediate levels (i.e. 5 in total)

D.1   Spatial Abilities

This section asks you to rate how good you are at certain everyday activities.

There are 20 questions in this section.

Move the slider to the appropriate position, and click on the button to continue to the next question.

1. How good are you at building simple objects using wood or cardboard?

2. How good are you at passing an elementary physics course?

3. How good are you at arranging objects in a balanced, space-efficient manner?

4. How good are you at sketching rough plans and designs?

5. How good are you at making minor repairs on a house?

6. How good are you at passing an elementary laboratory course in Chemistry?

7. How good are you at packing and storing boxes and luggage quickly and efficiently?

8. How good are you at sketching geometric designs?

9. How good are you at following instructions to assemble furniture or toys?

10. How good are you at understanding graphs and charts in textbooks and magazines?

11. How good are you at judging accurately how many objects can fit on a shelf?

12. How good are you at drawing simple objects showing depth perspective?

13. How good are you at following plans for building models?

14. How good are you at solving practical problems using maths?

15. How good are you at efficiently arranging equipment or furniture?

16. How good are you at passing an introductory course in mechanical drawing?

17. How good are you at working successfully with hand tools?

18. How good are you at using rulers, micrometers and other measuring devices?

19. How good are you at arranging displays of objects tastefully?

20. How good are you at drafting and designing things?

Appendix E

Test on Data Set

This appendix contains the questions given to participants to test their knowledge of the data set. The instructions at the top are those as given to the participants. The correct answers are given in italics below each question.

E.1 Family Relationships

This section asks questions about the family relationships represented by the virtual environment.

There are 14 questions in this section.

1. Who had the most children?
   Correct answer was Gaea and Uranus.

2. Name two gods without any children
   Correct answer was any two of Hades, Demeter, Hestia, Poseidon, Artemis, Apollo, Aphrodite.

3. Was Coeus related to Phoebe?
   Correct answer was "yes".

4. Was Gaea related to Chaos?
   Correct answer was "yes".

5. Was Dione related to Phoebe?
   Correct answer was "yes".

6. Was Apollo from the same generation as Artemis?
   Correct answer was "yes".

7. Was Gaea from the same generation as Uranus?
   Correct answer was "no".

8. Were Phoebe and Rhea from the same generation?
   Correct answer was "yes".

9. Was Iapetus Clymene's child or partner?
   Correct answer was "partner".

10. Was Uranus Gaea's child or partner?
    Correct answer was both child and partner - half marks were awarded for either one.

11. Name 2 children (or grandchildren) of Aether.
    Any two of Gaea, Uranus, Coeus, Phoebe, Cronus, Rhea, Iapetus, Clymene, Leto, Hades, Demeter, Zeus, Hestia, Poseidon, Epimetheus, Artemis, Apollo, Aphrodite, Dione.

12. Was Leto of the same generation as the children of Cronus and Rhea?
    Correct answer was "yes".

13. Name two people who had only one parent.
    Correct answer was any two of Nyx, Uranus, and Dione.

14. Name one person who had children with their own child.
    Correct answer was any two of Chaos, Nyx, and Gaea.