CHAPTER IV

TRAVEL

4.1 Introduction and Definitions

Travel, or *Viewpoint Motion Control*, is one of the most basic and universal interactions found in virtual environment applications. We define travel as the control of the user's viewpoint motion in the three-dimensional environment. This is distinguished from *wayfinding*, which is the cognitive process of determining a path based on visual cues, knowledge of the environment, and aids such as maps or compasses. Together, travel and wayfinding make up the overall interaction called *navigation*. In our work, then, we are studying the techniques which allow a user to move from place to place in a VE, and not the displays or other aids which help the user to find her way.

Travel is almost certainly the most common interaction in VE applications, apart from simple head motion. In most VE systems, the user must be able to move effectively about the environment in order to obtain different views of the scene and to establish a sense of presence within the 3D space. Therefore, it is essential that travel techniques be well-designed and well-understood if VE applications are to succeed. In most cases, travel is not an end unto itself; rather, it is simply used to move the user into a position where he can perform some other, more important task. Because of this, the travel technique should be easy to use, cognitively simple, and unobtrusive. It is not obvious whether a given technique meets these criteria, so formal evaluation and analysis are important.

In this chapter, then, we will explore interaction techniques for viewpoint motion control in immersive VEs, beginning with prior work in the area. Next, an initial evaluation framework and four experiments will be presented. These experiments analyzed some common techniques and had important results, but fell short in other areas. This led to the development of an alternate framework. Another experiment is discussed which shows the relative advantages of the expanded approach. Finally, this chapter will discuss the testbed evaluation we performed for the task of viewpoint motion control.

4.2 Related Work

A number of researchers have addressed issues related to navigation and travel both in immersive virtual environments and in general 3D computer interaction tasks. It has been asserted (Herndon et al, 1994) that studying and understanding human navigation and motion control (e.g. Schieser, 1986, Warren & Wertheim, 1990) is of great importance for understanding how to build effective virtual environment travel interfaces. Although we do not directly address the cognitive issues surrounding virtual environment navigation, this area has been the subject of some prior investigation (e.g. Wickens, 1995). Wayfinding issues have been the subject of studies by Darken and Sibert (1996a, 1996b). Also, a system has been proposed (Ingram & Benford, 1995) which attempts to replicate the classic urban wayfinding cues identified in "The Image of the City" (Lynch, 1960).

Various metaphors for viewpoint motion and control in 3D environments have also been proposed. Ware et al. (1988, 1990, 1996) identify the "flying," "eyeball-in-hand," and "scene-in-hand" metaphors for virtual camera control. As an extension of the scene-in-hand metaphor, Pausch et al. (1995) make use of a "World-in-Miniature" representation as a device for navigation and locomotion in immersive virtual environments. Another interesting metaphor uses head motion to control the position of the viewpoint (Kheddar, Chellali, and Coiffet, 1995, Koller, Mine, and Hudson, 1996).

Numerous implementations and studies of non-immersive 3D travel techniques have been described. Strommen compares three different mouse-based interfaces for children to control point-of-view navigation (Strommen, 1994). Mackinlay et al. describe a general method for rapid, controlled movement through a 3D environment (Mackinlay, Card, and Robertson, 1990), and a similar technique is used immersively in the Cosmic Explorer application (Song and Norman, 1993). Ware and Slipp assessed the usability of different velocity control interfaces for viewpoint control in 3D graphical environments (Ware and Slipp, 1991).

Mine (1995) offers an overview of motion specification interaction techniques. He and others (Robinett & Holloway, 1992) also discuss issues concerning their implementation in immersive virtual environments. Several user studies concerning immersive travel techniques have been reported in the literature, such as those comparing different travel modes and metaphors for specific virtual environment applications (e.g. Chung, 1992, Mercurio et al., 1990). Physical motion techniques have also been studied (e.g. Iwata and Fujii, 1996), including an evaluation of the effect of a physical walking technique on the sense of presence (Slater, Usoh, and Steed, 1995).

4.3 Original Evaluation Framework

4.3.1 Categorization of Techniques

Given techniques for travel in immersive virtual environments, one could perform many experiments involving those techniques and come to some understanding of their effect on performance in certain applications. However, it is not entirely clear what determines the "performance" of a travel technique. Moreover, it would be difficult or impossible to determine which components of the techniques were significant in improving or lessening performance, and results from one application or task would not necessarily transfer to another. For this reason, we have devised a more formalized framework within which to evaluate virtual travel techniques. Stanney (1995) proposes that a taxonomy of interaction techniques is needed for "imposing order on the complex interactions between user, task, and system phenomena." The evaluation framework presented here includes such a taxonomy and an emphasis on outside factors which can influence user performance.

In order to understand travel techniques and their effects more deeply, we need to categorize them and break them down into their lower-level components. Toward this end, we have developed a taxonomy of immersive travel techniques, which is presented in

Figure 4.1. The taxonomy splits a technique into three components, which apply regardless of the type of travel being done (exploration, search, maneuvering, etc.).

Direction/Target Selection refers to the method by which the direction or object of travel is specified. Depending on whether control of direction is continuous or not, the user may either "steer" (choose a direction), or simply choose a target object. Gaze-directed steering, in which the user moves in the direction she is looking, and pointing, where the user points in the direction she wants to go, are two popular steering techniques. This section also lists techniques for discrete selection of a target object.

Velocity/Acceleration Selection techniques allow the user to vary the speed of travel. Many VE applications dispense with this entirely, and use a constant travel velocity. However, several techniques have been proposed, including continuous gestures to select velocity, the use of props such as foot pedals, or adaptive system-controlled speed.

The final component of a travel technique is the *Conditions of Input*. This refers to the input required by the system in order to begin, continue, and end travel. The user may be in constant motion, in which case no input may be required. Alternately, the system may require continuous input to determine the user's state, or simple inputs at the beginning and/or end of a movement. Again, this component may be under system control.

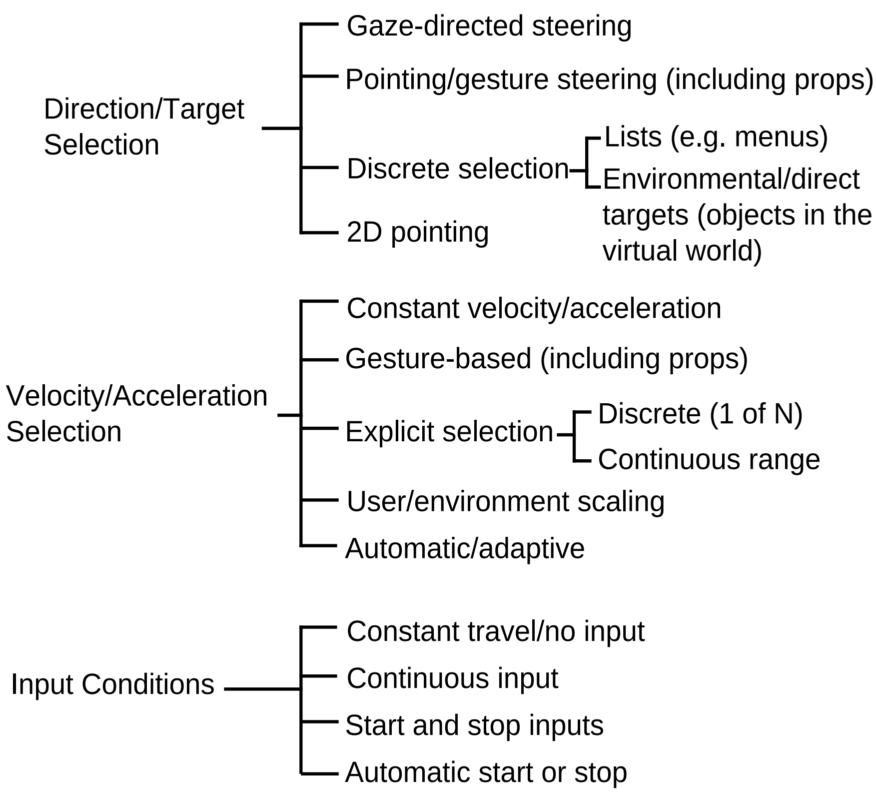

Figure 4.1 Taxonomy of Travel Techniques for Immersive Virtual Environments

We do not claim that this taxonomy is complete, since many new techniques for controlling user motion are being designed. However, most current techniques fit into the taxonomy, at least at a high level. More importantly, by breaking a technique into three components, we can study them separately, and gain a greater understanding of differences in performance. A technique which is performing poorly may be improved by changing only one of the components, but this might not be recognized unless techniques are divided into their constituent elements.

This taxonomy also encourages the guided design of new techniques. By choosing a component (and an implementation of that component) from each section of the taxonomy, a travel technique may be created from its parts, and useful new combinations may come to light. Not all components will work with all others, but there are many opportunities for interesting designs.

For example, one might combine environmental target selection with gesture-based velocity selection, explicit start inputs, and explicit or automatic stop inputs. This would produce a technique that would allow a user to travel along a path from the current position to a specified object, using a high velocity on the less interesting parts and a slower speed at places of interest. The user could stop moving at any point along the path, or be stopped automatically when the target object was reached. Such a technique might be a natural fit for an immersive “tour” application, where there are certain known places that users wish to visit, but designers also desire that movement be under some degree of user control.

We limit the scope of our design and evaluation to travel techniques implementing virtual movement. That is, we will not consider techniques which use physical motions such as walking in place or walking on a treadmill. Such techniques may be quite natural and useful, but are not generally applicable to VE applications, especially when three-dimensional motion is needed.

4.3.2 Performance Measures

There are few categories of virtual environment applications that are currently in use for productive, consistent work, but the requirements of these applications for travel techniques cover a wide range. Further, there are many new applications of VEs being researched, which also may require travel techniques with different characteristics. It is therefore impractical to evaluate travel techniques directly within each new application. Instead, we propose a more general methodology, involving a mapping from travel techniques to a set of performance metrics. These are measurable characteristics of the performance of a technique. With this indirect mapping, application designers can specify desired levels of various metrics, and then choose a technique which best fits those requirements.

Our list of performance metrics for immersive travel techniques includes:

1. *Speed* (efficient task completion)
2. *Accuracy* (proximity to the desired target)
3. *Spatial Awareness* (the user’s knowledge of his position and orientation within the environment during and after travel)
4. *Ease of Learning* (the ability of a novice user to use the technique)
5. *Ease of Use* (the complexity or cognitive load of the technique from the user’s point of view)
6. *Information Gathering* (the user’s ability to actively obtain information from the environment during travel)

7. *Presence* (the user's sense of immersion or "being within" the environment due to travel)
8. *User Comfort* (lack of physical discomfort, including simulator sickness (e.g. Hettinger and Riccio, 1992))

Again, this list may not be complete, but it is a good starting point for quantifying the effectiveness and performance of virtual travel techniques. In particular, we emphasize that speed and accuracy are not the only characteristics of a good travel technique, and in many applications are not the most important. For example, the designer of an architectural walkthrough application might be most interested in high levels of spatial awareness, information gathering, and presence. By doing experiments that relate travel technique components to performance metrics, we can identify techniques that meet those needs, and the results of the experiments will also be generalizable and reusable by designers of other applications.

Some of the metrics, such as speed and accuracy, are simple to measure quantitatively. Others, however, are difficult to measure due to their inherent subjective nature. To quantify these metrics, standard questionnaires for factors such as ease of use (e.g. Chin, Diehl, & Norman, 1988), presence (e.g. Slater, 1995), and simulator sickness (e.g. Kennedy et al., 1993) should be part of the experimental method.

4.4 Initial Experiments

Using this framework, we designed and ran three initial experiments on common VE travel techniques (These experiments are described in more detail in Bowman, Koller, & Hodges, 1997). We wanted to show that generalizable results could be obtained without knowing the target application. These experiments produced useful data which is applicable in a variety of situations.

4.4.1 Spatial Awareness Experiment

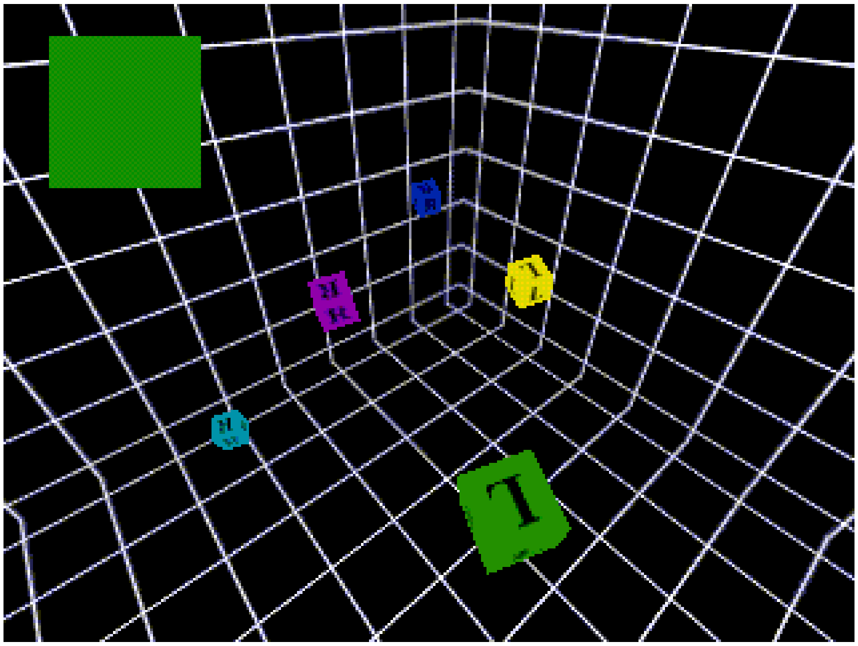

Figure 4.2 Environment for the Spatial Awareness Experiment

Our first experiment focused on one of the more abstract performance metrics: spatial awareness. We were interested in how immersive travel techniques would affect the user's awareness of the three-dimensional environment around him. Specifically, we tested how various velocity and acceleration schemes altered the user's level of spatial awareness.

The virtual environment for this experiment consisted of a set of cubes of contrasting colors, as seen in Figure 4.2. Users learned the locations of the cubes within the space, from both stationary and moving positions. In an experimental trial, the user was taken from the starting location to a new location, then shown a colored stimulus, matching the color of one of the cubes. We measured the user's spatial awareness by the time required to find the cube of that color. The subject proved she had found the correct cube by pressing either the left or right mouse button depending on the letter ("L" or "R") printed on the cube.

We contrasted four different velocity/acceleration techniques, each of which was system-controlled. The first two techniques used a constant velocity, one quite slow, the other relatively fast. We also implemented and tested a "slow-in, slow-out" technique, in which travel starts and ends slowly, with acceleration and deceleration in between. Finally, we tested an infinite velocity (also called "jumping" or "teleportation") technique, where users are taken immediately to the target location.

The results of the experiment showed that the level of spatial awareness was significantly decreased with the use of a jumping technique ($p < 0.01$). In fact, users were generally reduced to a simple search of the space after jumping from one location to another. This is a significant result, since many application designers might be tempted to use teleportation because of its speed and accuracy. The experiment shows that this is unwise unless some degree of user disorientation is acceptable in the target application. Surprisingly, none of the other techniques showed significant differences in performance: even up to relatively large velocities, users could maintain spatial awareness.

4.4.2 Absolute Motion Experiment

In the second experiment, we wanted to obtain some basic information about the speed and accuracy of two common steering techniques: gaze-directed steering, in which the direction of motion is determined by the user's gaze, and pointing, in which the user's hand orientation determines the direction of travel. Even though speed and accuracy are not always the most important considerations in a travel technique, they are still widely desirable. Once a target has been chosen, it is usually unacceptable to the user to move there slowly or imprecisely. We chose to compare gaze-directed steering with pointing because they seem to be quite different in their focus: gaze-directed steering is simple but constraining, while pointing is expressive but also more complex.

The experimental task was quite simple. Users traveled using one of the techniques from a starting location to a target sphere. We varied the size of the sphere and the distance to the sphere. We hypothesized that gaze-directed steering might produce greater speed and accuracy than pointing, because of its simplicity and the relative stability of the head compared to the hand.

Although gaze-directed steering did produce slightly better times for this task, we found that there was no statistically significant difference between the two techniques. Users were able to travel very close to the optimal straight-line path between the starting and target locations whether gaze-directed steering or pointing was used. This was useful information given the advantages of pointing shown by our next experiment.

4.4.3 Relative Motion Experiment

Rather than moving directly to an object in the environment, in this experiment users were required to move to a point relative to an object in the 3D space. This task is commonly used in applications such as architectural walkthrough. For example, suppose the user wishes to obtain a head-on view of a building so that it fills his field of view. There is no specific target object; rather, the user is moving relative to the building. In this experiment, the target was located on a line defined by a three-dimensional pointer, at a known distance from the tip of the pointer. Figure 4.3 shows the pointer and the target, although the target was not visible during experimental trials.

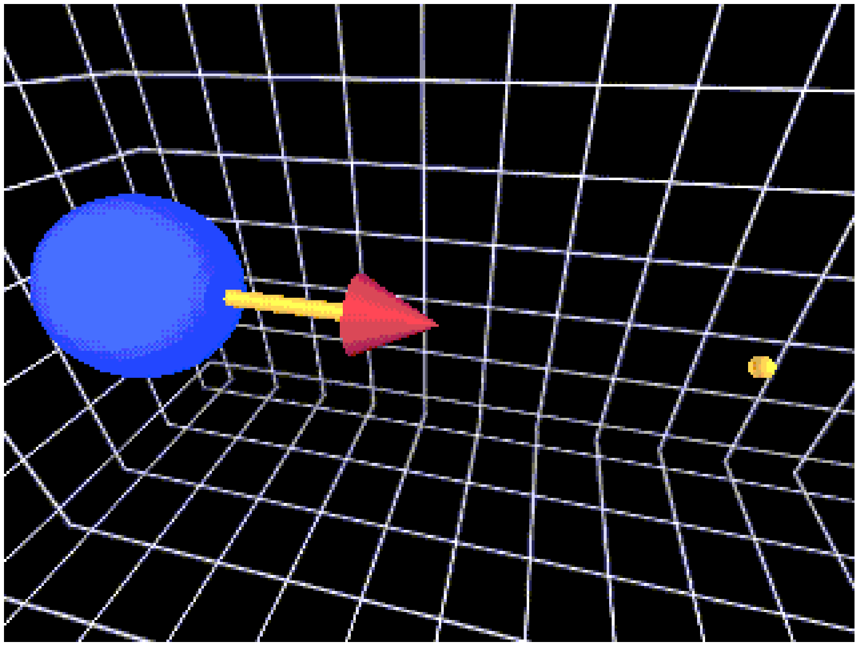

Figure 4.3 The Relative Motion Experiment Environment

Again, we measured speed and accuracy for the gaze-directed steering and pointing techniques. With this task, however, we highlighted the main difference between the two techniques: that gaze-directed steering requires the user to be looking in the direction of motion, while pointing allows gaze and travel to be in different directions. Thus, users of the pointing technique could look at the pointer to judge their travel to the target location, while gaze-directed steering required users to look at the pointer, then look in the estimated target direction to travel, then look back to check their progress, and so on.

Indeed, the experiment showed that the pointing technique was significantly faster for the relative motion task ($p < 0.025$). When combined with results from the absolute motion experiment, we can conclude that pointing is a good general-purpose technique where speed and accuracy are important performance measures.

4.5 Expanded Evaluation Framework

Although our initial set of experiments produced significant results in evaluations of some common VE travel techniques, we also noted that we were not able to capture a complete picture of the techniques from simple experimental designs. The problem was that our experiments studied the effects of a single factor only (travel technique), and did not consider other factors that might have an important effect on performance.

This is illustrated well by the absolute and relative motion experiments. Though they tested the same techniques (gaze-directed steering and pointing) and measured the same performance variables (speed and accuracy), they produced quite different results. In the case of absolute motion, the two techniques performed equally, but for a relative motion task, pointing showed more speed and accuracy. There was, therefore, an interaction between technique and task. This illustrates the fact that a technique cannot in general be considered in isolation from the task for which it is to be used.

Similarly, characteristics of the environment may affect the performance of a travel technique. Consider the absolute motion experiment. In the environment that we used, there was only a single object (the target), visible at all times, with a straight-line path between it and the user. In this environment, gaze-directed steering and pointing produced the same results. However, if the environment had been full of distracter objects and obstacles that the user had to avoid to reach the target, the two techniques might have exhibited significantly different performance characteristics. Techniques cannot be considered in isolation from the environments in which they are to be used.

For these reasons, we felt it necessary to expand our evaluation framework to include the multitude of other factors that can affect performance of virtual travel techniques. Rather than attempting to discern these dependencies in an ad hoc fashion for each experiment that is run, our expanded framework formalizes the notion that many variables contribute to the performance metrics. By explicitly including these variables in the framework, we can more easily choose what factors to control in an experimental setting, and choose values wisely for those variables which will be held constant. The expanded framework includes variables related to task, environment, user, and system characteristics.

4.5.1 Task Characteristics

For immersive travel, there are many factors related to the task that could conceivably affect performance. Some of these characteristics come directly from a consideration of the performance values that we wish to measure. Some of the task characteristics that we consider are:

- Distance to be traveled
- Amount of curvature or number of turns in the path
- Visibility of target
- Number of degrees of freedom of motion required
- Accuracy required
- Complexity of the task; cognitive load induced on the user
- Information required of the user

For example, we could distinguish between the absolute and relative motion tasks described above using the visibility characteristic. The target is invisible in the relative motion task, meaning that other objects in the environment must be used to determine the location of the target.

4.5.2 Environment Characteristics

As we have noted, the environment in which the user travels can also have an effect on performance. The same task in different environments may produce strikingly different results on one or more of the performance measurements. We have identified characteristics such as:

- Visibility within the environment
- Number of obstacles or distracters
- Activity or motion within the environment
- Size of the environment
- Level of visual detail and fidelity
- Homogeneity (amount of variation) in the environment
- Structure
- Alignment with the standard axes

Varying one or more of these environment variables may have allowed us to see some significant differences between the gaze-directed steering and pointing techniques in the absolute motion experiment. For example, adding more distracter objects or greater activity in the environment may have caused the more cognitively simple gaze-directed steering technique to perform better.

4.5.3 User Characteristics

It is also important to consider the differences in users of VE applications when evaluating performance. This can be a significant factor in the performance of various techniques, because the designers of techniques often assume something implicitly about users. Work in the field of user modeling (Kobsa & Wahlster, 1989) is quite relevant to this part of our framework. We are considering, among others, the following user characteristics:

- Age
- Gender
- Visual acuity
- Height
- Reach
- Ability to fuse stereo images
- Experience with VEs
- Experience with computers
- Technical / non-technical background
- Spatial ability

The importance of taking user characteristics into account became quite evident during a study we performed comparing various techniques for selecting and manipulating virtual objects (Bowman & Hodges, 1997). Our implementation of one technique (Poupyrev et al., 1996) mapped the user's physical arm extension to a more lengthy virtual arm extension, so that the number of objects that could be selected depended on the user's reach.. In the user study, most people liked this technique, but a few of our users had very short arms, and could not reach many of the objects at all. This caused them to become quite frustrated with this technique and to prefer other techniques that did not rely on physical arm length.

4.5.4 System Characteristics

Finally, we have extended our framework to include aspects of the hardware or software used to realize the virtual environment application. It is quite possible that design decisions made by system developers or hardware designers may affect the performance of techniques for virtual travel. However, just because these factors are not always under the control of the technique designer does not mean that they should not be considered in the design. For best performance, designers may need to create techniques that perform in a robust manner under a wide variety of system conditions. The system characteristics we have identified include:

- Rendering technique
- Lighting model
- Frame rate
- Latency
- Display characteristics (stereo/mono, field of view, resolution, brightness, etc.)
- Collision detection
- Virtual body representation

These factors can cause differences in the usefulness of many interaction techniques. Studies on the effects of varying frame rate and latency for various tasks have been performed (e.g. Ware & Balakrishnan, 1994), but there is still much work to be done.

4.5.5 Information Gathering Experiment

In order to validate our evaluation methodology, we designed and ran a new experiment within our expanded framework. We hoped to isolate some important and general results, and to show the usefulness of considering a larger number of experimental variables simultaneously.

Our focus was the effect of various steering techniques on the performance metric of information gathering. Information gathering is an important goal in many situations, and it is especially applicable to immersive virtual environments. Many of the major categories of VE applications, such as architectural walkthrough (e.g. Brooks, 1992), information visualization (e.g. Ingram & Benford, 1995, Bolter et al., 1995), simulation and training (e.g. Tate, Sibert, & King, 1997), and education (e.g. Allison et al., 1997), have a strong informational component. If the user is not able, for whatever reason, to focus on and remember important information, then the utility of the VE application is questionable.

There are many possible reasons why a user might not be able to gather as much information as is desirable, but a major factor is cognitive load. A famous result from cognitive psychology (Miller, 1956) shows the severe limitations on the capacity of working memory. When other influences force the person to use part of his working memory or other cognitive resources, information may be lost, or displaced (Baddeley, 1983). We wondered whether travel techniques induced cognitive load, and could therefore affect the amount of information that could be recalled by the user.

We chose to focus on the direction selection portion of the taxonomy, and to again study gaze-directed steering and pointing techniques. We also added a third technique, torso-directed steering, in which a tracker is attached to the user’s torso, so that she travels in the direction her body is facing. We felt that these three represented a useful crosssection of commonly used techniques, and that there were some interesting tradeoffs among them.

For example, both pointing and torso-directed steering have the advantage that the user can look in one direction and move in another. This could be important when gathering widely scattered information. However, these techniques are also cognitively more difficult than gaze-directed steering, in which head orientation is the only thing the user must control. Torso-directed steering might be more natural (since it simulates the way we walk) and thus produce less cognitive load than pointing, but it also has the disadvantage that it can only be used to move in a horizontal plane, as the torso cannot comfortably be pointed up or down. We were quite interested to see how these tradeoffs affected a user's ability to gather information.

Looking at our expanded framework, however, we felt that there were several other factors that could influence performance on this task. Therefore, we also chose one environment characteristic and one system characteristic to vary along with the travel technique. First, we felt that the complexity of the path through the environment might be quite important in the cognitive load induced upon a user. We captured this complexity characteristic in the *dimensionality* of the path. That is, some paths would be one-dimensional: straight and horizontal; others would be two-dimensional: still horizontal, but with turns; and still others would be three-dimensional: having turns and also vertical components.

Second, we hypothesized that the presence or absence of a collision detection feature might affect information gathering. If a user is focusing on information and not on the path he is traveling, he may move through a wall or other object. The effort required to move back through the object and back onto the desired path may use cognitive resources and displace information. With collision detection available, the user is kept near to the path, and is free to gather information without paying as much attention to the direction of motion. On the other hand, the use of collision detection may violate the mental model of the user, if the user has been told that he will keep moving as long as a button is pressed, for example. This also may induce cognitive loading. Therefore, we were interested to see how the use of collision detection would affect performance.

4.5.5.1 Method

To measure the user's ability to gather information, we decided to use a memory task. Subjects traveled through corridors, using one of the three steering techniques. Corridors were used so that the user would have only a single, directed path through the environment, with no choices as to which path to take. The experiment used one-, two-, and three-dimensional corridors, 3x3 meters in size, made up of straight segments, and employing only 90 degree turns. An outside view of a three-dimensional corridor is seen in figure 4.4. Signs, each containing a single word, were located on the walls, ceilings, and floors of the corridors, as seen in figure 4.5. The words used were common, short, non-proper nouns and were randomly scattered through the corridor. Each corridor contained 12 signs. Subjects were instructed to minimize the amount of time spent in the corridor (the maximum time was 60 seconds, but a trial also ended if the subject reached the end of the corridor), and also maximize the number of words and locations of words that they could remember.

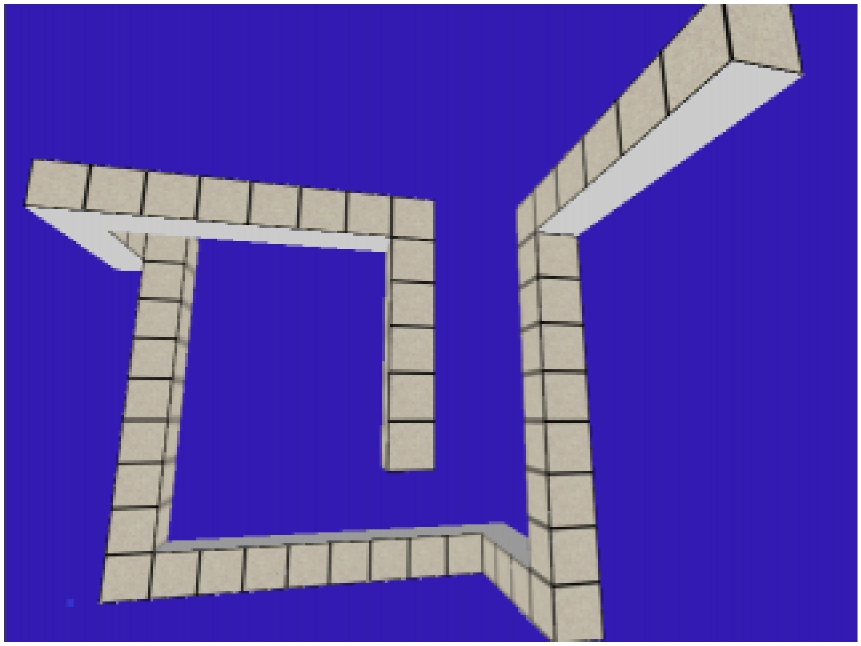

Figure 4.4 Outside View of a Three-Dimensional Corridor

Thus, we presented subjects with a very difficult, memory-overloading task. It has been shown that the limit of working memory is generally seven plus or minus two chunks of information (Miller, 1956), and we were presenting 12 words and associated sign locations to the subject. Even if subjects could store both the word and location as a single chunk, and even if some words could be chunked together semantically or in some other way, the amount of information should still fill working memory. Therefore, if cognitive load is induced because of the travel technique, the dimensionality of the corridor, or the presence or absence of collision detection, we should observe that the amount of remembered information should decrease.

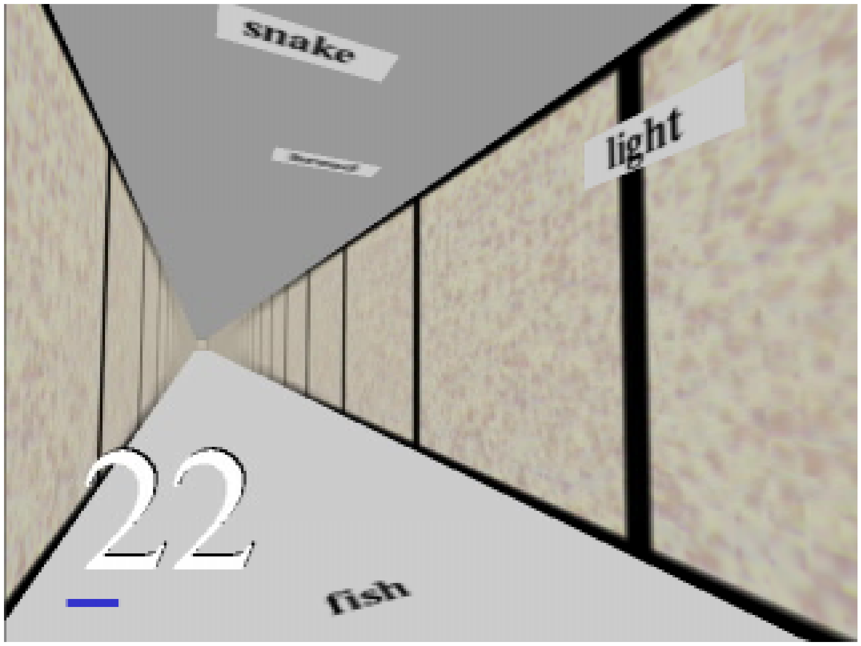

Figure 4.5 Interior of a Corridor from the Information Gathering Experiment

In order to demonstrate their memory of the corridor, subjects indicated words and locations on a paper map of the corridor immediately after each trial. An example map is shown in figure 4.6. Subjects indicated the position of the sign along the corridor, the surface on which the sign was seen, and the word printed on the sign. If words were remembered without locations, or vice-versa, these could also be listed on the map.

For each of the steering techniques, the other two components of a complete travel technique were held constant. Velocity was 3.0 meters per second while traveling; subjects began travel by pressing and holding a button, and stopped by releasing the button.

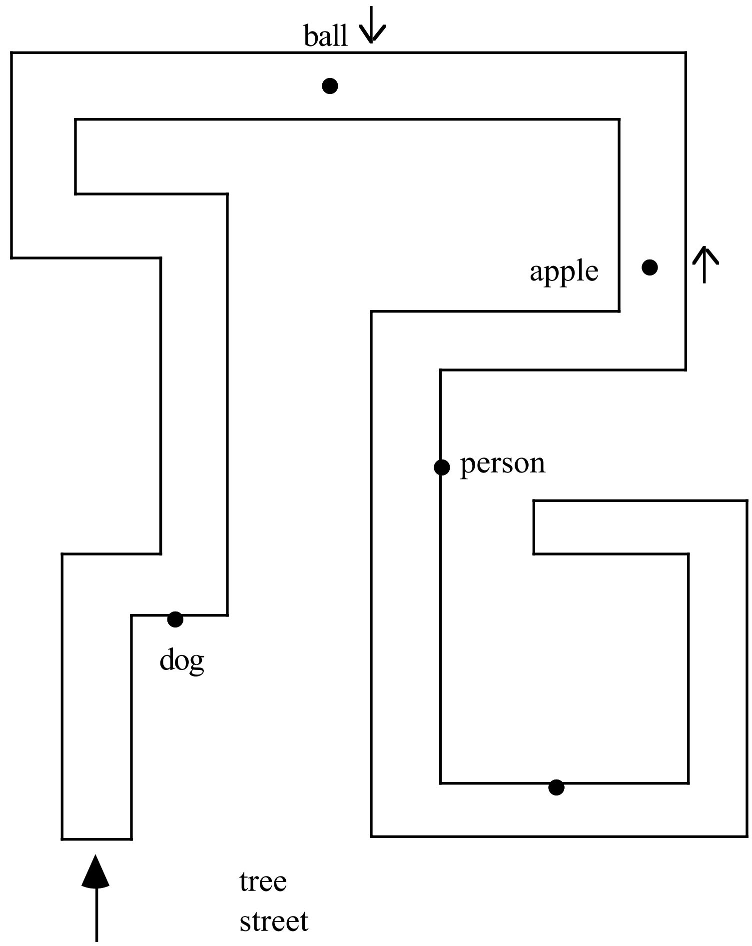

Figure 4.6 Example Completed Corridor Map with Four Word/Location Pairs, One Unpaired Location, and Two Unpaired Words

Each subject completed 16 trials: six each with the gaze-directed steering and pointing techniques, and four with the torso-directed steering technique. Within each technique, there were two trials of each dimensionality (the torso-directed technique can only be used in one- and two-dimensional environments), where one of the trials used collision detection and the other trial did not. Thus, each combination of the three variables (steering technique, dimensionality, and collision detection) was encountered once by each participant. Each subject traveled through each corridor exactly once, and the order of the corridors was different for each subject. To be less confusing for the subject, trials using a given technique were grouped together; however, we counterbalanced the order in which the techniques were seen. To eliminate effects of learning the techniques, subjects spent time in a “practice room” before each set of trials, where they practiced the use of the next steering technique.

Twenty-six student volunteers (twenty-three males and three females) participated in the study. Two subjects quit the experiment before completion due to dizziness or nausea induced by the VE system. Each subject completed a pre-session questionnaire (Appendix A) in which we gathered demographic data such as age, gender, eyesight, technical background, computer knowledge, and experience with immersive VEs. Subjects wore a Virtual Research VR4 head-mounted display (HMD), and were tracked using Polhemus Fastrak or Isotrak II electromagnetic trackers. Input was given to the system with a three-button joystick. The system maintained a constant rate of 30 frames per second.

4.5.5.2 Results

The experiment measured various response variables related to the information gathering task. We measured the time spent in each corridor, the number of word/location pairs the subject got exactly right, and several variations of partially correct words and locations. Since we had instructed subjects to maximize several things simultaneously, we desired a single response variable that would encompass all of these values. The formula used for this overall score was: $1/3\ (60-t) + 3a + 2\ (b+c+d) + e + f + g$, where t=seconds spent in the corridor, a=number of word/location pairs exactly correct, b, c, and d represent responses that have two of three aspects (word, position, and surface) correct, and e, f, and g are responses where only one of the aspects are correct. This formula gives higher weight to the most correct responses, and rewards moving quickly through the corridor.

Using this metric as our response variable, we performed a three-factor analysis of variance (ANOVA). Results were quite clear: the dimensionality of the corridor was extremely significant in affecting the score ($p < 0.01$), but travel technique and collision detection did not have a significant effect. Further analysis using Duncan's test for comparison of means showed that the average score for each dimensionality was significantly different than the averages for the other two dimensionalities. Table 4.1 presents the average scores for each condition.

Table 4.1 Average Values of Overall Score for Each Tested Treatment Combination in the Information Gathering Experiment; Higher Scores are Better

	1-Dimensional		2-Dimensional		3-Dimensional	
	Collision Off	Collision On	Collision Off	Collision On	Collision Off	Collision On
Gaze-directed	16.90	16.51	11.85	11.21	10.21	9.57
Pointing	15.57	16.68	10.36	10.85	9.33	9.38
Torso-directed	15.50	15.92	10.63	12.15	N/A	N/A

We also performed further analysis of the data in order to find other relationships between our three independent variables and performance of the information gathering task. First, we wondered whether any learning was occurring during the trials themselves. We plotted learning curves for each of the orderings of techniques (necessary since the number of trials depended on the technique), and found no significant improvement in score over time for any of the orderings, implying that neither the use of the technique nor the task strategy changed much as the trials progressed.

Second, we also performed a three-factor ANOVA with total time per trial as the response variable, in order to see which variables had an effect on the speed with which users moved through the corridors. The results here were synonymous with the previous

ANOVA: dimensionality was the only significant factor ($p < 0.01$). Thus, as the dimensionality of the path increased, time spent in the corridor increased. Most subjects finished the one-dimensional corridors quickly, while two- and three-dimensional corridors often took the entire 60 seconds.

Finally, we examined the demographic data collected in the questionnaire for any trends relating this information to performance of the information gathering task. There was a fairly even split between those who had never experienced immersive VEs (16 subjects) and those who had used a VE system previously (10 subjects). Among those who completed the experiment, the more experienced participants had a slightly higher average score per trial (13.2 vs. 11.5). This is not a statistically significant result, but may show that users with even a single experience using a VE application were more focused on the task and not distracted by the technology itself or the feel of the system.

4.5.5.3 Discussion

The results of the information gathering experiment were somewhat surprising, as we had expected that different steering techniques would produce different levels of cognitive load, and thus significantly affect overall scores. We found, though, that path dimensionality was the only significant variable, and that it dominated the results. Our intution regarding the techniques was not sufficient to predict the results (hypothesis 1) However, this does not mean that we learned nothing about the nature of the travel techniques in question.

On the contrary, we noted many important characteristics of the various techniques that help us to explain the lack of significant differences from the experiment. First of all, as we noted previously in our absolute motion experiment (Bowman, Koller, & Hodges, 1997), novice users tend to emulate gaze-directed steering with pointing (by keeping their hands pointed in the direction of their gaze) unless there are large rewards for doing otherwise (as in the relative motion experiment). We saw this again in the current experiment, and also noted the same characteristic with the torso-directed steering technique. This fact quite possibly led to the lack of significant differences between the techniques. We hypothesize that users more familiar with the techniques would be able to use them more advantageously (e.g. look to the side as you move forward with the pointing or torso-directed steering techniques). Given enough expert users of the techniques, it would be interesting to include the experience level of users as another independent variable.

Also, as we stated at the beginning of this section, each technique contains certain tradeoffs. Intuitively, gaze-directed steering should produce the least cognitive load of the three techniques. However, it also provides fewer affordances for information gathering (one must stop moving in order to look to the side for information). The opposite is true of pointing: it should be more cognitively complex but should better afford information collection. Since we have only one measure of information gathering ability, these tradeoffs may have balanced out, producing no visible differences between techniques. In order to further examine these tradeoffs, we would need experimental tasks that test the limits of both sides.

This experiment also showed the usefulness of our evaluation framework. Before the experiment began, it was not clear what factors would lead to significant performance differences. However, because of our expanded framework, we were able to identify three different factors which we felt could be important in an information gathering task. Had we considered only travel techniques in isolation, this experiment might not have revealed any significant results. Because we varied several factors, however, we were able to identify a characteristic with an extremely significant effect on performance.

We found no statistically significant information about the effects of the use of the collision detection feature. However, several subjects did comment to the experimenter that they felt that it was easier to move through the space and perform the task when this feature was enabled. This in itself should encourage designers to include this characteristic in their systems.

Finally, we observed that our subjects had several different strategies for performing the experimental task. Some focused on time, and raced through the environment as quickly as possible, memorizing a few words and locations along the way. Others were much more deliberate, stopping at each sign or cluster of signs to try to commit them to memory. Still others developed hybrid schemes. Subjects also differed in what they attempted to remember. Some consistently recalled the first three or four words and locations (the primacy effect), while others focused on the last things they saw in the corridor (the recency effect). A third group simply wrote down as many words as they could, then tried to match them to locations on the map.

All of these dissimilar strategies may have affected our ability to get significant results. We could have imposed a strategy on the user by instructing them explicitly to perform the task a particular way, and perhaps seen less variability. On the other hand, users will also have differing methods in real applications, and we should be searching for interaction techniques which perform in a robust manner under a variety of strategies. In this sense, it is correct to allow the user flexibility in determining the most appropriate tactic for the task at hand. Formal evaluation of the effects of user strategy proved important in a later experiment.

4.6 Alternate Evaluation Framework

4.6.1 Taxonomy

The expanded framework for design and evaluation of travel techniques provides a great deal of power in explaining performance differences. However, we do not feel that our initial taxonomy of techniques is as complete and general as it should be. One of the reasons for this is the “force-fit” of techniques for quite different travel tasks into the same category. Therefore, we have developed an alternate taxonomy (figure 4.7) based on a very simple task analysis.

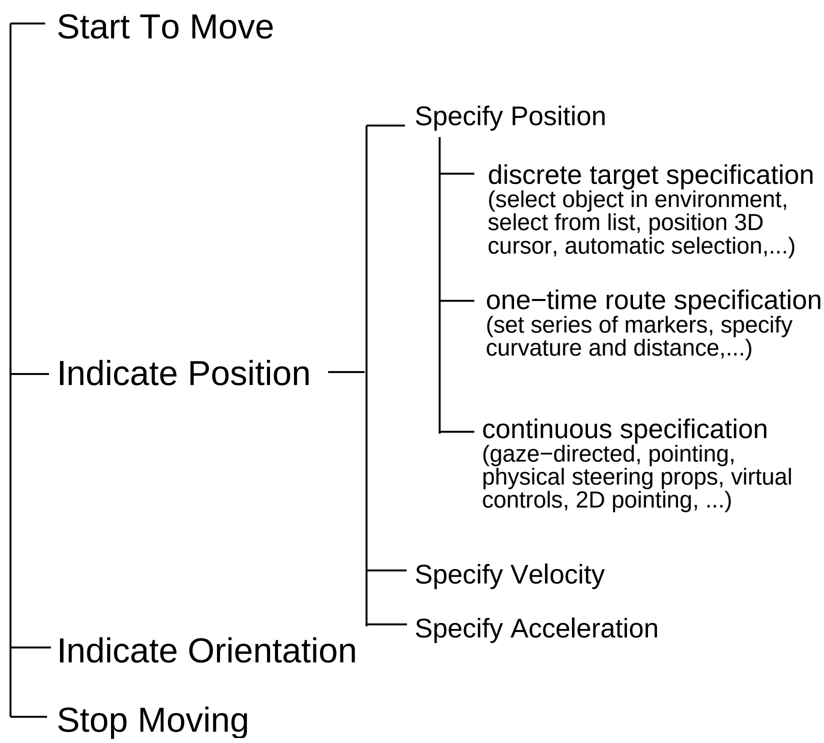

Figure 4.7 Alternate Taxonomy for Travel Techniques with Detail on Position Indication Subtask

We recognize that the task of viewpoint motion control consists mainly of two parts: setting viewpoint position and setting viewpoint orientation. Within each of these parts, there are two quite different methods one might employ. One might specify the destination or target position/orientation. This is a discrete, one-time selection. On the other hand, one might not have a target in mind at all, and might wish to simply specify a continuous trajectory for the position and/or orientation of the viewpoint. In between these two extremes are techniques we call “one-time route specification,” which allow the user to specify not only the endpoints of the motion, but also intermediate points or an entire path. The path specification is done prior to any actual movement of the viewpoint.

By separating these strategies in our task analysis, we can more closely and accurately fit the various technique components from our original taxonomy. This ensures that we compare techniques with tasks for which they are suitable, and that the design of new techniques using the framework will follow a logical progression.

4.6.2 Guided Design

This new taxonomy inspired the design of two new interaction techniques for travel, showing the usefulness of the concept of guided design (designing based on a taxonomy and not intuition alone).

First, we noticed in our review of published travel techniques that almost all of them fall into the continuous steering category, with a few discrete target selection techniques as well. However, the one-time route specification category was not represented. Therefore, we developed a simple route-planning technique which we tested experimentally against the other two metaphors (section 4.6.3). With our technique, the user holds a three-dimensional scale model of the environment, and places markers in this model using a stylus (figure 4.8). When a button is pressed, the system does a simple interpolation and takes the user along a piecewise linear path connecting these markers in the full-scale environment . Such a technique has the potential to be a good compromise between the amount of user control over travel and the amount of cognitive load placed on the user while moving, since all route specification is done prior to movement.

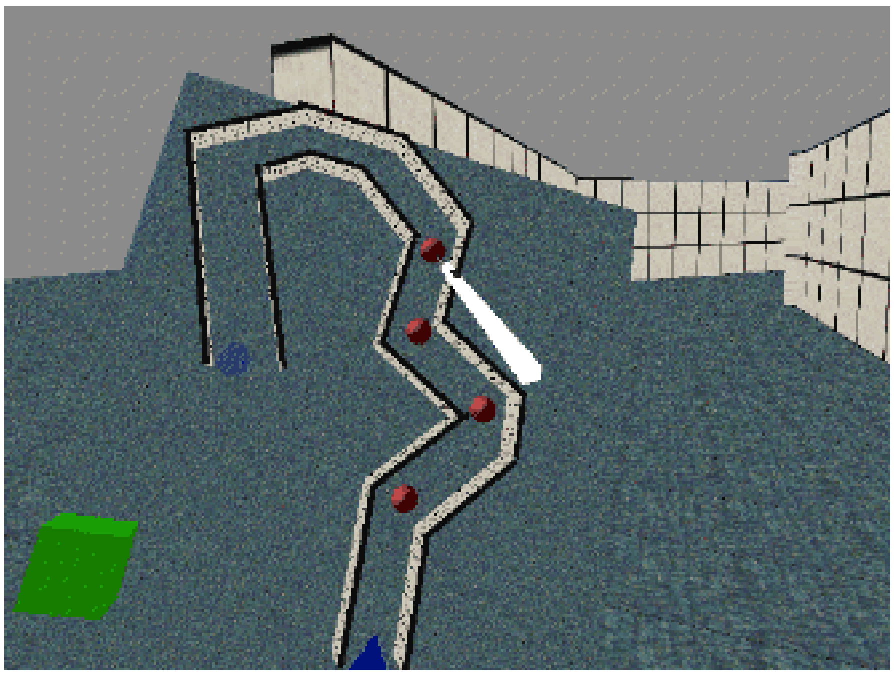

Figure 4.8 Route-planning Technique Using Virtual Map and Stylus

Second, the taxonomy shed new light on the continuous steering metaphor. Most of the existing techniques let the user specify a direction, and potentially a velocity (as in our earlier taxonomy). However, the subtask in the new taxonomy is “specify position,” which implies that viewpoint motion can also be thought of as a manipulation task. It is a simple step from this realization to a large set of potential travel techniques which are based on object manipulation techniques (see chapter 5). This is not a completely new idea (Ware & Osborne, 1990, Pierce et al, 1997), but is made more precise because of the taxonomy. We implemented a travel technique based on the HOMER object manipulation technique (section 5.3.3), in which the selected object is used as a pivot point for viewpoint movement.

4.6.3 Spatial Orientation Experiment

The new taxonomy also lent itself to a new evaluation of techniques. Using one technique from each of the three metaphors discussed above (steering, discrete target selection, route planning), we conducted a test to see which one would produce the highest levels of spatial orientation in users during and after travel. Spatial orientation refers to the knowledge that people have of their own location and orientation (direction) within a space.

Chance et al (1998) found that using a physical (walking) translation technique produced better spatial orientation, although the absolute error measurements they report were still relatively high. As in Chance et al, we chose to use a maze traversal task followed by a pointing task to measure spatial orientation. The mazes were actually corridors – they presented no choice points – and each contained three easily recognizable objects (Figure 4.9). At the end of each corridor, the subject was "virtually blindfolded" (the corridor and objects disappeared from view), and asked to point in the direction of one of these objects. The response variable was the angular error, in degrees, for this pointing task.

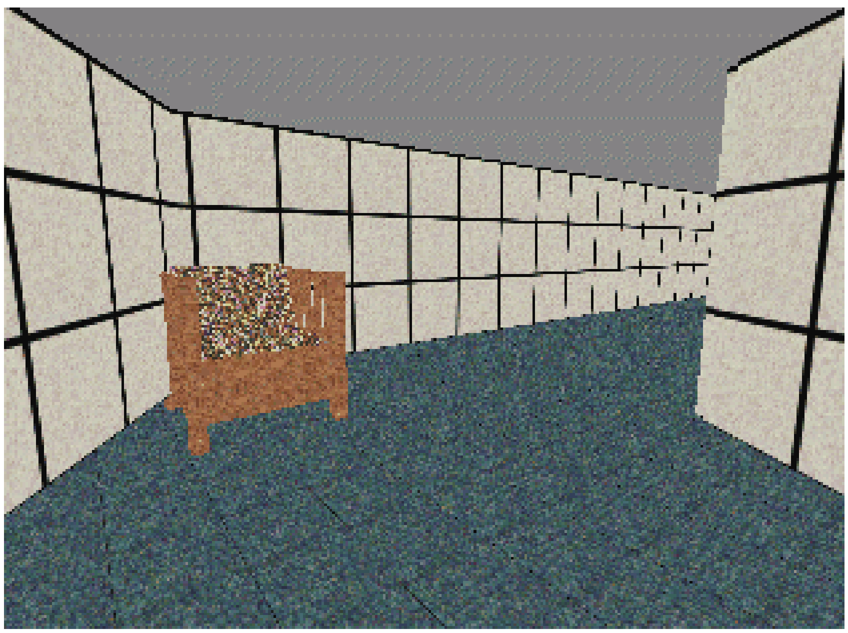

Figure 4.9 Inside View of a Corridor with a Target Object

The experiment, then, required users to pay close attention to the environment through which they were moving. The task might be performed using only route knowledge, along with the positions of the objects along the route, but survey knowledge of the corridor would make the task much easier. In order to maximize subjects' chances to acquire knowledge about the environment, we did not place any time restrictions on the corridor traversal, but rather allowed subjects to stop at any point and take as much time as needed.

The travel techniques were chosen as representatives of the three types of position-specification techniques from our new taxonomy. First, the *system-automated* technique gave users no control over their path. The system simply moved the subjects from the beginning of the corridor to its end on a route approximately in the center of the corridor at all times. The *pointing* technique allows users to continuously specify the direction of motion. Users point in the desired direction of motion. Finally, we chose the *route-planning* technique, in which users set a path before moving, and then are moved along that path by the system.

These three techniques represent different levels of user vs. system control of motion, and we hoped to discover which metaphor produces the highest level of spatial orientation. The extreme techniques are analogous to a driving example: the pointing technique lets the user "drive," while the system-automated technique simply makes the user a "passenger." The route-planning technique represents a compromise between the two.

We also included other variables that could potentially affect spatial orientation. Two factors relate to the complexity of the environment. As in the information-gathering experiment, we varied the dimensionality of the corridor (two or three dimensions). Two-dimensional corridors replicate the experience of moving through building hallways, while three-dimensional corridors also require ascending and descending. Second, some corridors had only ninety-degree (right angle) turns, while others turned at arbitrary angles. See Figure 4.10 for examples of the four corridor types. Finally, we examined conditions in which a three-dimensional map of the corridor was given to subjects before traversal versus trials with no map available*.

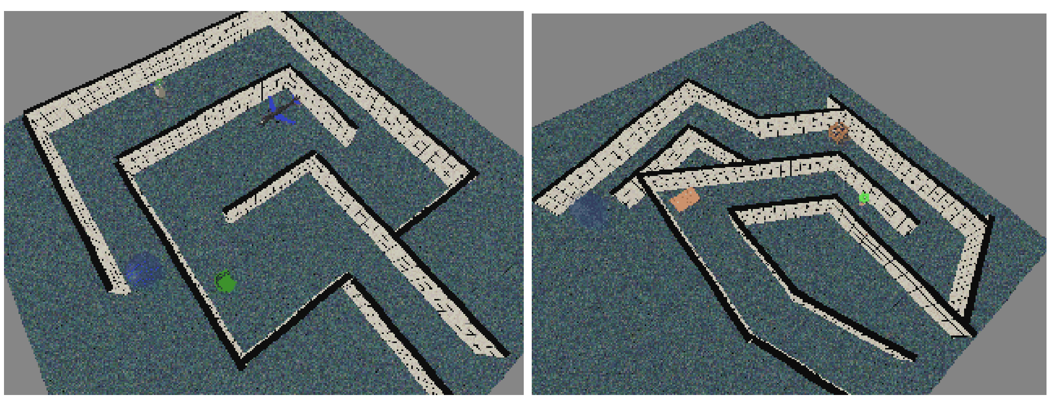

* We initially included a fourth factor: the presence or absence of a velocity control feature with which the user could speed up or slow down his rate of travel. Pilot testing, however, indicated that this factor was insignificant for the task, and it was dropped from the experiment.

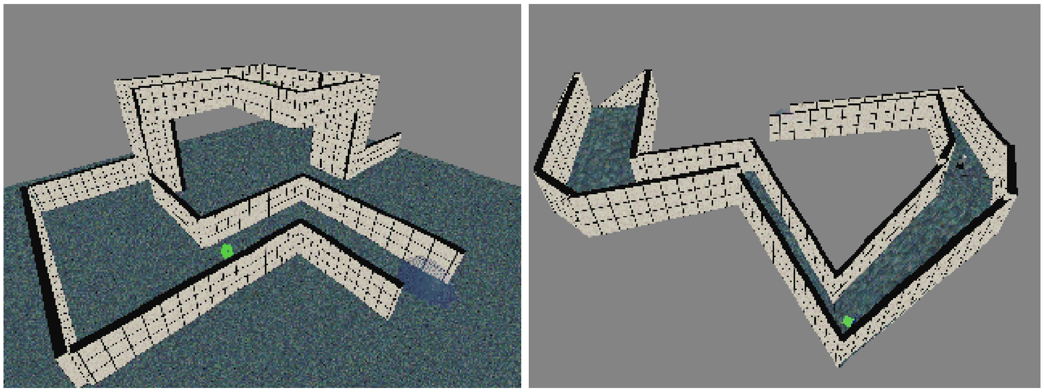

Figure 4.10 Views of Four Corridor Types used in the Spatial Orientation Experiment: Top left: 2D, right angles, top right: 2D, non-right angles, bottom left: 3D, right angles, bottom right: 3D, non-right angles

4.6.3.1 Method

The subjects for the experiment were 29 university students (23 males, 6 females), ranging in age from 18 to 24, with a mean age of 21.14. Eleven subjects reported some previous experience with immersive VEs. Subjects received extra credit in psychology or computer science classes for their participation. Three additional subjects did not complete the experiment due to simulator sickness.

Each subject completed a demographic questionnaire (Appendix A) before taking a standardized spatial ability test (the "cube comparison test" from the Educational Testing Service). This test measures 3D visualization and rotation skills, which are crucial to the experimental task. Before beginning the experiment, subjects were also shown a simple virtual world containing 24 common objects (chairs, tables, lamps, etc.) that would be used as stimuli during the experiment. This served the dual purpose of acquainting users with the head-mounted display (HMD) and tracking system, and introducing them to the objects they would need to know later.

The HMD used was a Virtual Research VR4, with a biocular display (same image to both eyes). The Polhemus Fastrak tracking system tracked the subject's head and two hands. Experimental software was built using the Simple Virtual Environment (SVE) library (Kessler, Kooper, & Hodges, 1998) and ran on a Silicon Graphics (SGI) Indigo2 MaxImpact at a near-constant frame rate of 25 frames per second.

Subjects also completed a set of preliminary VE tests designed to provide a benchmark for their ability to point to object locations in virtual space. In both the benchmark and main experimental tasks, pointing was accomplished using a tracked stylus, which is simply a tracker receiver embedded inside a pen. Users see a virtual representation of the stylus in the VE (see Figure 4.8) that moves in sync with the physical stylus so that the direction of pointing can be visualized. The stylus button is used to record answers. Two other receivers are used – one for head tracking and the other in the user's non-dominant hand where the 3D corridor maps may be viewed.

The benchmark tasks first measured the subjects' ability to point to visible objects in a sparse virtual world (each environment contained a "home" object which users looked at to begin a trial, and three target objects). Users pointed the stylus at one of the targets in response to an aural cue played through headphones. The second set of benchmark tasks required subjects to first study object positions and then turn away. When the stylus button was pressed, the objects disappeared and the subject would be asked to point in the direction of one of the objects. This more closely mimicked the main experimental task, which would require users to point blindly in the direction of a previously seen object. For both types of tasks, we presented trials in which all objects were on the same horizontal plane as the user, and trials in which objects might be anywhere in the 3D space surrounding the user. Subjects completed five trials of each type, for a total of 20 trials.

The main experiment compared the three travel techniques, the two corridor dimensionalities (2D & 3D), the two turn conditions (right angles vs. non-right angles), and the two map conditions (present or absent). All of these were within-subjects variables. There are only 20 valid combinations of these variables, as the route-planning technique uses a 3D map on every trial as a fundamental component of specifying a path through the corridor.

Subjects completed one trial for each valid treatment combination, with all trials using the same travel technique grouped together. Before each group, subjects were allowed to use the travel technique in a practice corridor as long as they wished. The order of travel techniques was counterbalanced between subjects. Within a set of trials using the same travel technique, the order of treatment combinations was randomized.

Corridors were chosen from a set of 16 (four corridors for each combination of the dimension and angle variables). Three objects were placed in each corridor at one of several pre-defined locations within the corridor. Subjects might encounter the same corridor layout more than once during the experiment, but never during the same travel technique group, and never with the same objects or object positions.

On trials where a map was present, the user was allowed to study and manipulate the map using a tracker in his non-dominant hand. Subjects were given as much time as they desired to study the map before movement started. When the subject began moving, the map disappeared.

The pointing technique was implemented using the tracked stylus. Users pressed the button to begin moving, then pointed the stylus in the desired direction of travel. In the system-automated technique, the user simply pressed the button to begin moving down the pre-defined path through the corridor.

For the route-planning technique, subjects used the stylus to place markers on the 3D map of the corridor to define a path (Figure 4.8). The path began at the corridor entrance, then moved in a straight line to the first marker, from there to the second, and so on. The last segment of the path took users from the location of the last marker to the end of the corridor. Subjects began motion by clicking the stylus button while touching a green box on the edge of the map.

While using any of the techniques, users could click the button while moving to stop, then click again to start moving. No collision detection was provided, so that subjects could travel through corridor walls. On two-dimensional corridors, users were constrained to a constant height above the floor, but they were allowed to move anywhere in virtual space while in three-dimensional corridors. Subjects always had complete control of their head orientation and gaze direction, and could look in any direction while moving or stopped.

At the end of each corridor, the visual representation of the corridor was removed, and subjects were presented with an audio stimulus instructing them to point in the direction of one of the objects seen in that corridor. Users estimated the direction by pointing the stylus and recorded their answer by pressing the button. We measured the angular error between the direction pointed and the actual direction to the object. Secondary response variables included time spent in each corridor, the number of times the user stopped in each corridor, and the strategies used to manage the spatial orientation task. We recorded subjects' strategies by observation only. This aspect of the experiment, though exploratory, proved quite interesting, and is discussed in detail below.

4.6.3.2 Results

In this section we present results of statistical analyses on the experimental data. The following section will explain and expand upon each of these results. Analysis of the experiment was split into two full factorial designs, since not all combinations of all factors were valid. Analysis 1 considered two techniques (pointing and system-automated), two dimensionalities, two turn conditions, and two map conditions, for a total of 16 treatment combinations. Analysis 2 considered all three techniques, two dimensionalities, and two turn conditions, with the 3D map always present, for a total of 12 combinations. The map variable could not be used in analysis 2 since a map was available on all trials with the route-planning technique.

We performed a repeated-measures analysis of variance for both of these experimental designs, on both the main dependent variable (angular pointing error) and the secondary dependent variable (time spent in each corridor). Results of both analyses for the error metric are summarized in Figure 4.11. Analysis 1 showed a significant main effect for dimension (mean 2D error: 32.47, mean 3D error: 38.62, $p < 0.005$), and a marginally significant main effect for the map variable (mean map absent error: 33.29, mean map present error: 37.80, $p < 0.075$). No significant differences between travel techniques or the turn conditions were found. Analysis 2 showed a marginally significant main effect for dimension (mean 2D error: 36.012, mean 3D error: 41.254, $p < 0.1$), but no other significant effects.

We performed the same analyses on the amount of time spent by subjects in each corridor. Both analysis 1 and analysis 2 showed significant main effects for technique ($p < 0.075$), dimension ($p < 0.005$), and angle condition ($p < 0.001$). Subjects spent longer amounts of time while using the pointing technique, when in 3D corridors, and when in corridors with right angles only. These results are summarized in Figure 4.12.

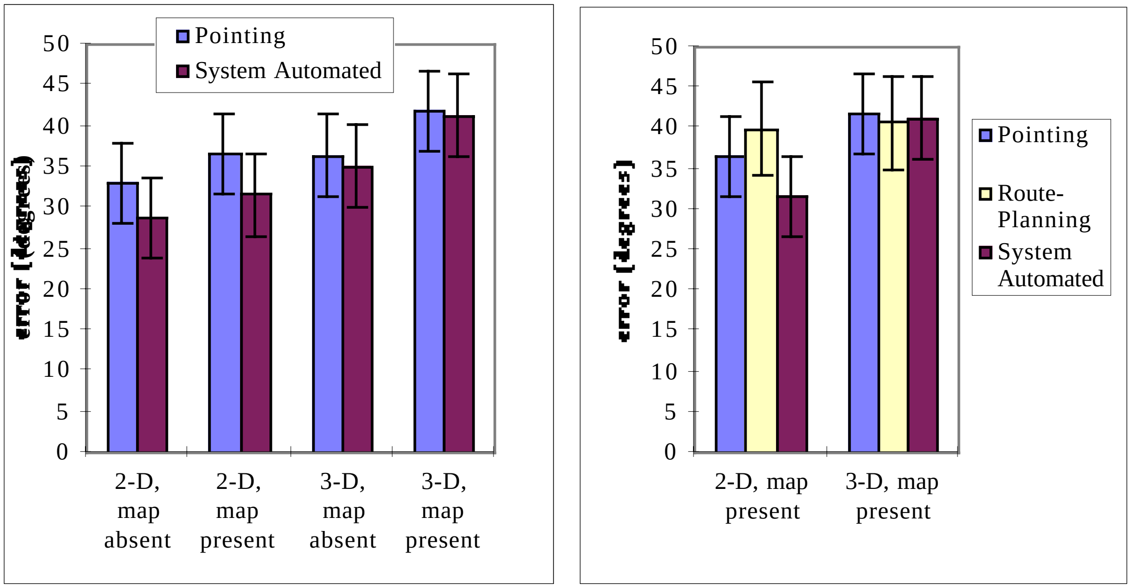

Figure 4.11 Mean Error in Various Treatment Combinations for Analysis 1 (left) and Analysis 2 (right)

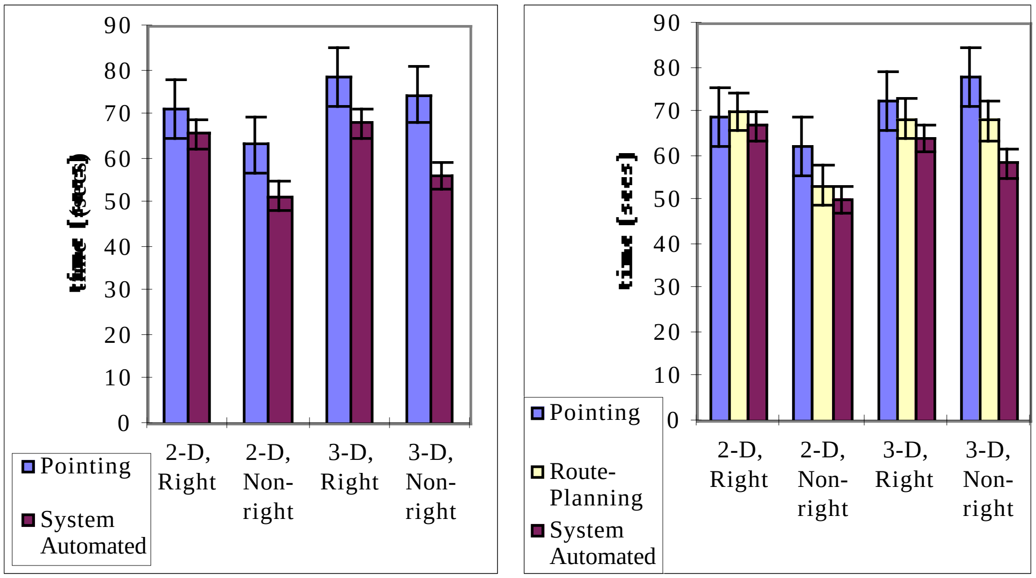

Figure 4.12 Mean Time in Various Treatment Combinations for Analysis 1 (left), and Analysis 2 (right)

We also included user strategies in our analysis of the experimental results. Our subjects were quite creative in their methods for minimizing error in the pointing task. We gave no suggestions to subjects about how to approach the task. We told them only the capabilities of the various travel techniques and that they should do whatever they felt necessary to point accurately at the target objects. We observed six main strategies during the experiment:

- *Stop & Look:* The simplest strategy, in which users simply stop moving at various points within the corridor and turn to look in other directions than the direction of motion.
- *Proprioceptive Pointing*: Users physically point in the directions of objects they have already seen, to give themselves a proprioceptive cue for later recall.
- *Backing in*: Users turn around just before the end of the corridor and move backwards to the end, ensuring that they see the corridor just before it disappears. This is possible with all techniques, but is quite difficult to do properly using the pointing technique.
- *Path retracing*: Users go back along the path they have just traversed through the corridor, both to remind themselves of what they have seen, and to see the corridor from another viewpoint. This strategy cannot be used with the system-automated technique, and requires careful thought to be used with the route-planning technique.

- *Moving through walls*: Users do not follow the corridor exactly, but instead move through corridor walls in order to better understand the relationships between adjacent passageways. A "lawnmower" strategy, in which the user simply travels along parallel lines through the space, is one example. Again, the system-automated technique does not allow this strategy.
- *3D overview*: Users fly up above the corridor to obtain a single view of the complete corridor and the objects within it, which might encourage a survey representation. This strategy is only available on 3D corridors using the pointing or route-planning techniques.

Table 4.2 shows the number of subjects using a strategy with a particular travel technique. One subject did not use any of the six strategies; another used every available strategy (six each for pointing and route-planning, three for system-automated). Subjects averaged 2.5 strategies for the pointing technique, 1.6 for the system-automated technique, and 2.2 for the route-planning technique.

For our analysis of the relation of strategy to error, we defined three between-subjects variables corresponding to the "level" of strategy sophistication for each subject on each technique. For pointing, this was a value between zero and three indicating the number of technique-specific strategies used (3D overview, moving through walls, path retracing). For the system-automated technique, the level was either zero or one indicating whether or not the subject used the backing in strategy. The route-planning strategy level ranged from zero to three indicating the use of the 3D overview, moving through walls, and/or backing in strategies.

Table 4.2 Number of Subjects Observed Using Common Strategies for Each Travel Technique

	Stop & look	Prop. pointing	Backing in	Path retracing	Through walls	3D overview
Pointing	22	12	5	10	7	15
System-automated	22	10	13	0	0	0
Route-planning	21	10	8	2	7	15

With the strategy level variables included in analysis 2, we found a large number of significant interactions indicating that the use of technique-appropriate strategies made a difference in the user's spatial orientation (error metric). For example, we found a significant ($p < 0.05$) interaction between technique, dimension, and the pointing technique strategy level. Subjects who had a pointing technique strategy level of zero had better scores with the system-automated technique than with the pointing technique, and did equally well on 2D and 3D corridors. Subjects with a pointing technique strategy level of one or two had approximately equal scores using all three travel techniques. Subjects with a pointing technique strategy level of three had better scores using the pointing technique than other techniques, and performed better on 3D corridors than 2D. This interaction suggests that strategy sophistication is significant in determining user performance.

We also analyzed the demographic and spatial ability information that we collected. The spatial ability test has a maximum score of 42, and our subjects averaged 25.862, with scores ranging from 4 to 41. This average is higher than reported means sampled from the general population and from college students. Regression analysis showed that spatial ability score was a significant predictor of average error in the experiment, and average error on each technique, corridor complexity, and map condition. Subjects with higher spatial ability performed better on the pointing task. Previous VE experience did not significantly predict these values. We also found that males performed significantly better than females, but are reluctant to draw conclusions from this due to our low number of female subjects. These results are consistent, however, with prior work. For example, Waller et al (1998) also found a significant effect due to gender in their experiment.

Finally, we analyzed the benchmark tests run on each subject before the main experiment. Both benchmark variables (visibility of objects and location of objects) produced significant differences ($p < 0.001$). Subjects averaged 12.8 degrees of error when objects were visible versus 26.7 degrees when object locations had to be remembered. Trials in which all objects were on the same horizontal plane had an average error of 14.6 degrees, while trials in which objects could appear anywhere had an average error of 24.9 degrees. Regression indicated that the error on trials with visible objects in the same horizontal plane was a significant predictor of error in the main experiment.

4.6.3.3 Discussion

The results presented above confirm that the spatial orientation of a user traveling through an immersive virtual environment depends on the complex interactions of many factors. None of the variables we studied proved solely responsible for the subjects' performance; rather, they all contributed in subtle ways. User strategy played an unexpected role in determining performance. In this section we will revisit and explain each of the major results.

The analyses of the angular error response variable showed that the dimension of the corridor was a significant effect – that subjects performed better on 2D corridors than 3D. Such a result is to be expected since the added complexity of the third dimension makes the corridor layout more difficult to comprehend and remember. Such 3D corridors are not familiar to subjects, but 2D corridors are seen in everyday life. This result also replicates our earlier finding in the information gathering experiment.

Interestingly, we found no significant differences between the performance of the three travel techniques. The system-automated technique produced the lowest average error, but the differences were not statistically significant. The overall mean error was 37.2 degrees for all conditions, which is lower than the mean error for the *best* technique (physical motion) reported in Chance et al (1998). This indicates that virtual travel techniques may indeed allow maintenance of good spatial orientation, although the error values are not directly comparable due to differences in the experiments.

We also found no main effect of the angle condition variable, though we expected that corridors with right angles only would be less complex and therefore produce lower error. However, it is also possible that the non-right angles served as more unique landmarks for subjects, allowing them to visualize their position in a corridor more effectively. These characteristics may have balanced each other so that we saw no significant effect from the angle condition.

In analysis 1 (considering technique, dimension, angle condition, and map), we also found a marginally significant main effect for the map variable, showing that subjects performed better when they were not given a map of the corridor before traveling through it. This result seems counterintuitive, since one would think a map would allow the user to form a better mental representation of the corridor layout. This result could be explained in three ways. First, subjects may have felt that the map gave them an advantage, and therefore did not concentrate as deeply when traveling through the corridor. The map did not show the locations of the three target objects, so subjects needed to integrate the object information with the corridor representation while traveling. Second, the map itself may have been another source of cognitive load, distracting subjects from the task rather than aiding them in it. Finally, we noted several subjects who did not make use of the map or gave it only a cursory glance before beginning through the corridor.

The time response variable also provided some useful information. Subjects using the pointing technique spent significantly longer in each corridor than they did using the other two travel techniques. However, this longer amount of time did not result in any performance gains for the pointing task, indicating that the pointing technique is more complex and requires more effort on the part of the user to maintain spatial orientation. 3D corridors were also shown to require more user time, again proving the difficulty of understanding such corridors. Even with the extra time spent in these corridors, the error was still significantly higher than in 2D corridors. Finally, we showed that subjects spent more time in corridors containing only right angles than in those with non-right angles. This is easily explained due to the fact that the corridors with right angles only were on average 20 meters longer than non-right angle corridors. When times are normalized by corridor distance, time in right-angle corridors is actually slightly lower.

The most interesting results are those pertaining to the strategies subjects used to maximize their performance on the pointing task. They are too numerous to go through one by one, but the example given above illustrates the importance of strategy. Those subjects who used no sophisticated strategies with the pointing technique (such as 3D overview or moving through walls) had better scores using the system-automated technique. However, those subjects with a high level of sophistication for pointing technique strategies actually performed better with pointing than the other two techniques. Furthermore, these subjects reversed the effect of the corridor's dimension by performing better on 3D corridors than 2D.

This gets at the heart of the contrast between the active pointing technique and the passive system-automated technique. Subjects who use the pointing technique naively, to take them directly through the corridor, will experience more cognitive load and thus will perform better with the system-automated technique, where the distraction of choosing a path is absent. On the other hand, subjects who take advantage of the unique characteristics of the pointing technique (the ability to move through walls, the ability to fly in three dimensions, the ability to retrace one's path) give themselves more and better opportunities to comprehend the layout of the space and thus will perform better with the pointing technique. Better performance on 3D corridors for these sophisticated users is explained by the fact that subjects were constrained to a constant height above the floor on 2D corridors, and therefore could not use the powerful 3D overview strategy.

Other significant interactions indicated the importance of the user's strategy when using the other two travel techniques, as well. It is overly simplistic, then, to say that one interaction technique outperforms another, although this may be the case in some situations. In general, though, it is more correct to say that one interaction technique *affords* the user more *opportunities* for high performance levels. Whether or not the user takes advantage of those opportunities is a major factor in determining user performance. For the travel techniques we studied here, the pointing and route-planning techniques give users more control, meaning more opportunities to understand corridor layout and object placement. It is more difficult, in general, to use the sophisticated strategies with the route-planning technique because the entire path must be specified in advance – the user cannot decide halfway along the path that she would like to go somewhere else. Therefore, in cases where subjects were highly sophisticated in all three techniques, performance should be highest using the pointing technique. Indeed, among the two most sophisticated groups of subjects, the average error for pointing was lower (20.32 degrees) than mean error values for the route-planning (21.76 degrees) and system-automated (25.71 degrees) techniques, although these subjects did extremely well using all three techniques.

These results provide two important guidelines (hypothesis 2) for developers of interaction techniques for VE applications. First, the techniques should provide affordances for high performance on an application's main tasks. Second, the users must be trained to take advantage of the opportunities – to use strategies that will help them achieve the desired performance levels. For tasks where spatial orientation is especially important, it appears that a travel technique giving users complete control over their position, such as pointing, can produce high performance levels given that appropriate strategies are used. If it is not possible or practical to train users in those strategies, it may be more beneficial to use a passive travel technique inducing lower cognitive load.

4.7 Travel Testbed

After all of these preliminary experiments, we implemented and ran a testbed evaluation for the task of travel. This testbed is based both on our alternate taxonomy and the formal framework presented earlier, including the lists of outside factors and multiple performance metrics. The testbed is designed to allow experimentation with any travel technique on a wide variety of travel tasks. However, we implemented only two search tasks that were especially relevant to our target application.

These tasks are simple and general, being found in a wide variety of VE applications. Darken (1995) characterizes the two as *naïve search* and *primed search*. Naïve search involves travel to a target whose location within the environment is not known ahead of time. Primed search involves travel to a target which has been visited before – if the user has developed a good cognitive map of the space and is spatially oriented, he should be able to return to the target. We would also like to test *exploration*, in which the user is simply moving about with no specific target, but it would be very difficult to quantify performance on such an open-ended task.

4.7.1 Method

We created a medium-sized environment (one in which there are hidden areas from any viewpoint, and in which travel from one side to the other takes a significant amount of time). The size and type of the environment could be variable if this was deemed an important outside factor on performance, but we left it constant in our implementation. We also built several types of obstacles which could be placed randomly in the environment. These included fences, sheds, and trees (figure 4.13).

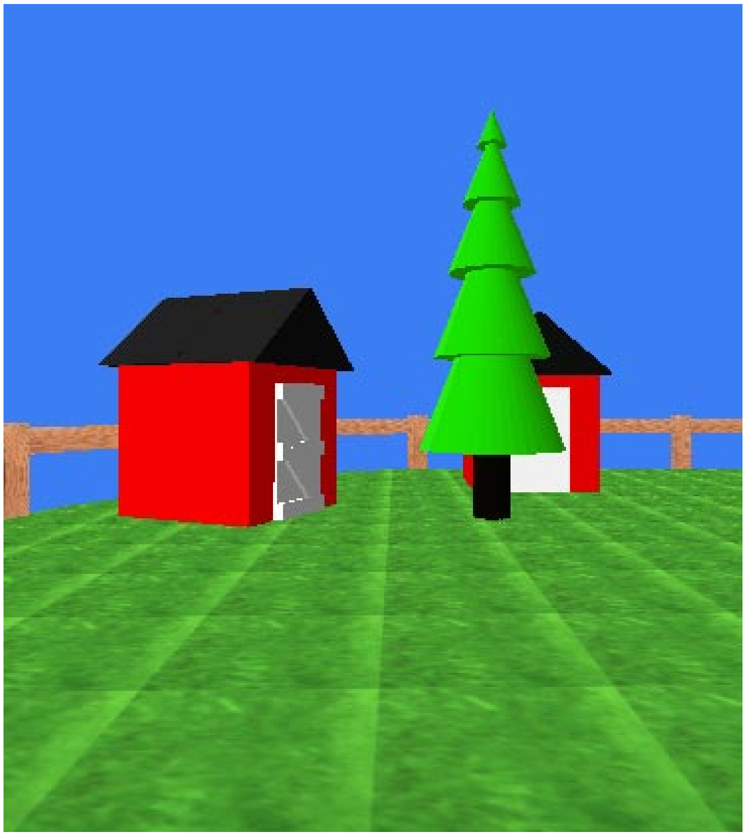

Figure 4.13 Example Obstacles from the Travel Testbed Experimental Environment

Targets for the search tasks were flags mounted on poles. The targets were numbered one to four in each instance of the environment, and each number was associated with a flag color so that the user would be able to identify the targets from a distance. Each target also had a circle painted on the ground around it, indicating the distance within which the user would have to approach to complete the search task (figure 4.14). There were two sizes of this circle: a large one (ten meter radius) corresponding to low required accuracy, and a small one (five meter radius) corresponding to high required accuracy.

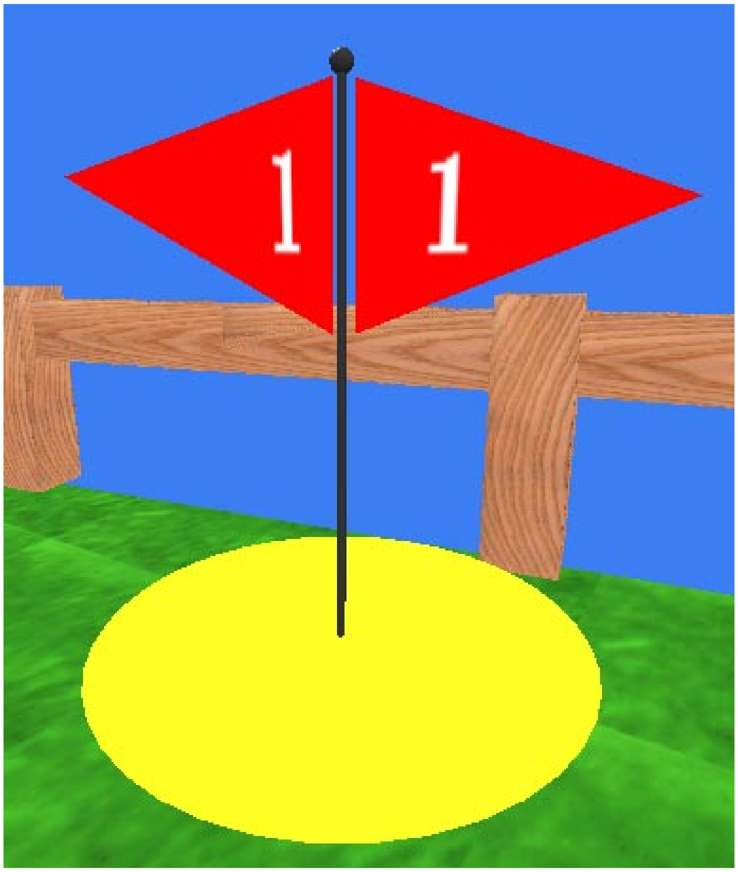

Figure 4.14 Target Object from the Travel Testbed Experimental Environment Including Flag and Required Accuracy Radius

Each subject completed 24 trials – eight trials in each of three instances of the environment. Each environment instance had the same spatial layout, but different numbers and positions of obstacles, and different target positions. In each environment instance, the user first completed four naïve search trials, in which the four targets (numbered one to four) were to be found in order. Before each trial, the flag number and color were presented to the user. During this phase, targets only appeared as they were needed. That is, during the first trial only the first target was visible, during the second trial, targets one and two were visible, and so on. This was to ensure that subjects would not see a target before its trial, thus changing a naïve search to a primed search. The first trial began at a predefined location, and subsequent trials began at the location of the previous target. In each of these trials, the target was not visible from the starting location, and the required accuracy radius was at its low level.

In the second phase, subjects completed four primed search trials where they returned to each of the four targets once, not in numerical order. Again, the flag number/color stimulus was presented to users before each trial. During these trials, all targets were present in the environment at all times, since the subjects had already visited each target and these were therefore primed search trials. Two factors were varied (within-subjects) during these trials. First, we varied whether the target could be seen from the starting position of the trial (visible/invisible). Second, we varied the required accuracy using the radii around each target. Each of these variables had two levels, and therefore there were four possible combinations, and one trial of each of these combinations during each environment instance.

Seven travel techniques were implemented and used in this experiment. Travel technique was a between-subjects variable. Three were steering techniques: pointing, gaze-directed, and torso-directed, as described in the information gathering experiment.

We also implemented two target-specification techniques. In the ray-casting technique, the user pointed a virtual light ray at an object to select it, and then was moved by the system from the current location to that object. The second target-specification technique involved dragging an icon on a two-dimensional map held in the non-dominant hand. The map shows the layout of the environment and an icon indicating the user's position within the environment (figure 4.15, top). Using a stylus, the user can drag this icon to a new location. When the icon is released the user is flown smoothly from the current location to the corresponding new location in the environment. Both the stylus and the map have both physical and virtual representations (figure 4.15, bottom). This technique was one of the travel metaphors used in our target application at the time (section 6.3.2). With both the ray-casting and map techniques, the user could press a button during movement to stop at the current location.

Finally, we studied two manipulation-based travel techniques, as described in section 4.6.2: one based on the HOMER technique (section 5.3.3) and another on the Go-Go technique (section 5.3.1). With the HOMER technique, the user selects an object using ray-casting, then uses hand movements to move the viewpoint around that object. The Go-Go technique uses a non-linear mapping to allow the user to stretch his virtual hand far away from his body. The user clicks a button to "grab the air" at the current location of the virtual hand, and then uses hand motions to move the viewpoint around the environment.

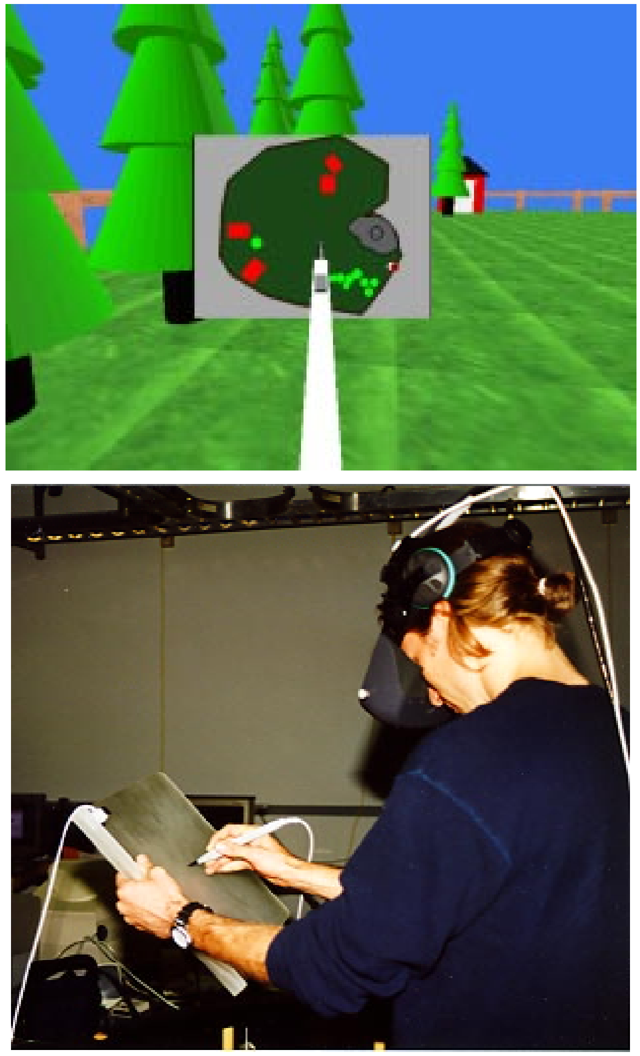

Figure 4.15 Virtual (top) and Physical (bottom) Views of the Map Dragging Travel Technique

For each subject, we measured the total time taken to complete each trial (broken into two parts: the time between the onset of the stimulus and the beginning of movement, and the actual time spent moving). We assumed that this first time would correspond to the time spent in mental processing (perception of the stimulus and environment, and cognitive effort to remember where a target was last seen in the primed search task). This is not entirely accurate, as wayfinding activities undoubtedly continue after a subject's travel has begun. Therefore, the absolute measurements here are not meaningful, but the relative differences between techniques may give some indication of the amount of perceptual/cognitive processing necessary to move to a certain location or in a certain direction using a technique. We have labeled this measure "think time" in the analyses to follow. We also obtained subjective user comfort ratings in the areas of arm strain, hand strain, dizziness, and nausea. After each environment instance, the subjects gave a rating for each of these factors on a ten-point scale (Appendix B). Each subject also took a standardized test of spatial ability (the ETS cube comparison test). Finally, we gathered demographic information about our subjects, including age, gender, handedness, technical ability, and VE experience (questionnaire in Appendix A).

Forty-four undergraduate students were recruited from the department of Psychology subject pool to participate in the experiment. Four subjects did not complete the experiment due to sickness or discomfort, and two subjects did not complete the experiment due to computer problems. Thus, 38 subjects (32 males, 6 females, mean age 19.7) completed the evaluation, meaning that each technique was used by at least five subjects.

4.7.2 Results

In this section, we will present the most relevant results from the travel testbed experiment. For complete tables of results, see Appendix C.

We performed a one-way analysis of variance (ANOVA) on the results for the naïve search task, with travel technique as a between-subjects variable. Table 4.3 gives the results for the naïve search task for each technique.

Table 4.3 Mean Completion Times (seconds) for Naïve Search Task (Standard Deviation in Parentheses)

Technique	Think Time	Travel Time	Total Time
Gaze-directed	2.16 (1.10)	18.28 (4.63)	20.44 (5.24)
Pointing	2.20 (0.92)	22.33 (6.98)	24.53 (7.88)
Torso-directed	2.77 (0.63)	27.00 (6.27)	29.77 (6.49)
HOMER	4.20 (1.00)	37.66 (5.65)	41.86 (6.31)
Map dragging	29.54 (11.58)	52.39 (13.11)	81.93 (18.61)
Ray-casting	1.86 (0.48)	34.95 (8.89)	36.81 (8.43)
Go-Go	3.29 (2.43)	21.48 (6.86)	24.77 (8.73)

For each of the three time measures (think time, travel time, and total time), the travel technique used had a statistically significant effect ($p < 0.001$). We also performed post-hoc comparisons of techniques (LSD and Bonferroni), and found that for the think time measure the map dragging technique was significantly slower than all other techniques. This makes intuitive sense, since the map technique is based on the target-specification metaphor, where movement must be planned before it is carried out. The ray-casting technique also has this property, but selection of a single object is much faster than dragging an icon through an entire route. With the other techniques, movement could begin immediately. However, because the difference is so large, we feel that there may be another factor at work here. The map technique requires users to mentally rotate the map so that it can be related to the larger environment. This mental rotation induces cognitive load on the user, which may cause them to be unsure of the proper direction of movement. The increased cognitive load can be seen directly in increased thinking time.

In the travel time measure, using the same post-hoc tests, we found that the pointing and gaze-directed steering techniques and the Go-Go technique were significantly faster than HOMER, ray-casting, and map dragging. The torso-directed steering technique was significantly faster than HOMER and map dragging. In general, then, steering techniques performed well at this task because of their directness and simplicity. Users could look at the environment, determine where they wanted to search next, and then go there with little or no thought required. The torso-directed technique performs slightly worse, as we found in the information gathering experiment. We believe this is purely a function of mechanics. The user of the torso-directed technique must move his feet and whole body to change directions, while the other steering techniques require only movements of the head or hand. It is also interesting that the Go-Go technique performed well here, but HOMER did not, since they are both manipulation-based travel techniques. The difference seems to be that HOMER requires an object to move about, while the Go-Go technique allows the user to simply grab the air and pull himself forward. Again, the map dragging technique performed poorly. It is simply not suited for exploration and naïve search, because it assumes the user has a distinct target in mind.

Total time results for the naïve search task were almost identical to the results for the travel time measure, since most of the time was spent moving. Again, pointing, gaze-directed steering, and to a lesser degree, Go-Go, performed significantly better than other techniques.

For the primed search task, we performed a multi-variate analysis of variance (MANOVA), with technique as a between-subjects variable and visibility (two levels) and required accuracy (two levels) as within-subjects variables. Travel times were normalized relative to the distance between the starting point and the target (this was not necessary for the naïve search task since subjects in that task had no knowledge of the location of the target and thus did not move in straight lines). Table 4.4 presents a summary of results for this task. We do not list results for the two levels of required accuracy independently, because this factor was not significant in any of our analyses. Results for think time mirrored the naïve search task. Again, technique was significant ($p < 0.001$), with the map dragging technique significantly slower in post-hoc comparisons (LSD and Bonferroni) than all other techniques, for the same reasons given above. Neither of the within-subjects factors was significant in predicting think time.

Technique was also significant for the travel time measure ($p < 0.001$). Here, post-hoc tests showed that pointing and gaze-directed steering were significantly faster than HOMER, ray-casting, and the map technique. Again, these techniques allow the user to form a direct mapping between the desired direction of motion and the action that needs to be taken (look or point in that direction). The map technique performed badly, but it was only significantly worse than gaze-directed steering, pointing, and Go-Go. We had expected that the map would be useful for the primed search, since it allows users to specify the location of the target and not the direction from the current location to the target. However, this assumes that the user understands the layout of the space, and that the technique is precise enough to let the user move exactly to the target. In the experiment, the size of the target was not large enough, even in the low required accuracy condition, to allow precise behavior with the map technique. We observed users moving directly to the area of the target, but then making small adjustments in order to move within the required range of the target. However, the best results with the map did occur in trials with low required accuracy and a target not visible from the starting location. We also found that visibility of the target from the starting location was significant here ($p < 0.001$). Trials in which the target was visible averaged 12 seconds, as opposed to 23 seconds for trials in which the target was hidden.

Total time for the primed search task produced similar results. Again, technique was significant ($p < 0.001$), with the gaze-directed steering and pointing techniques performing best, according to the post-hoc comparisons. Visibility also significantly affected total time ($p < 0.001$). Another technique that we expected to perform well in the primed search task was ray-casting, since it allows the user to move directly to a target. This should especially hold in cases where the target is initially visible. We believe these results were not found due to our implementation of targets as flags. The flagpoles were very thin, and thus impossible to select at any distance. The flags themselves were larger, but due to the size of the environment might appear very small from the starting location. Thus, users of the ray-casting technique often had to select an intermediate target in order to get close enough to select the flag.

Table 4.4 Mean Completion Times (seconds) for Primed Search Task, with Targets not Within View from Start Location (Invisible) or In View from Start Location (Visible)

*(*normalized times – seconds per 100 meters)*

Technique	Invisible think time / Visible think time	Invisible travel time* / Visible travel time*	Invisible total time* / Visible total time*
Gaze-directed	1.69	10.52	12.21
	1.49	4.70	6.18
Pointing	2.30	10.20	12.49
	2.03	5.61	7.63
Torso-directed	2.95	22.87	25.82
	1.40	5.81	7.21
HOMER	3.85	26.34	30.19
	2.67	13.81	16.48
Map dragging	20.58	25.07	45.65
	14.01	18.97	32.98
Ray-casting	2.09	29.69	31.78
	1.92	13.72	15.64
Go-Go	2.66	17.55	20.21
	1.72	7.36	9.09

We also performed an analysis that compared the two types of tasks. For this analysis, technique was again a between-subjects variable, while task was a within-subjects factor. We only considered the trials in which the target was initially visible and the required accuracy was low, since these were the conditions in all of the naïve search trials. For the travel time measure, we found that task was significant ($p < 0.001$), with the naïve search taking 30 seconds on average vs. 23 seconds for the primed search. For the think time measure, task was not significant, but we did find a significant interaction between task and technique ($p < 0.025$). This interaction is due to the fact that the amount of think time for the map technique drops significantly for primed search trials (figure 4.16 – error bars have been omitted in the figure for readability), while think time for the other techniques remains approximately the same. This indicates that subjects had learned the layout of the space and were more confident in the map dragging task because they knew the area in which the target lay. For each of the significant results reported above, the observed statistical power was 0.987 or greater, with alpha = 0.05.

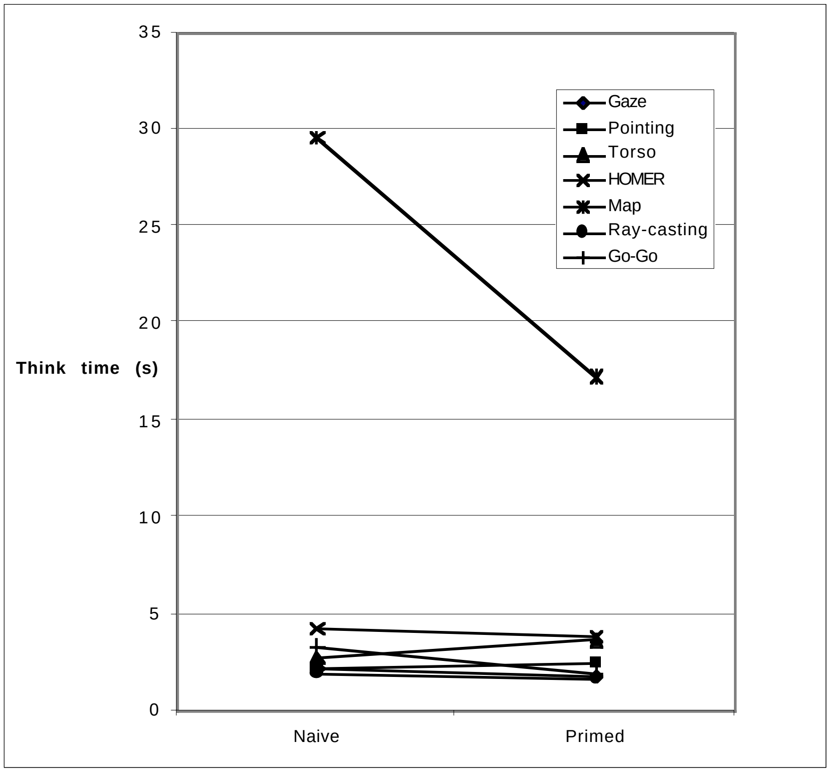

Figure 4.16 Interaction Between Task Type and Technique for Think Time on Search Tasks

Our evaluation showed that if the most important performance measure is speed of task completion, that steering techniques are the best choice. Users also seem to prefer these techniques over others. Of the steering techniques, pointing is clearly the most versatile and flexible, given our earlier results comparing it to gaze-directed steering and torso-directed steering. The Go-Go technique also performed well in this study with respect to speed. However, upon analysis of our comfort rating measures, we found that Go-Go produced arm-strain, dizziness, and nausea in some users. These problems were also seen with the HOMER technique, suggesting that viewpoint movement using hand-based manipulation may be discomforting to users because it is so different from the normal methods of movement. Gaze-directed steering also produced some significant discomfort (mainly dizziness), likely because it requires rapid and repeated head movements. The visual scene lags behind head movements due to tracker latency, so these could be the cause of discomfort. Of the seven techniques, only pointing and ray-casting produced no significantly high discomfort levels.

We also analyzed the demographic data and found no correlation between task performance and age, gender, VE experience, or spatial ability. These appear to be tasks whose speed is largely determined by the physical interaction technique used rather than individual differences.

As discussed above, the map technique was the most disappointing technique in this study. It seems to be well-suited for low precision, goal-directed travel. We believe that this technique would have performed better if the required accuracy had been lower on certain trials. It would probably also benefit from the use of a “view-up” map as opposed to a standard “north-up” map (Darken & Cevik, 1999). Performance on the primed-search would likely increase because of its egocentric nature. However, we have other reasons for using a north-up map, including the fact that it is a fixed frame of reference within a dynamic environment, and thus may facilitate learning of the spatial layout more quickly (Wickens & Baker, 1995). The map technique is also useful for other tasks, such as object manipulation (see chapter six), and so we do not believe that this technique should be removed from consideration as a result of its performance in this evaluation.

Finally, we also noted a reoccurrence of the theme of user strategies in this evaluation. No collision detection was implemented in the experimental environments, so users could move through objects if desired. In certain cases, this was highly advantageous, for example when the flag was just on the other side of a large fence. We noted that subjects using this strategy performed better on the primed search task, because they could take a straight-line path to the target. We also observed that certain techniques afford this strategy more than others. Steering techniques in general do not afford this, as they more closely mimic natural movement. Subjects using steering techniques generally went around obstacles. More unnatural techniques such as map dragging, Go-Go, HOMER, and ray-casting seem to suggest to the user that this environment does not work in the same manner as the physical world, and that therefore moving through objects is allowed. This represents another benefit of so-called “magic” techniques.

4.8 Summary

In this chapter, we have presented the results of our design and evaluation of viewpoint motion control, or travel, techniques for immersive VEs. Because of its pervasiveness, it is essential that we understand this task and the space of techniques for it. To this end, we have presented a formal framework for the design and evaluation of travel techniques, including two alternate taxonomies, performance metrics, and outside factors that could influence performance. Within this framework, we have designed new techniques and evaluated a wide range of techniques in six experiments. These evaluations, in particular the testbed experiment, have produced guidelines and empirical results that will allow application developers to choose appropriate travel techniques. We present such a practical application of these results in chapter six.

We learned several important lessons from our evaluation of travel techniques. From the relative motion experiment, we learned that techniques that do not couple the user’s head and the direction of motion are more efficient for relative motion tasks. The spatial awareness experiment showed that teleportation can cause disorientation in users. The information gathering experiment indicated that path complexity affects a user’s ability to obtain information from an environment. The spatial orientation experiment showed that users’ strategies, in conjunction with the affordances of the travel technique, can affect spatial orientation performance. Finally, the testbed evaluation indicated that steering techniques are generally efficient for search tasks.