

Travel in Immersive Virtual Environments: An Evaluation of Viewpoint Motion Control Techniques

Doug A. Bowman, David Koller, and Larry F. Hodges
Graphics, Visualization, and Usability Center
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280
{bowman,koller,hodges}@cc.gatech.edu

Abstract

We present a categorization of techniques for first-person motion control, or travel, through immersive virtual environments, as well as a framework for evaluating the quality of different techniques for specific virtual environment tasks. We conduct three quantitative experiments within this framework: a comparison of different techniques for moving directly to a target object varying in size and distance, a comparison of different techniques for moving relative to a reference object, and a comparison of different motion techniques and their resulting sense of "disorientation" in the user. Results indicate that "pointing" techniques are advantageous relative to "gaze-directed" steering techniques for a relative motion task, and that motion techniques which instantly teleport users to new locations are correlated with increased user disorientation.

1. Introduction

Virtual environment (VE) user interfaces have not been the focus of a great deal of user testing or quantitative analysis. *Travel*, by which we mean the control of user viewpoint motion through a VE, is an important and universal user interface task which needs to be better understood and implemented in order to maximize users' comfort and productivity in VE systems. We distinguish travel from navigation or wayfinding, which refer to the process of determining a path through an environment to reach a goal. Our work attempts to comprehend and categorize the techniques which have been proposed and implemented, and to demonstrate an experimental method which may be used to evaluate the effectiveness of travel techniques in a structured and logical way.

There are several restrictions we place on our consideration of VE travel techniques. First, we examine only immersive virtual environments, which use head tracking and head-mounted displays or spatially immersive displays (SIDs), and use 3D spatial input devices for

interaction. Secondly, we study only first-person travel techniques, or those in which the user's view is attached to the camera point in the VE (techniques have been proposed in which the user's view is temporarily detached from this position for a more global view of the environment [e.g. 11]). Also, we do not include techniques using physical user motion, such as treadmills or adapted bicycles. Finally, we consider only techniques which are predominantly under the control of the user, and not those in which travel is carried out automatically or aided significantly by the system.

The following sections of this paper review related research in the area of VE travel interaction, and present a taxonomy of travel techniques and a framework for their evaluation. Three relevant experiments illustrating this framework and their results are then described.

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 [5] that studying and understanding human navigation and motion control is of great importance for understanding how to build effective virtual environment travel interfaces [13,19]. Although we do not directly address the cognitive issues surrounding virtual environment navigation, this area has been the subject of some prior investigation and discussion [3,20].

Various metaphors for viewpoint motion and control in 3D environments have been proposed. Ware et al. [17,18] identify the "flying," "eyeball-in-hand," and "scene-in-hand" metaphors. A fourth metaphor, "ray casting," [6] has been suggested, which can be used to select targets for navigation. Others make use of a "World-in-Miniature" representation as a device for navigation and locomotion in immersive virtual environments [11,15].

Numerous implementations of non-immersive 3D travel techniques have been described. Strommen compares three different mouse-based interfaces for children to control point-of-view navigation [16]. Mackinlay et al.

describe a general method for rapid, controlled movement through a 3D environment [8].

Mine [10] offers an overview of motion specification interaction techniques. He and others [e.g. 12] 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 [2,9]. Physical motion techniques have also been studied, such as the effect of a physical walking technique on the sense of presence [14], and the use of a “lean-based” technique [4].

3. Evaluation framework

3.1 Taxonomy

After reducing the space of viewpoint movement control techniques that have been proposed for immersive VEs (by applying the restrictions described in the Introduction), we are able to categorize these techniques in an organized design space (similar to [1]). Figure 1 shows the high-level entries in our taxonomy. There are three components in a travel technique, each of which corresponds to a design decision that must be made by the implementor. *Direction/Target Selection* refers to the method by which the user “steers” the direction of travel, or selects the goal position of the movement. *Velocity/Acceleration Selection* methods allow the user/system to set speed and/or acceleration. Finally, *Input Conditions* are the ways in which the user or system specifies the beginning time, duration, and end time of the travel motion.

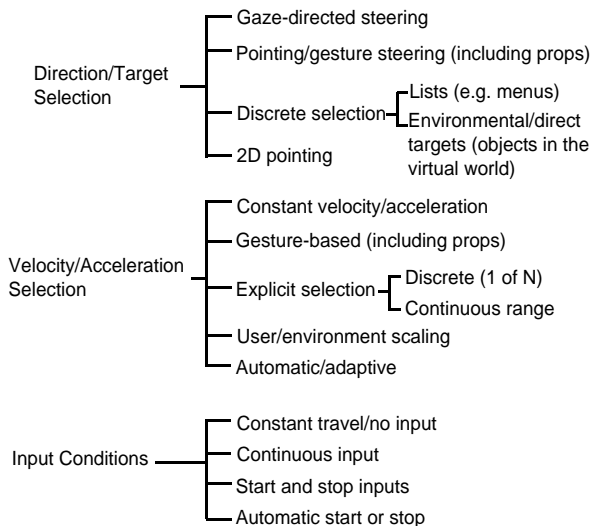


Figure 1. Taxonomy of virtual travel techniques

Note that some branches of the taxonomy may be combined to form new methods. For example, under velocity selection, a gesture-based technique may also be adaptive (the user’s gestures may cause different velocities in different system states). Also, some combinations of methods may not work together at all. In general, however, a travel technique is designed by choosing a method from each of these three branches of the taxonomy. For example, in one common technique the user holds a mouse button and moves with constant speed in the direction she is looking. In the taxonomy, this corresponds to gaze-directed direction selection, constant velocity, and continuous input conditions.

3.2 Quality factors

Explicit, direct mappings of the various travel techniques to suitable applications are not obvious, given that applications may have extremely different requirements for travel. Instead, we propose a list of quality factors which represent specific attributes of effectiveness for virtual travel techniques. These factors are not necessarily intended to be a complete list, and some of them may not be relevant to certain applications or tasks. Nonetheless, they are a starting point for comparing and measuring the utility of various travel techniques.

An effective travel technique promotes:

1. *Speed* (appropriate velocity)
2. *Accuracy* (proximity to the desired target)
3. *Spatial Awareness* (the user’s implicit 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)

The quality factors allow a level of indirection in mapping specific travel techniques to particular virtual environment applications. Our method involves experiments which map a travel technique to one or more quality factors, rather than to a specific application or task. Application developers can then specify what levels of each of the quality factors are important for their application, and choose a technique which comes closest to that specification.

For example, in an architectural walkthrough, high levels of spatial awareness, ease of use, and presence might be required, whereas high speeds might be unimportant. On the other hand, in an action game, one might want to maximize speed, accuracy, and ease of use, with little attention to information gathering. Because

applications have such diverse needs, we find it most efficient to relate experimental results first to specific quality factors and then allow designers to determine their own requirements and weighted importance for each quality factor.

4. Experiments

Even considering the aforementioned constraints on the techniques we are studying, our space of travel techniques is still large. It would be difficult to test every technique against every other technique for each quality factor. Therefore, we present three example experiments to produce preliminary results and illustrate the experimental method which may be used for such evaluations. These experiments were chosen because of their relevance and relate to travel techniques which are being implemented in some contemporary immersive virtual environments. The first two tests compare two direction selection techniques for absolute motion (travel to an explicit target object) and relative motion (travel to a target located relative to a “reference” object). The third experiment measures the spatial awareness of a user after using a variety of velocity/acceleration techniques.

In each of these experiments, the subjects were undergraduate and graduate students, with immersive VE experience ranging from none to extensive. A Virtual Research VR4 head-mounted display, Polhemus Isotrak trackers, and a custom-built 3-button 3D mouse were used. The test applications were run on an SGI Crimson workstation with RealityEngine graphics, and frame rates were held constant at 30 frames per second. Times were measured to within 0.001 second accuracy.

4.1 Comparing steering techniques

Perhaps the most basic of the quality factors listed above are speed and accuracy. These are simple to measure, generally important in most applications, and vary widely among different VE travel techniques. When a user wishes to move to a specific target location, it is not acceptable to move there slowly or inaccurately. Users can quickly become fatigued from holding input devices steady, pressing buttons, or looking in a certain direction for a lengthy period of time.

Clearly, the fastest and most accurate techniques will be those which allow the user to specify exactly the position to move to, and then automatically and immediately take the user to that location. For example, in our taxonomy, the direction/target selection technique might be discrete selection from a list or using direct targets (select an object to move to that object). Lists, however, require that the destinations be known in advance, while direct targets only allow movement to objects, not to arbitrary positions.

Therefore, a more general direction/target selection technique is needed that still maintains acceptable speed

and accuracy characteristics. Two of the most common techniques used in VE applications are gaze-directed steering and hand-directed steering (or “pointing”) [10]. In gaze-directed steering, the user’s view vector (typically the orientation of the head tracker) is used as the direction of motion, whereas the direction is obtained from the user’s hand orientation in the pointing technique. Our first set of experiments compares these two techniques in the absolute and relative motion tasks.

4.2 Absolute motion experiment

Our study of *absolute motion* compared these techniques for the task of traveling directly to an explicit target object in the environment. Subjects were immersed in a sparse virtual environment containing only a target sphere. A trial consisted of traveling from the start position to the interior of the sphere, and remaining inside it for 0.5 seconds. The radius of the sphere and the distance to the target were varied, and subjects’ time to reach the target was recorded.

Besides varying the travel technique between gaze-directed steering and pointing, we also studied another factor: constrained vs. unconstrained motion. In half of the trials, users could move about the environment with six degrees of freedom. In the constrained trials, however, the user was not allowed to move vertically (the target sphere appeared on the horizontal plane in all trials). Thus, there were four travel techniques tested in all.

We hypothesized that gaze-directed techniques and constrained techniques would produce lower times, because these techniques should be more accurate than pointing and unconstrained methods. It is clear that the 2D constraint should produce more accuracy, because there are fewer degrees of freedom to control. It may not be as obvious that gaze-directed steering should be more accurate than pointing, but consider two comparisons:

First, gaze-directed steering uses the muscles of the neck, while pointing uses the arm and wrist muscles. The neck muscles seem more stable than the arm or wrist muscles; therefore one can hold the head in a fixed position easier than the arm or hand. Second, with gaze-directed steering, there is a more direct feedback loop between the sensory device (the eyes) and the steering device (the head). The user looks in a direction and sees travel in that direction. With pointing, the user may look in one direction and travel in another. More interpretation of the visual input must occur to pick the correct direction, and the hand must be made to point in that direction.

Subjects performed 80 trials with each of the four techniques. There were four values of the sphere radius (0.4, 0.8, 1.5, and 2.5 meters) and four target distances (10, 20, 50, and 100 meters); subjects thus performed 5 trials with each of these 16 combinations within a technique block. The travel velocity was kept constant, and a mouse button was used to effect travel (using a continuous input technique). Eight subjects participated,

and there were four different orderings for the travel techniques used, so that the effect of ordering was counterbalanced.

The time required for the subject to satisfy the goal condition was measured for each trial, and the results were analyzed using a standard 3-factor analysis of variance (ANOVA). The travel technique was shown to be non-significant for the experimental conditions, while target distance and target size were significant ($p < 0.01$). These results were somewhat surprising, since we hypothesized that gaze-directed steering and 2D constraints would produce lower response times due to greater accuracy. Figure 2 compares the times obtained by the four techniques at different distances, while figure 3 plots time against the target radius.

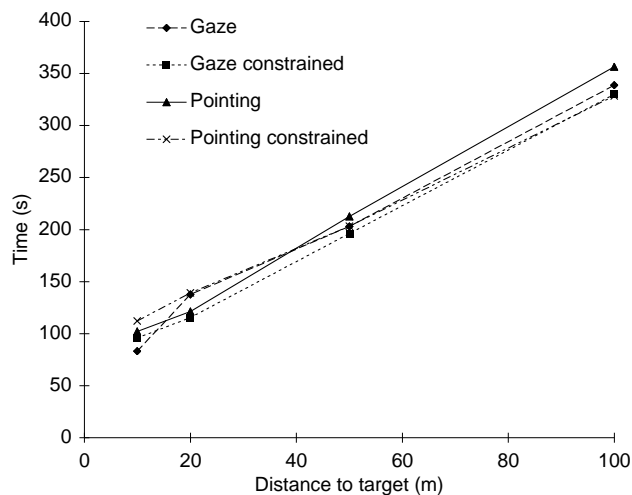


Figure 2. Absolute motion results for various target distances

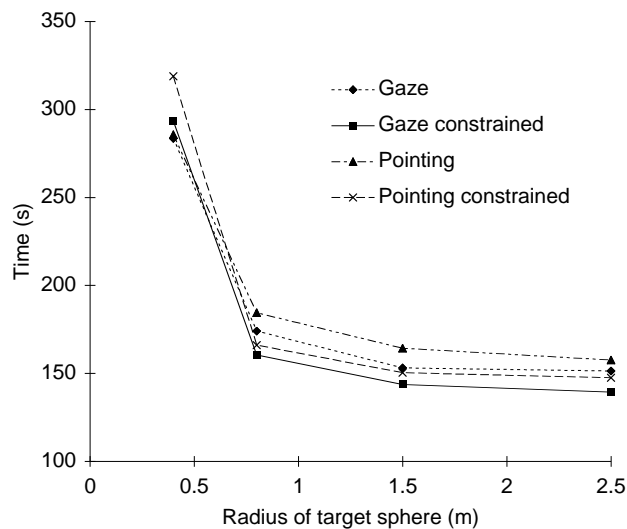


Figure 3. Absolute motion results for various target sizes

One possible reason for the lack of a statistically significant difference between gaze-directed techniques and pointing techniques in this experiment is that many subjects emulated gaze-directed steering during the pointing trials. That is, they both gazed and pointed in the desired direction, so that their head motions were mimicked by their hand motions. Also, because the desired trajectory in the experimental trials was always a straight line, with no obstacles, it was fairly easy for subjects to quickly find the right direction and lock their hand position. More significant differences between the techniques might be found with a more complex steering task.

Overall, this experiment suggested that both gaze-directed steering and pointing could produce accuracy in an absolute motion scenario. With the advantages of pointing that we will show in the second experiment of this set, we have strong evidence that it is a useful, general technique for direction/target selection when speed and accuracy are important.

The use of 2D constraints did not show a statistically significant performance gain in this experiment, but we still believe constrained motion to be an important technique for many applications where users do not need the extra freedom of motion. It allows users to be more lazy in their direction specification, so that more attention can be paid to the other tasks or features of the virtual environment. Although this reduced cognitive loading was not a factor in this experiment due to the sparseness of the environment and simplicity of the task, it would prove interesting to study performance of constrained vs. unconstrained motion in a dense virtual environment, perhaps with the addition of distractor tasks.

4.3 Relative motion experiment

In the second of this set of experiments, we again contrasted gaze-directed steering with pointing. Subjects were asked to travel from the starting position to a point in space a given distance and direction away from a reference object in the environment. This task was designed to measure the effectiveness of the techniques for traveling relative to another object in the environment.

This task is actually frequently used in such applications as architectural walkthrough. For example, suppose the user wishes to obtain a head-on view of a bookshelf which fills her field of view. There is no object to explicitly indicate the user's destination; rather, the user is moving relative to the bookshelf.

The environment for this experiment again consisted of a single object, in this case a three-dimensional pointer (see figure 4). This pointer defined a line in space, and the subject's goal was to travel to a position on that line which is a reference distance away from the pointer. In order to help the user learn this distance, which was constant for each trial, there were five initial practice trials at the beginning of each set in which a sphere was placed at the target position (as in the figure). During normal

trials, the sphere was not visible. The trial ended when the subject had reached the target point, within a small radius. After each trial, the pointer moved to a new position and orientation in space for the succeeding trial.

The capability of traveling in reverse was added as a second factor in this experiment. By pressing a mouse button, the user toggled between forward mode and reverse mode. In reverse mode, the user traveled in the opposite direction (the direction obtained by negating each value in the direction vector) from the one specified by the head or hand position. Each trial began in forward mode, and subjects were free to use reverse mode as often or as little as they liked. In total, then, we tested four techniques: gaze-directed steering with and without reversal capability, and pointing with and without reversal capability.

Nine subjects participated in the experiment. Each subject completed four blocks of trials. Within each block, there were four sets, corresponding to the four travel techniques, and each set consisted of 20 trials. The sets were ordered differently within each block for counterbalancing purposes. Since we anticipated a significant learning effect for this difficult task, only the last 5 trials were counted toward the overall time. Travel time was measured from the moment the subject initiated motion to the moment when the task was completed. For each trial, the distance from the starting position to the target was either 5, 10, 15, or 20 meters. As in the absolute motion experiment, constant velocity and continuous input conditions were used. Median travel times collected in the experiment are shown in table 1.

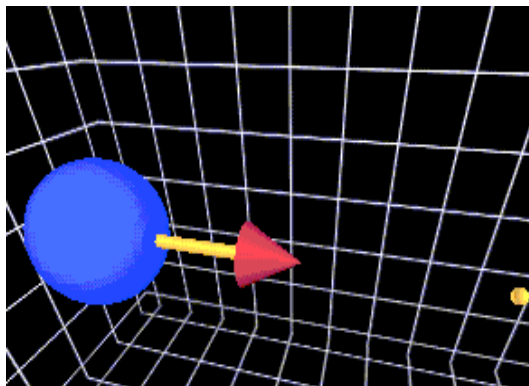


Figure 4. Relative motion environment

A standard single-factor ANOVA was performed on the median times of each of the subjects to analyze the results of this experiment. Median times were used here in order to minimize the effect of very short or very long times. Short trials could occur if the subject simply “got lucky” in hitting the target, and long trials occurred when the subject made several passes at the target, missing it by a little each time. Since we were interested in the normative case, we did not wish these very small or large times to have a large influence on the dependent measure.

The analysis showed that the travel technique used did indeed have a significant effect on time ($p < 0.025$), and further analysis of the individual means (using Duncan’s test for comparison of means) revealed that both pointing techniques were significantly faster than each of the gaze-directed techniques ($p < 0.05$). There were no significant differences between gaze-directed steering and gaze-directed steering with reversal, or between pointing and pointing with reversal.

	Without reverse	With reverse
Gaze-directed	12.36	12.15
Pointing	9.60	9.75

Table 1. Relative motion experiment median times by technique (in seconds)

The reason that pointing techniques were superior for this task is clear both theoretically, and from observation. In order to move relative to an object, especially in this sparse environment, the subject needs to look at the object while traveling. Therefore, except in the case where the subject is already on the line connecting the target and the object, gaze-directed steering requires this cycle of actions:

1. Look at the reference object
2. Determine direction toward target
3. Look in this direction
4. Move in this direction for an estimated amount of time
5. If the target has not been reached, repeat

On the other hand, with pointing techniques, one can look at the object while travel is taking place, making directional corrections “on the fly.” Most subjects discovered this right away, and would often point off to the side while gazing straight ahead at the object.

Gaze-directed steering becomes especially painful when the subject gets too close to the object, because then each check of the object requires that the head be turned 180 degrees as the user travels out along the reference line.

This situation shows the utility of the reversal capability. Subjects often complained about the physical difficulty of the gaze-directed technique, since it required so much head motion, but they did not complain when the reversal capability was added. However, the directional accuracy of most subjects suffered greatly when in reverse mode. Reverse mode requires users to turn the head or hand to the left in order to back up to the right; the fact that the virtual environment allows travel in three dimensions adds to the complexity. A few users became expert at this, but overall it did not improve times over simple gaze-directed steering.

In the same way, the addition of the reversal capability to pointing added cognitive load and complexity to the technique. It is somewhat useful (less useful than with gaze-directed steering, though), since going backwards

with simple pointing requires that the arm be pointed straight back or that the wrist be turned completely around, both of which are physically difficult. The gain in ease of use, however, is not significant.

This experiment highlights the advantages that pointing techniques have over gaze-directed steering; pointing is clearly superior for relative motion. Since pointing and gaze-directed steering showed no significant difference in the absolute motion task, we would recommend pointing as a direction/target selection technique for almost all general purpose applications which require speed and accuracy. This is not to say that gaze-directed steering should never be used. It has significant advantages in its ease of use and learning, and its direct coupling of the steering mechanism and the user view. Table 2 outlines some of the major advantages and disadvantages of the two techniques that we have seen both in controlled experiments and observation of VE application users.

Gaze-Directed Steering	
Advantages	Disadvantages
<ul style="list-style-type: none"> steering and view are coupled ease of use/learning easier to travel in a straight line slightly more accurate 	<ul style="list-style-type: none"> requires much head motion less comfortable can't look at object & move another direction
Pointing	
Advantages	Disadvantages
<ul style="list-style-type: none"> user's head can stay relatively still more comfortable can look and move in different directions 	<ul style="list-style-type: none"> can lead to overcorrection more cognitive load harder to learn for most users slightly less accurate

Table 2. Comparison of two direction selection techniques

4.4 Directional disorientation due to velocity and acceleration

Our final experiment deals with another of the quality factors, *spatial awareness*. For travel, we define this term to mean the ability of the user to retain an awareness of her surroundings during and after travel. The opposite of spatial awareness would be *disorientation* due to travel. Users may become disoriented because of improper motion cues, lack of control over travel, or exposure to large velocities or accelerations.

For this experiment, we focused on the second branch of our taxonomy, velocity/acceleration selection. We investigated the effect of various velocity and acceleration techniques on the spatial awareness of users. Specifically, we were interested in infinite velocity techniques, which we will refer to as “jumping,” since the user jumps from

one position in the virtual environment to another. Our previous experience with VE applications had led us to believe that such techniques could be quite disorienting to the user. Jumping techniques are often paired with a discrete target selection technique, such as when the user picks a location from a list or selects an object in the environment to which he wishes to travel.

To test the user's spatial awareness, we created a simple environment consisting of several cubes of contrasting colors (see figure 5). The subject was instructed to form a “mental map” of the environment from the starting position, and to reinforce that map as the experimental session continued. For each trial, the user was taken to a new location via a straight-line path using one of the velocity/acceleration techniques. Upon arrival, a colored stimulus (seen in the corner of figure 5) corresponding to one of the cubes was presented to the user. The user located this cube in the environment, and pressed either the left or right button on a mouse, depending upon whether an “L” or “R” was displayed on the cube.

By measuring the amount of time it took the user to find the cube and make this simple choice, we obtained data on how well the user understood the surrounding environment after travel. In other words, were they still spatially aware after travel, or were they disoriented? If complete disorientation had taken place, the time to complete the task should be about the same as a random visual search. On the other hand, if the subject were still spatially aware, the response time should be much lower.

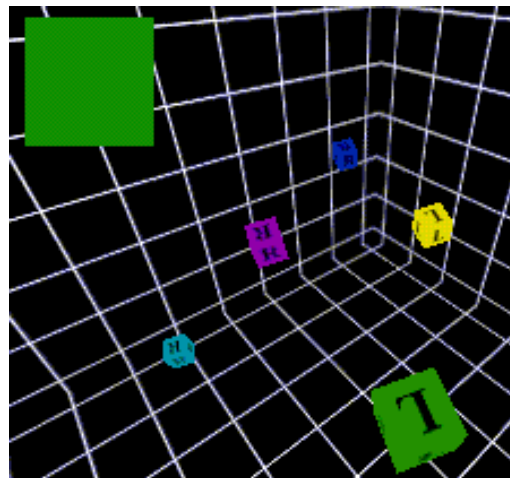


Figure 5. Spatial awareness environment

We tested four different velocity/acceleration techniques in this experiment. Two constant velocity techniques were used, with the fast velocity ten times greater than the slow velocity. A third technique was infinite velocity, where the user is taken directly to the destination. Finally, we implemented a “slow-in, slow-out” (SISO) technique (similar to [8]) in which the user begins slowly, accelerates to a maximum speed, then decelerates as the

destination is reached. This technique was implemented in such a way that the time to travel to the destination was always equal to the time it would take to travel the same path using the fast constant velocity technique.

Ten subjects participated in the experiment. Each subject completed four blocks of trials, and there were four sets of trials (one for each technique) within each block. Each set consisted of 20 trials, the first 10 of which were considered practice trials. These practice trials allowed the subjects to learn the task, and also gave them a chance to build an accurate mental map of the environment by viewing it from many different locations (the positions of the cubes in the environment were different for each set of trials). Within each block, the order of the techniques was different to eliminate any effect of ordering.

To analyze the results, we performed a standard single-factor ANOVA on the average times of the subjects. We found that the differences in time for the various velocity and acceleration techniques was significant ($p < 0.01$). Further analysis on the individual means, using Duncan's test with $p < 0.05$, showed that the times for the infinite velocity (jumping) technique were significantly greater than times for each of the other techniques. There were no other significant differences, however. Table 3 presents the average times for each technique by subject. For 7 of 9 subjects, the largest time was for the jumping condition.

	Slow	Fast	SISO	Jumping
Subj. 1	3.13	4.24	6.09	5.82
Subj. 2	2.01	2.83	3.25	4.88
Subj. 3	2.38	2.59	2.69	3.63
Subj. 4	2.94	2.71	2.48	4.31
Subj. 5	3.56	2.60	3.02	3.97
Subj. 6	3.28	2.67	2.90	3.23
Subj. 7	3.44	4.39	4.84	4.97
Subj. 8	2.75	3.73	3.27	5.19
Subj. 9	2.71	2.32	2.91	3.15
Average	2.91	3.12	3.49	4.35

Table 3. Spatial awareness experiment average times by subject and technique (in seconds)

These results support our main hypothesis: that jumping techniques can reduce the user's spatial awareness. We frequently observed subjects perform a visual search of the entire space for the target when using the jumping technique, even though they supposedly had all the information they needed to find the target. That is, they knew the starting position, the time of travel and the direction they were facing (travel did not change the viewer's orientation). However, they were unable to process this information accurately enough to know the target direction.

Our observations suggest that the problem lies in the lack of continuity of travel. With jumping techniques,

there is no sensation of motion, only that the world has somehow changed around the user. It is a technique whose motion has no analog in the physical world. Of course, if the speed required to reach the target is the only consideration, infinite velocity techniques are optimal. However, they sacrifice the spatial awareness of a user, and our observations lead us to believe that these techniques reduce the sense of presence as well.

We were surprised that there were no significant differences between other pairs of techniques. We had expected that the slow constant velocity would produce the least disorientation (it did have the lowest time, but the differences were not significant), and hypothesized that our slow-in, slow-out technique would be less disorienting than the fast constant velocity.

The problem with slow-in, slow-out may have been in our implementation. In order to ensure that this technique would produce the same travel times as the fast constant velocity technique, it was necessary that the acceleration function change dynamically for each trial under slow-in, slow-out. It is possible that users were simply not able to build an accurate mental model of their velocity and acceleration, meaning that they would not know how far they had traveled for a given trial. We noted that subjects generally turned in the general direction of the target, but were not sure of its exact location.

These results may be taken as encouraging to the designers of VE travel techniques, in that they suggest that the amount of user disorientation may not be significantly affected by the velocity/acceleration technique, at least up to a relatively high velocity. We would like to perform a follow-up experiment in which we attempt to find the velocity at which user disorientation becomes a significant factor in user spatial awareness.

5. Conclusions and future work

These experiments only scratch the surface in investigating the design space of travel techniques for virtual environments. However, we believe that we have isolated some important results in this area with our current work. Our first set of two experiments showed that pointing techniques are faster than gaze-directed steering techniques for the common relative motion task, and that the two techniques perform equally for absolute motion. In an application needing a general technique with speed and accuracy, therefore, pointing is a good choice. It requires more time to become expert, however, so if the application will be used only rarely or a single time by a user, a more cognitively simple technique may be called for. The spatial awareness experiment showed that infinite velocity techniques can significantly increase user disorientation and may lead to reduced presence.

Also, we have presented an experimental methodology and framework that can be a common ground for discussion and further testing in this area. A more completely developed taxonomy which is orthogonal and

comprehensive is desired. Particular VE travel techniques in this taxonomy may then be mapped to levels of the quality factors experimentally, in the manner described. Application designers may then specify the weight given to each of the quality factors for their specific needs and goals and choose techniques accordingly.

In addition to the follow-up experiments discussed above, we would like to create a more general testbed for VE travel techniques. Our plans call for creation of a test environment similar to the Virtual Environment Performance Assessment Battery (VEPAB) [7]. This environment would be instrumented to collect data on any or all of the quality factors we discussed. Specific travel techniques would then be used in these environments and assigned an overall score for each of the quality factors. Such a system would provide an objective measure for a travel technique that could be compared to the scores from other techniques under consideration for an application.

Acknowledgments

The authors would like to thank Neff Walker, Ben Watson, and Drew Kessler for their help and advice, and the experimental subjects who volunteered their time. This work was supported in part by the National Science Foundation.

References

- [1] S. Card, J. Mackinlay, and G. Robertson. The Design Space of Input Devices. *Proceedings of CHI* (Seattle, WA, 1990), pp. 117-124.
- [2] J. Chung. A Comparison of Head-tracked and Non-head-tracked Steering Modes in the Targeting of Radiotherapy Treatment Beams. *Proceedings of 1992 Symposium on Interactive 3D Graphics* (Cambridge, MA, 1992), pp. 193-196.
- [3] R. Darken and J. Sibert. A Toolset for Navigation in Virtual Environments. *Proceedings of ACM Symposium on User Interface Software and Technology* (Atlanta, GA, 1993), pp. 157-165.
- [4] K. Fairchild, L. Hai, J. Loo, N. Hern, and L. Serra. The Heaven and Earth Virtual Reality: Designing Applications for Novice Users. *Proceedings of IEEE Symposium on Research Frontiers in Virtual Reality* (Seattle, WA, 1993), pp. 47-53.
- [5] K. Herndon, A. van Dam, and M. Gleicher. The Challenges of 3D Interaction. *SIGCHI Bulletin*, vol. 26, no. 4, pp. 36-43, 1994.
- [6] K. Hinckley, R. Pausch, J. Goble, and N. Kassell. A Survey of Design Issues in Spatial Input. *Proceedings of ACM Symposium on User Interface Software and Technology* (Marina del Rey, CA, 1994), pp. 213-222.
- [7] D. Lampton, B. Knerr, S. Goldberg, J. Bliss, J. Moshell, and B. Blau. The Virtual Environment Performance Assessment Battery (VEPAB): Development and Evaluation. *Presence*, vol. 3, no. 2, pp. 145-157, 1994.
- [8] J. Mackinlay, S. Card, and G. Robertson. Rapid Controlled Movement Through a Virtual 3D Workspace. *Proceedings of SIGGRAPH* (Dallas, TX, 1990), in *Computer Graphics*, vol. 24, no. 4, pp. 171-176.
- [9] P. Mercurio, T. Erickson, D. Diaper, D. Gilmore, G. Cockton, and B. Shackel. Interactive Scientific Visualization: an Assessment of a Virtual Reality System. *Proceedings of INTERACT* (Cambridge, UK, 1990), pp. 741-745.
- [10] M. Mine. Virtual Environment Interaction Techniques. UNC Chapel Hill Computer Science Technical Report TR95-018, 1995.
- [11] R. Pausch, T. Burnette, D. Brockway, and M. Weiblen. Navigation and Locomotion in Virtual Worlds via Flight into Hand-Held Miniatures. *Proceedings of SIGGRAPH* (Los Angeles, CA, 1995), pp. 399-400.
- [12] W. Robinett and R. Holloway. Implementation of Flying, Scaling, and Grabbing in Virtual Worlds. *Proceedings of Symposium on Interactive 3D Graphics* (Cambridge, MA, 1992), pp. 189-192.
- [13] E.R. Schieser. Principles of Navigation. In *Handbook of Modern Electronics and Electrical Engineering*, C. Belove (Ed.), Wiley, 1986.
- [14] M. Slater, M. Usoh, and A. Steed. Taking Steps: The Influence of a Walking Technique on Presence in Virtual Reality. *ACM Transactions on Computer-Human Interaction*, Vol. 2, No. 3, pp. 201-219, September 1995.
- [15] R. Stoakley, M. Conway, and R. Pausch. Virtual Reality on a WIM: Interactive Worlds in Miniature. *Proceedings of CHI* (Denver, CO, 1995), pp. 265-272.
- [16] E. Strommen. Children's Use of Mouse-Based Interfaces to Control Virtual Travel. *Proceedings of CHI* (Boston, MA, 1994), pp. 405-410.
- [17] C. Ware and D. Jessome. Using the Bat: a Six-Dimensional Mouse for Object Placement. *IEEE Computer Graphics and Applications*, vol. 8, no. 6, pp. 65-70, 1988.
- [18] C. Ware and S. Osborne. Exploration and Virtual Camera Control in Virtual Three Dimensional Environments. *Proceedings of Symposium on Interactive 3D Graphics* (Snowbird, UT, 1990), in *Computer Graphics*, vol. 24, no. 2, pp. 175-183.
- [19] R. Warren and A. Wertheim (Eds.). *Perception and Control of Self-motion*. Hillsdale, NJ: Erlbaum, 1990.
- [20] C. Wickens and P. Baker. Cognitive Issues in Virtual Reality. In *Virtual Environments and Advanced Interface Design*, W. Barfield and T. Furness (Eds.), Oxford University Press, 1995.