

Doug A. Bowman
Department of Computer
Science
Virginia Tech
Blacksburg, VA 24061
bowman@vt.edu

Elizabeth T. Davis
Larry F. Hodges
Albert N. Badre
Georgia Institute of Technology

Maintaining Spatial Orientation during Travel in an Immersive Virtual Environment

Abstract

An experiment is presented comparing the effects of various virtual travel techniques in an immersive virtual environment (VE) on the spatial orientation of users. The experiment was designed and implemented in the context of a formal framework for the design and evaluation of VE travel techniques. Subjects traveled through virtual corridors, noting the location of objects along the way, and were asked to point in the direction of one of these objects when the end of the corridor was reached. Results indicate that virtual travel techniques, in which users do not physically translate their bodies, can allow the maintenance of a user's spatial orientation as measured by a pointing task. The experiment also replicates an earlier result, showing that path dimension significantly affects user performance. Finally, the strategies used by subjects to perform the task were shown to be significant, indicating that performance depends not only on the technique, environment, and task, but also on the sophistication of the user.

1 Introduction and Related Work

Users of large-scale or highly occluded virtual environments (VEs) often experience difficulty in maintaining their spatial orientation (sense of position and orientation) within the virtual world. VEs do not in general provide the same rich set of cues for distance, motion, and direction found in the physical environment. Also, immersive VEs using head-mounted displays (HMDs) generally have a small field of view, further reducing the visual cues used to maintain spatial orientation. This problem becomes important for certain applications of VEs (such as training for navigation in a real-world environment), because the maintenance of spatial orientation is essential for the task at hand: acquiring knowledge of the layout of the virtual world that can be applied to navigation in the physical world.

Previous research in this area has looked at various augmentations to the VE that aid the user in remaining spatially oriented. For example, Darken and Sibert (1996a, 1996b) used view-up maps (which rotate to match the orientation of the world), grid lines, and other organizational structures drawn directly over virtual terrain, and also directional cues such as signposts to aid users in a wayfinding task. Such cues provide advantages to users navigating through large-scale VEs. They also discuss the influence of spatial ability and other individual differences on navigation.

Other experiments have compared training in a VE to training in the real

Presence, Vol. 8, No. 6, December 1999, 618–631
© 1999 by the Massachusetts Institute of Technology

world or with a map (for example Waller, Hunt, and Knapp (1998)). Such experiments have generally found VE training to be inferior for the acquisition of spatial knowledge. For example, Waller et al. found that VE training proved equal to real-world training on a simple path only after a long VE exposure period (five times longer than the real-world training period). They also state that VE training allows the user to obtain *route knowledge* (information about a specific path through the environment), but does not allow the acquisition of *survey knowledge* (overall information about the structure and layout of the environment). However, there do not seem to be any fundamental characteristics that would prevent VE users from obtaining survey knowledge or would restrict the efficacy of VE training.

A potential difficulty with experiments of both types is that they generally do not consider the method of locomotion, or *travel*, through the virtual world. We define *navigation* as the complete process of moving through an environment. Navigation has two parts: *wayfinding* (the cognitive decision-making process by which a movement is planned), and *travel* (the actual motion from the current location to the new location). Obviously, people navigating in both real and virtual worlds do not explicitly separate their navigation in this way; rather, they are constantly gathering environment information, planning movements, and executing them simultaneously. Since travel and wayfinding are so closely related and used in concert with one another, the method of travel may have an effect on the ability to perform wayfinding tasks, and thus on a user's spatial orientation.

Such effects can be easily observed in the real world. In most people's experience, the driver of a car maintains better spatial orientation in a new town than does a passenger, because the driver must attend to the spatial information in steering the car. On the other hand, if the passenger attends carefully with the intention of acquiring spatial knowledge, he may maintain better spatial orientation than the driver, because he is not distracted by the task of driving the car. In either case, the method of travel is important in determining spatial orientation.

The studies mentioned above all assumed a simple travel technique in the VE conditions, such as constant velocity gaze-directed steering (the user moves in the

direction she is looking). However, a recent study (Chance, Gaunet, Beall, & Loomis, 1998) addresses the importance of the travel technique in a path-integration task, by comparing techniques that differed in their level of similarity to movement (walking) in the physical world. They found that, in general, a technique more similar to real walking (physical translations and/or rotations) produced better spatial orientation in subjects than those in which self-motion was virtual (the subject's view translated or rotated while the subject remained still). The amount of angular error in this study was relatively large, however, even with physical translations and rotations. In a similar vein, Slater, Usoh, and Steed (1995) found that physical motion had a positive influence on reported levels of presence in a VE.

In this paper, we present an experiment that further explores the effects of travel techniques on spatial orientation in an immersive VE. In particular, our study distinguishes between different types of virtual travel techniques. In a typical immersive VE setup (using a tracked head-mounted display), it is simple to use the user's physical head rotations to specify gaze direction in the VE, but the use of physical translation is usually limited to the range of the tracking system. Therefore, a fully physical travel technique is not generally possible unless some motion scaling is performed. We will show that certain virtual travel techniques, in which translations are completely virtual, can produce good spatial-orientation scores in an object-pointing task. Our experiment also considers the effects of user strategy, 3-D virtual maps, path dimension, and turn complexity on spatial orientation.

The next section summarizes an evaluation and design framework within which the current experiment and others have been performed. We then present our study along with a detailed discussion of primary and secondary results.

2 Design and Evaluation Framework

The design and evaluation of interaction techniques (ITs) for immersive VEs seems as much art as it is science. Researchers have a great deal of anecdotal and

experiential evidence that users enjoy certain techniques, or that a given technique worked in an application. However, the systematic design and evaluation of ITs for VEs is a research topic only in its infancy. The interaction we consider in this paper, travel, is found in almost any VE application to allow the user to move his viewpoint to various positions within the 3-D scene. Because of this ubiquity, it is essential that we understand travel techniques and their effects on the performance of various user tasks.

For this reason, we have created a formal framework within which design and evaluation of travel techniques can take place. The framework—consisting of taxonomies, performance metrics, and outside factors—gives us a common reference point and structure to guide our thinking regarding VE travel techniques. The framework suggests design ideas and helps in the planning of experiments to measure the performance of techniques.

2.1 Taxonomies

The first aspect of our formal design and evaluation framework is a set of taxonomies, which categorize techniques and partition the design space into smaller, more easily understandable pieces. Our original taxonomy (Bowman, Koller, & Hodges, 1997) divided the task of travel into three subtasks: direction or target selection, velocity and acceleration selection, and input conditions. Several interaction options for each of these subtasks were listed. For example, one might specify velocity with a continuous gesture (for example, the distance between two fingers), or by choosing from a discrete set of possibilities (for example, choose one of five velocities from a menu).

This taxonomy was useful in three ways. It helped us understand the space of possible techniques; it guided us in designing new techniques (choose a component for each of the three subtasks; and it helped generate ideas for experimental evaluation of techniques.

However, no taxonomy can make clear all of the fundamental differences among interaction techniques. Rather, they are tools that each provide a different view and structure for the same space of techniques. Our current experiment was inspired by a new taxonomy, shown

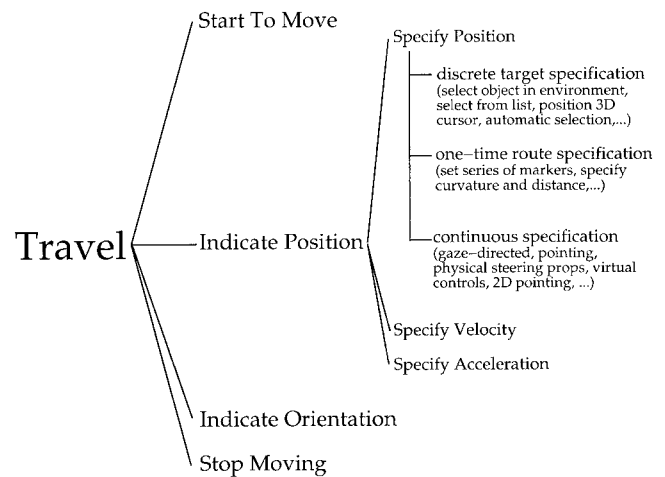


Figure 1. Taxonomy of virtual travel techniques, with detail shown for the 'specify position' subtask.

in part in Figure 1. This taxonomy considers travel to be the task of setting one's viewpoint position and orientation in a 3-D space. Specification of position can be accomplished in one of three ways: discrete target specification (the user specifies the goal of the movement and the system moves the user there in some way), continuous specification (the user is constantly steering—setting trajectory or direction of motion), or one-time route specification (the user defines a path through the space and the system moves the user along that path). The taxonomy lists some ways these general techniques could be implemented.

For immersive VEs, our interest is mainly in this position-specification task, because viewpoint orientation is generally implemented using head-tracking. Desktop VEs must also be concerned with the specification of viewpoint orientation, adding additional complexity to travel techniques. We do not consider desktop VEs in this work.

This taxonomy covers approximately the same space of travel techniques as the original, but addresses different distinctions between techniques. In particular, this taxonomy separates travel techniques by the amount of control the user has over her motion. With discrete target specification, the user controls only the two endpoints of motion, and leaves the path between those points up to the system. With continuous specification, the user maintains complete control over every part of

her path. One-time route specification represents a compromise: the user defines the path of viewpoint motion, but has no control over motion during motion itself. These distinctions lead to interesting questions about the attention and awareness of the user, as we will discuss below.

2.2 Performance Metrics

The second component of our formal design and evaluation framework is a set of performance metrics. Experimentally, these are the response variables by which one measures the performance of a travel technique in a given situation. Application designers could test various travel techniques by implementing them in the target application and then running usability studies. However, this is a time-consuming process whose results are valid only for the tested application. We propose instead that travel techniques can be evaluated more generally and their performance measured in terms of the metrics we provide. Application designers can then specify the desired levels of performance for one or more of the metrics, and choose a travel technique that has been shown experimentally to achieve those levels. This indirect mapping of techniques to applications allows the results of a single evaluation to be used many times.

For travel, a large set of possible performance metrics represent different application goals and different uses for the travel interaction. Some possibilities are

1. **Speed:** time for task completion
2. **Accuracy:** proximity to the desired target or path
3. **Spatial orientation:** user's knowledge of position and orientation within the environment
4. **Ease of learning:** the ability of a novice user to learn the technique
5. **Ease of use:** the ability of a user to form accurate mental models with low cognitive load
6. **Information-gathering potential:** the user's ability to acquire information while traveling
7. **Presence:** the user's sense of immersion within the environment
8. **User comfort:** the lack of simulator sickness, dizziness, or discomfort while traveling

Some of these metrics are simple to understand and measure quantitatively. Others, like ease of use, are generally based on qualitative, subjective ratings. The VE community has attempted to standardize questionnaires for important subjective metrics such as presence and comfort. It is tempting to dismiss qualitative metrics, but these often have the same or greater importance than measures like speed and accuracy. For example, if an interaction technique causes moderate discomfort in 25 percent of users, this subjective measure will outweigh any other performance advantages the technique may have.

2.4 Outside Factors Affecting Performance

Finally, our framework includes sets of outside factors besides travel technique that might affect one or more of the performance metrics. By explicitly modeling these factors, we can separate their effects from the performance effect due to the travel technique.

We have partitioned these factors into four groups.

1. **Task characteristics** are elements of the user task that can change user performance, such as the distance to be traveled, the number of turns in the path, or the visibility of the target.
2. **Environment characteristics** can also be important, including the presence of obstacles, the size of the environment, or the amount of structure (as opposed to randomness) within the space.
3. **User characteristics** are performance-affecting factors such as the user's level of visual acuity, his spatial ability, or his technical background.
4. **System characteristics** include the frame rate, rendering style, or the use of a virtual body representation.

For more-complete lists of outside factors, see Bowman et al. (1998).

2.5 Previous Related Experiments

We have previously completed four experiments comparing various travel techniques within the context

of this formal evaluation framework. Two of them are particularly relevant to the discussion of our spatial-orientation experiment below.

Our first spatial awareness experiment (see Bowman et al. (1997) for a detailed description) measured the user's knowledge of the position of her viewpoint and other objects within a virtual environment, after travel using certain velocity/acceleration schemes. The environment was simple: the user moved within a large black cube with white grid lines, and this space contained four, small, colored cubes. After pressing a button, the user would be moved from the current location to a new location using one of four velocity/acceleration techniques. At the end of travel, a colored stimulus was presented, and the user was required to find the cube of the same color. The results were clear: there were no significant differences in spatial awareness due to different levels of velocity or acceleration, except in the case where velocity was infinite (the user was instantly "teleported" to the target location). This "jumping" technique caused significantly more disorientation in users.

A second experiment focused on the user's ability to gather information while traveling through a VE. (See Bowman et al. (1998) for a detailed description.) We compared three steering techniques (methods for selecting the direction of motion), which used the user's gaze direction, pointing direction, or torso direction. We simultaneously varied the dimensionality of the environment (travel was in one, two, or three dimensions). Subjects traveled through corridors in which they saw signs containing simple words. At the end of travel they indicated their level of information gathering by marking words and locations on a paper map of the corridor.

We found that the dimensionality of the corridor was highly significant, with three-dimensional corridors (those moving along all three principal axes, including the vertical axis) producing the lowest scores. However, we found no significant differences between travel techniques. A possible explanation for this is that the techniques that afford information-gathering (for example, one can look in any direction while using the pointing technique) also impose the most cognitive load on the user, which may displace or interfere with some of the information. Thus, in the current experiment, we




Figure 2. Inside view of a corridor with a target object.

searched for techniques that would encourage spatial awareness but also reduce cognitive load on the user.

3 Experiment

The goal of this evaluation was to compare the spatial orientation of users after traveling through an immersive virtual environment (using an HMD) using certain travel techniques. As noted in the introduction, Chance et al. (1998) found that using a physical translation technique produced better spatial orientation, although the absolute error measurements they report were still relatively high. As they did in their study, 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 2). At the end of each corridor, the subject was "virtually blindfolded" (the corridor and objects disappeared from view), and asked to point (in three-dimensional space) in the direction of one of these objects. The response variable was the angular error, in degrees, for this pointing task.

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.

Like the previous experiment, this evaluation was done in the context of the framework discussed above. We compared three types of travel techniques from our taxonomy for their performance on a spatial-orientation metric. We also explicitly included outside factors that might affect performance.

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. In all cases, the translation of the user's viewpoint was virtual, but the user's gaze direction was controlled by physical head rotations.

These three techniques represent different levels of user versus system control of motion, and we hoped to discover which metaphor produces the highest level of spatial orientation. The extreme techniques are analogous to our 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 path. Two-dimensional corridors replicate the experience of moving through building hallways, while three-dimensional corridors also require ascending and descending. (Note that dimensionality refers to the user's motion path only; all environments in the experiment

were three-dimensional models.) Second, some corridors had only 90 deg. (right-angle) turns, while others turned at arbitrary angles. See Figure 3 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.¹

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 before taking a standardized spatial ability test. This cube-comparison test measures 3-D 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 (such as chairs, tables, and lamps) that would be used as stimuli during the experiment. This served the dual purpose of acquainting users with the 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

1. 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 3. Views of four corridor types used in the spatial orientation experiment. Top left: 2-D, right angles, top right: 2-D, non-right angles, bottom left: 3-D, right angles, bottom right: 3-D, non-right angles.

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 (Figure 4) 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 nondominant hand where the 3-D 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 three target objects and a "home" object that users looked at to begin a trial). Us-

ers 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 to 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 3-D space surrounding the user. Subjects completed five trials of each type, for a total of twenty trials.

The main experiment compared the three travel techniques, the two corridor dimensionalities (2-D and 3-D), the two angle conditions (right angles versus non-right angles), and the two map conditions (present or absent). There are only twenty valid combinations of these variables, as the route-planning technique uses a 3-D 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. The purpose of practice was to ensure that differences in performance were not due to VE experience or expertise using the travel techniques. All subjects felt comfortable with the techniques after either one or two traversals of the practice corridor for each technique. 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 sixteen (four corridors for each combination of the dimension and angle variables). Three objects were placed in each corridor at one of several predefined 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. Corridors were of variable length.

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. This technique allows the user to look in one direction while moving in another.

In the system-automated technique, the user simply pressed the button to begin moving down the defined path through the corridor. Although the system controlled the translation of the user's viewpoint, the user was still free to look in any direction while moving.

For the route-planning technique, subjects used the stylus to place markers on the 3-D map of the corridor to define a path. (Figure 4 shows a user in the process of defining a path.) When the user was satisfied with the




Figure 4. Route-planning technique using virtual map and stylus.

markers, he clicked the stylus button within a green button on the lower-left corner of the map. 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, regardless of whether or not this was along a direct line of sight (the path might travel through the corridor walls). Again, the user could look in any direction while being moved along this path by the system.

While using any of the techniques, users could click the button while moving to stop. No collision detection was provided, so subjects could travel through corridor walls. On two-dimensional corridors, users were constrained to a constant height above the floor, but were allowed to move anywhere in virtual space while in three-dimensional corridors.

At the end of each corridor, subjects were presented with a recorded voice stimulus (presented via headphones) instructing them to point in the direction of one of the objects seen in that corridor (for example, "Please point to the chair"). Users estimated the direction by pointing the stylus and recorded their answer by pressing the button. We measured the angular error between the 3-D direction pointed and the actual direction to the object. We assumed that all objects could be remembered equally well and did not analyze the differences in angular error results for specific objects, al-

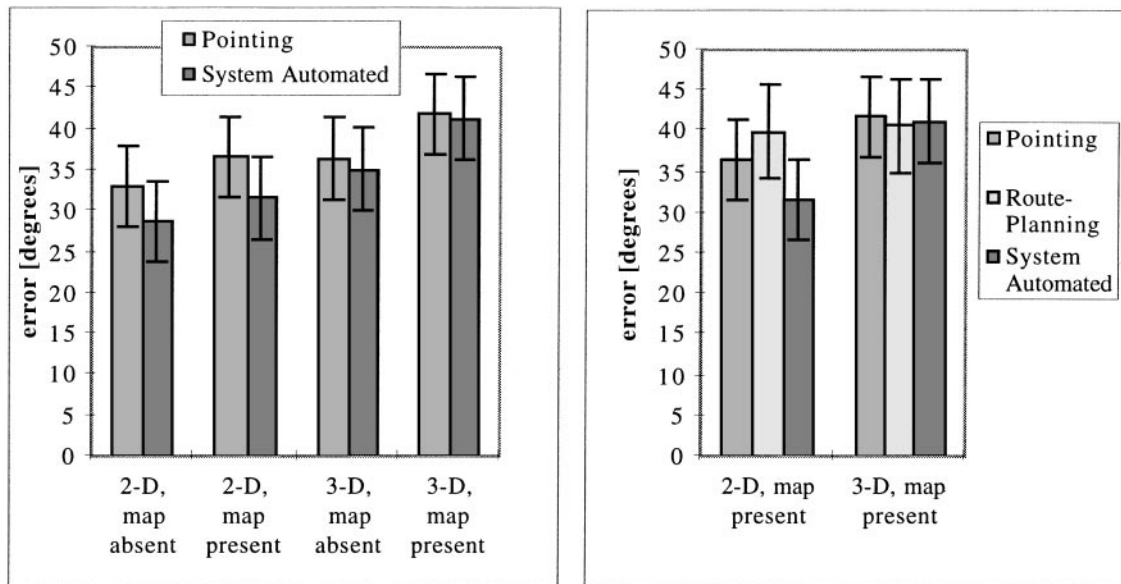


Figure 5. Mean error in various treatment combinations for Analysis 1 (left) and Analysis 2 (right).

though this would be interesting future work. Secondary response variables included time spent in each corridor (speed is not the primary focus of this experiment, but, if interaction techniques exhibit equal spatial-orientation performance, the faster technique might be preferable), 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, although exploratory, proved quite interesting and is discussed in detail below.

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 angle conditions, and two map conditions, for a total of sixteen treatment combinations. Analysis 2 considered all three techniques, two dimensionalities, and two angle conditions, with the 3-D map always present, for a total

of twelve combinations. The map variable could not be used in analysis 2 because 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 5. Analysis 1 showed a significant main effect for dimension (mean 2-D error: 32.47, mean 3-D 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 angle conditions were found. Analysis 2 showed a marginally significant main effect for dimension (mean 2-D error: 36.012, mean 3-D 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 significantly longer amounts of time while using the pointing technique, when in 3-D corridors, and when in corridors

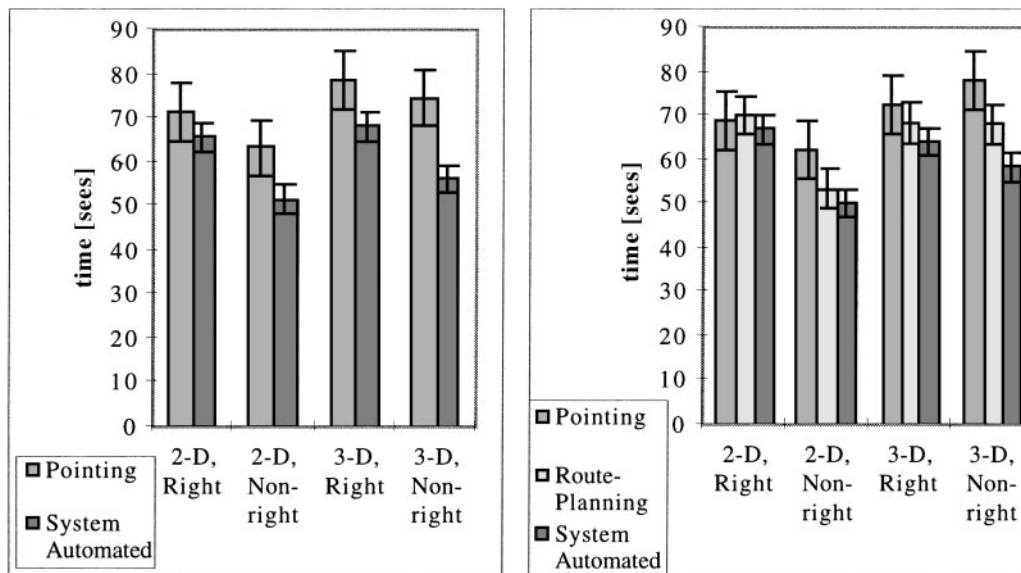


Figure 6. Mean time in various treatment combinations for analysis 1 (left), and analysis 2 (right).

with right angles only. The result for the angle variable is easily explained due to the fact that the corridors with right angles only were on average 20 m longer than non-right angle corridors. When times are normalized by corridor distance, time in right-angle corridors is actually slightly lower. Two-dimensional corridors were also longer on average than 3-D corridors, strengthening the result that more time is spent in 3-D corridors. Results for the time metric are summarized in Figure 6.

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 and 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 (since users always have control over their gaze direction via head-tracking), but it 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 it 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.
- **3-D overview:** Users fly above the corridor to obtain a single view of the complete corridor and the objects within it, which might encourage a survey

Table 1. Number of Subjects Observed Using Common Strategies for Each Travel Technique

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

representation. This strategy is available only on 3-D corridors using the pointing or route-planning techniques.

Table 1 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 0 and 3 indicating the number of technique-specific strategies used (3-D overview, moving through walls, path retracing). For the system-automated technique, the level was either 0 or 1 indicating whether or not the subject used the backing-in strategy. The route-planning strategy level ranged from 0 to 3 indicating the use of the 3-D overview, moving through walls, and/or backing in strategies.

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 0 had better scores with the system-automated technique than with the pointing technique, and did equally well on 2-D and 3-D corridors. Subjects with a pointing-technique strat-

egy level of one or two had approximately equal scores using all three travel techniques. Subjects with a pointing-technique strategy level of 3 had better scores using the pointing technique than other techniques, and performed better on 3-D corridors than 2-D. 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 deg. of error when objects were visible versus 26.7 deg. 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 deg. while trials in which objects could appear anywhere had an

average error of 24.9 deg. 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.

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 had a significant effect (that subjects performed better on 2-D corridors than 3-D). 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 3-D corridors are not familiar to subjects, but 2-D 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 deg. for all conditions, which is lower than the mean error for the best technique (physical motion) reported by 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, although 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 distinctive 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. None of these explanations is completely satisfactory, and further research is needed to determine the effects of map usage on spatial orientation in an immersive environment.

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 higher performance on the angular-error metric, perhaps indicating that the pointing technique is more complex and requires more effort on the part of the user to maintain spatial orientation. This may also show that subjects were using more-sophisticated strategies to maintain spatial orientation. (See below.) Three-dimensional corridors were also shown to require more user time, again demonstrating the difficulty of understanding such corridors. Even with the extra time spent in these corridors, the error was still significantly higher than in 2-D corridors. Finally, we found that subjects spent slightly less time in corridors containing only right angles than in those with non-right angles when the effect of corridor length is normalized.

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 3-D 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 3-D corridors than 2-D.

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 3-D corridors for these sophisticated users is explained by the fact that subjects were constrained to a constant height above the floor on 2-D corridors, and therefore could not use the powerful 3-D 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 tech-

niques 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 in which 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 deg.) than mean error values for the route-planning (21.76 deg.) and system-automated (25.71 deg.) techniques, although these subjects did extremely well using all three techniques.

These results provide two important lessons for developers of interaction techniques for VE applications. First, the techniques must provide opportunities for high performance on the 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 in which 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 Conclusions and Future Work

In this paper, we have reported on the results of an experiment comparing the effects of various virtual travel techniques on the spatial orientation of users. The experiment showed that techniques using virtual translation along with physical rotation could produce reasonable maintenance of spatial orientation as measured by an object-pointing task. Pointing errors were significantly affected by the complexity of the path through the environment. Interestingly, the strategies used by

subjects to aid their spatial orientation had a measurable effect on performance.

This experiment was designed and developed in the context of our formal design and evaluation framework for VE travel techniques, and it validates the usefulness of such a methodology. In the future, we plan more experiments comparing technique performance on various tasks. We will also use the taxonomies for the design of new techniques based on the results of such experiments.

For example, this experiment indicated that, if the proper strategies are used, a technique giving the user more control produces better spatial orientation. However, the performance of the pointing technique is hampered because of the cognitive load it induces as well as the failure of some users to recognize important strategies. Therefore, we will be looking to design new techniques that retain the aspect of user control, but which induce lower levels of cognitive load and which have clear affordances (Norman, 1990) indicating the most efficient strategies for their use.

Acknowledgments

The authors would like to thank the subjects who participated in our experiment. They also thank the members of the GVU Center's Virtual Environments group for their feedback on early versions of the software and the members of the 3DUI mailing list for their input on experimental design and hypotheses.

References

- Bowman, D., Koller, D., & Hodges, L. (1997). Travel in immersive virtual environments: An evaluation of viewpoint motion control techniques. *Proceedings of the Virtual Reality Annual International Symposium*, 45–52.
- . (1998). A methodology for the evaluation of travel techniques for immersive virtual environments. *Virtual Reality: Research, Development, and Applications*, 3(2), 120–131.
- Chance, S., Gaunet, F., Beall, A., & Loomis, J. (1998). Locomotion mode affects the updating of objects encountered during travel: The contribution of vestibular and proprioceptive inputs to path integration. *Presence: Teleoperators and Virtual Environments*, 7(2), 168–178.
- Darken, R., & Sibert, J. (1996a). Navigating in large virtual worlds. *The International Journal of Human-Computer Interaction*, 8(1), 49–72.
- . (1996b). Wayfinding strategies and behaviors in large virtual worlds. *Proceedings of CHI*, 142–149.
- Kessler, D., Kooper, R., & Hodges, L. (1998). *The simple virtual environment library: Version 2.0 user's guide*. (Graphics, Visualization, and Usability Center Technical Report GIT-GVU-98-13.)
- Norman, D. (1990). *The Design of Everyday Things*. New York: Doubleday.
- Slater, M., Usoh, M., & Steed, A. (1995). Taking steps: The influence of a walking metaphor on presence in virtual reality. *ACM Transactions on Computer-Human Interaction*, 2(3), 201–219.
- Waller, D., Hunt, E., & Knapp, D. (1998). The transfer of spatial knowledge in virtual environment training. *Presence: Teleoperators and Virtual Environments*, 7(2), 129–143.