

Testbed Evaluation of Navigation and Text Display Techniques in an Information-Rich Virtual Environment

Jian Chen Pardha S. Pyla Doug A. Bowman

*Department of Computer Science and Center for Human-Computer Interaction
Virginia Tech
Blacksburg, Virginia, USA
{jichen8, ppyla, bowman}@vt.edu*

Abstract

Information-Rich Virtual Environment (IRVE) is a virtual environment (VE) that is enhanced with the addition of related abstract information. Such an environment conveys a rich information set that can make the VE more useful if provided with a useful interface and efficient interaction techniques. This motivates our current research goals to investigate the interface and interaction issues involved in IRVE. The fundamental question is how to access and display information in an effective way. This paper presents the first controlled experiment using a travel testbed based on our text layout taxonomy. We investigated two manipulation-based navigation techniques: Hand-Centered Object Manipulation Extending Ray-casting (HOMER) and Go-Go based navigation, and two text layout techniques: within-the-world display (WWD) and heads-up display (HUD). Four search tasks with repeated measures were performed to measure subjects' performance in a densely packed environment. We found that using Go-Go based navigation combined with HUD techniques is significantly better than the other three combinations for difficult naïve search tasks. HUD enabled better performance than WWD and Go-Go technique enabled better performance than HOMER technique for most of the tasks. Also, users preferred the combination of Go-Go and HUD for all tasks. Such results on their own, or combined with specific application characteristics, can provide valuable design guidelines for IRVEs.

1. Introduction and related work

An *Information-Rich Virtual Environment (IRVE)* is a virtual environment (VE) that is enhanced with the addition of abstract information, such as text, audio, and video [3, 9], for example, showing a tree coupled with related information in a virtual habitat environment (Figure 1). This contrasts with pure information visualization, which uses a VE to present abstract information alone, e.g., a 3D scatterplot of census data or other 3D work presented in [10].



Figure 1. Abstract Information in the Virtual Habitat

Previous efforts on IRVE applications, the ‘virtual venue’ and the ‘virtual habitat’ showed promise. The ‘Virtual Venue’ [6] is an IRVE with embedded text, audio, imagery, 3-D animation, and ‘experiential’ information within the perceptual space. Usability studies conducted on this environment proved that the experience is enjoyable, useful, and better than the printed-page or hypermedia based information. The “Virtual habitat” [7] was used to teach students environmental design issues by immersing them in an information-rich virtual zoo exhibit. These efforts show the potential of IRVEs for education. In our first proof-of-concept IRVE experiment [9], we received positive feedback on the scope of using IRVEs in general for education, training, etc.

Along these lines of research, our long-term goal is to formalize the design space and produce useful guidelines for the development of effective and usable IRVEs. Previously, we have presented a formal framework and research agenda to systematically research on IRVEs [9]. Currently, we are trying to answer the question: how to access and embed information in an IRVE in an effective manner? We feel the designing of interface and interaction with the information is more difficult because of the integration of perceptual and abstract data in one space. Within such an environment, we need a way to present the abstract information to ‘minimize’ the interface between the user and the information. We also

need to have effective navigation techniques for accessing the embedded information.

We used testbed evaluation [8] to tackle this problem because of its advantages of combining multiple tasks, multiple independent variables, and multiple response measures. It allows us to obtain a complete picture of the performance characteristics of the variables of a study. Previous works on applying testbed evaluation techniques include VEPAB project [14], VRMAT [17], and navigation and manipulation testbed [8].

This paper presents the first full scale experiment aimed at evaluating two manipulation-based navigation techniques and two text layout methods. The two navigation techniques are Hand-Centered Object Manipulation Extending Ray-casting (HOMER) and Go-Go based, and two text layout techniques: WWD and HUD.

Text, as one type of abstract information, has been used in many other applications in addition to the two just mentioned. AnthroGloss is a desktop VE that displays text information about human anatomy on the HUD. Landmarks were used to create a link between the perceptual information and text labels [19]. Mobile Augmented Reality Systems (MARS) testbed overlays computer generated text on real world objects to allow outdoor and indoor users to access and manage information [12]. However, the difference between augmented reality (AR) and IRVE is that IRVEs are purely synthetic, which gives IRVEs much more flexibility. The information can be easily animated, moved, scaled, and manipulated. Furthermore, navigation techniques are one of most essential interaction techniques between an IRVE and users, and they have been well studied in the past few years [8, 11, 15]. For the experiment, we also defined four search tasks that represent the tasks users are likely to perform in an IRVE, which are discussed in section 3.4.

The organization of this paper is as follows: we provide taxonomy of text layout techniques in section 2. In section 3, we describe the experiment design including the environment, hypothesis, tasks, and procedure. Results and analysis are presented in section 4. Impact is described in section 5. We conclude in section 6 with future work.

2. Text Layout Taxonomy

We created a taxonomy of text layout and display as a theoretical basis for our testbed implementation. Text display was classified based on *quantity* (amount of text embedded), *visual attributes* (qualitative aspects), *location* (spatial aspects), and *variation in time*. The high-level entries of our taxonomy are shown in Figure 2. We have not yet pursued the temporal classification in our

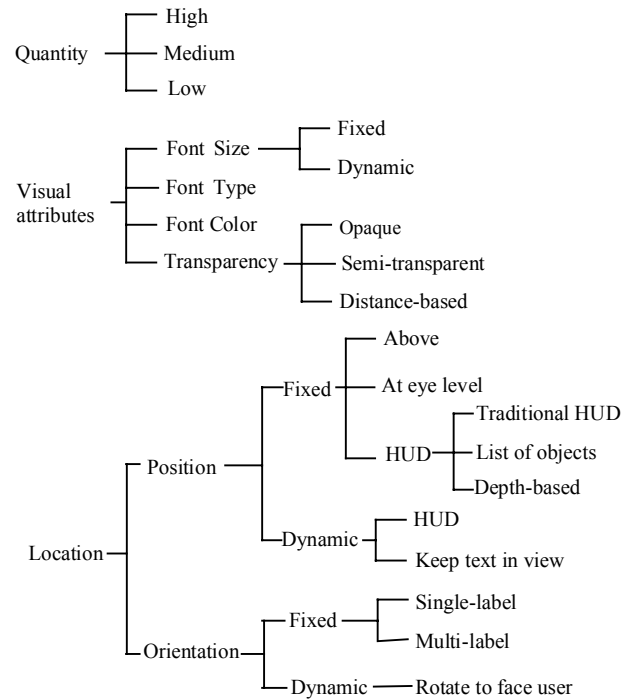


Figure 2. Taxonomy of Text Display Techniques

testbed implementation, but possibilities would include effects such as animation, marquee, and blinking.

Quantity is the amount of text embedded in an IRVE. In this classification, labels with few words are considered low in text content. A brief description of the objects was considered to be medium quantity, and a more detailed description as a high quantity of text content. For example, in our experiment, the low level information is the name of the tree like “Columbus pine”, the medium level information includes the properties of the tree, like the species name, height, diameter, and max age, and the high level information is composed of a short description of the geographic distribution of the tree, etc.

The visual attributes of the text are those aspects that directly affect the legibility of the text and the overall perception of the text in VE. These include font size, color, font type, and the transparency of the text frame. Some of these factors were further classified into lower-level attributes. The size of the text may be fixed or it may change dynamically. In the dynamic condition, text is scaled based on the user’s distance from it, allowing text to be read even when it is attached to a distant object. The transparency attribute refers to the panel on which the text is drawn (not to the text itself). This box may be opaque or semi-transparent, or the transparency may be varied according to distance from the user.

The first location attribute, position of the text, may be fixed or dynamic in space relative to the associated object. In the fixed case, we consider text displayed above

the object, fixed at the user's eye height, and fixed on a HUD. The fixed HUD condition was further classified into three categories. Traditional HUD is fixed-size text displayed on the image plane (e.g., Figure 4). "List of objects" is displaying the text as a list of all objects in view, so a fixed screen space is reserved for text display. Depth-based HUD is the variation of traditional HUD with dynamic font size to produce the illusion of depth. In the "keep text in view" case, the position of the text display is that the entire label is visible.

The second location attribute, orientation, was decomposed into fixed and dynamic categories. The fixed case includes a single label (like a billboard) and the use of multiple labels (e.g. a cube with all faces having the same text content, or two labels perpendicular to each other). In the dynamic orientation condition, the text is rotated to always face the user. It is important to note that the orientation issue only works with WWD, because HUD always displays information, which faces user.

3. Experimental Design

Our aim is to evaluate two navigation and two text layout techniques with four task sets. The two navigation techniques are HOMER and Go-Go based, and two text layout techniques are the WWD and HUD. We hypothesized that for both HOMER and Go-Go based navigation techniques, HUD would perform better than WWD for search tasks in terms of speed.

3.1 Environment

The environment is a medium sized environment with a large number of objects (i.e. hidden areas exist from any viewpoint, and travel from one side to the other requires significant amount of time). The environment consists of trees with different colors (perceptual information), sheds, a house, fence, and abstract text labels for the objects. The testbed was implemented using SVE [13], OpenGL and GLF [18], based on our previous navigation testbed [8].

3.2 Display techniques

The first display technique we evaluated was the WWD. We drew two perpendicular labels on each object and fixed all text labels at the eye-high with the fixed font size and color with transparent background (Figure 3). The second display technique is the traditional HUD, in which the text labels were displayed on the image plane with lines joining the labels and the center of the associated objects (Figure 4).

We implemented an algorithm to calculate the location to display text on the HUD in each frame. It is similar to Bell and coauthor's work [1] where the label is displayed as close to the object as possible. We also assumed that user's attention is on the center of the screen where we




Figure 3. Within-the-World Display Technique (WWD)



Figure 4. Heads-up Display Technique (HUD)

```

For each frame
{
  If (object is visible) {
    Calculate the distance to the user's eye.
    if(the distance < limit) {
      calculate its screen coordinate and slot
      If(the current slot has been taken) {
        if(slot is at the left side of the screen) {
          try its neighbors from right side
          until a free slot is found
        } else {
          try its neighbors from the left side
          until a free slot is found.
        }
      } else {
        display it there;
      }
    }
  }
}

```

Figure 5. Algorithm for Text Layout on the HUD

give a higher priority when searching for close-by vacant slot. If the nearest slots are occupied, the algorithm tries to find a slot in the next closest concentric layer with the slots in the center part of the screen having higher priority. We also calculate the distance between the user's eye and each object in the space. If this distance exceeds a pre-defined limit, the information of that object is not displayed. This strategy is employed to avoid clutter in

the environment. Also, only text labels for the objects within the view volume are displayed. Finally, no two labels overlap.

3.3 Navigation techniques

Both the navigation techniques, Go-Go and HOMER, tested in this experiment, use object manipulation metaphors to move the viewport. The Go-Go [16] uses a combination of linear and non-linear mapping between the virtual hand and the physical hand. The mapping is linear until a point (for nearby object manipulation) and then becomes exponential (to help user to grab far away objects). The HOMER [5] technique uses ray-casting to select objects. The movement is based on a linear mapping between the physical and virtual hands.

We chose these two because both techniques allow fast access to text information at the point-of-interest (POI). Users grab the world (Go-Go) or grab an object (HOMER), and move the viewpoint about that position using hand movements. Additionally, based on the comments from subjects of our proof-of-concept experiment [9], users are less likely to feel disoriented using these two techniques and can maintain spatial memory compared to gaze navigation technique.

3.4 User Tasks

Four types of tasks, representative of the different conditions a user is likely to experience in an IRVE, were considered (Table 1). All questions required searching and had varying levels of difficulty. Task types 1 and 2 required searching information in a single step. Task types 3 and 4 required at least 2 steps.

Type	Task
1	Search for abstract information and then search for perceptual information.
2	Search for perceptual information and then search for abstract information.
3	Search for perceptual information, followed by additional perceptual information, and then abstract information
4	Search for abstract information, followed by perceptual information, and then abstract information

Table 1. Task Types

Task type 1 requires that the user search the abstract text information first, and then find the corresponding perceptual information. An example of this task type is finding the color of the tree named “white fir.”

Task type 2 is conceptually reversed, in that the users find the spatial information first, and then navigate there to answer questions about the related abstract text information. An example of question type 2 is finding the name of the yellow-colored tree in the environment.

Task type 3 requires the user to first find the perceptual information, then to find other perceptual information that is related to the object in question, and to use that result to determine the corresponding abstract text information. We did not use this task type in the experiment because we wanted to test a more direct way to access abstract information. Since task 2 has been used, we skipped this one, but other researchers might find it useful for other investigation. An example of question type 3 would be finding the name of the pink-colored tree, and from the two trees nearest to it, the tree with a maximum age of 650 years.

Finally, task 4 asked the subjects to find abstract data, its corresponding perceptual data, and then abstract data again derived from the perceptual data. An example of this type is asking the subjects to find “Balsam fir” and to find the nearest shed from this tree and its dimensions. During the experiment, we ran two trials of this task, with different number of objects that subjects were to search, therefore with different difficulties. For example, we asked questions about the four sheds in the environment for task set 3 and questions about 28 trees for task set 4.

3.5 Design

Table 2 shows the summary of the experiment design. We used a 2x2 factorial, within-subjects design with 16 participants. The independent variables were the navigation techniques and the text layout techniques. The dependent variable was task completion time. The given treatments were Go-Go with WWD (GW), HOMER with WWD (HW), Go-Go with HUD (GU), and HOMER with HUD (HU). The subjects were treated with four repetitions of each of four adopted tasks (a total of 16 tasks). The order of treatments of navigation techniques, text display techniques, and tasks for each subject was balanced by Latin-square design. To avoid learning effect and to give the subjects a feeling of interest in the environment, the colors of the trees and text labels were changed with each of these combinations. The location of sheds and trees were also changed for task 4.

	GW	GU	HW	HU
Task 1 (Type 1)	<i>s1~s16</i>			
Task 2 (Type 2)				
Task 3 (Type 4 - Low amount)				
Task 4 (Type 4 - High amount)				

Table 2: Experimental Design Summary (Within-subjects)

3.6 Procedure

The experiment was performed in two sessions. During the first session, subjects filled out demographic forms and were trained to use the HOMER and the Go-

Go based navigation techniques. They practiced the techniques until they were comfortable with them, and this session lasted about 5-15 minutes. Text labels were displayed during the training session with the actual text replaced by a series of ‘x’ symbols. Both types of text display techniques were explained and the orientation issues with the WWD technique was specified (readable through limited viewing angles only). We explicitly informed subjects that there would be no two objects with the same name, and that each object has two labels for WWD with the same content.

Before the start of second session, the formal experiment, subjects were asked to complete the assigned tasks as fast and as accurately as possible. After completing all 16 tasks, subjects filled out a questionnaire that included questions for both qualitative and quantitative measures. Responses to these questions were used to determine the perceived difficulties of each technique, type of task and their combinations.

3.7 Equipment and participants

The experiment was HMD-based. The task completion time was logged by the computer. Participants used a wand to make selections in order to travel within the environment. Participants consisted of 8 graduate and 8 undergraduate computer science students. Twelve males and four females between the ages 20 and 34 participated. Two of the participants read the literature on HOMER and Go-Go navigation techniques previously, but neither of them used these techniques before. All participant finished tasks successfully.

4. Results and Discussion

We proved our hypothesis based on the results from the experiment based on task completion time and subjective responses. We found both factors, navigation techniques and display techniques, played a significant role in the completion times of tasks. In this section, we analyze each individual factor and their combinations based on the statistical results, subjects rating, and our observations.

4.1 Subjects performance

We performed a two-factor Analysis of Variance (ANOVA) on task completion time with repeated measure of navigation and text display factors based on each task type. Table 3 presents a summary of results of F and p values of the completion time of 16 subjects. Table 4 shows the mean values of the completion time for each task. From these tables, we can draw the conclusion that (1) Go-Go was significantly better than HOMER and HUD was significantly better than WWD; The HUD was significantly better than WWD in completion time for

tasks 1, 3, and 4. (2) The navigation and the display techniques had no effect on the completion times of task 2.

Subjects also rated WWD as 4.375 vs. HUD as 2.31 on a Likert scale (1 was very easy and 7 most difficult).

We believe Go-Go performed significantly better than HOMER for task 1, 3, and 4 for two reasons. First, Go-Go is simpler in selection, and requires shorter thinking time than HOMER [8]. Go-Go only requires the user to grab air and move about the held point, but HOMER requires object selection then movement. Second, based on our observation, the selection method used in HOMER could cost substantial amounts of time. Subjects did multiple selections most of the time for tasks 1, 3, and 4, but not for task 2. This observation was confirmed by our logging data, where the subjects used an average of 12.4 selections for task 1 and 4 selections for task 2.

We believe HUD performed significantly better than WWD for task 1, 3, and 4 for three reasons. First, the text information in a HUD is visible from all directions and always faces the user compared to the fixed and restricted view in WWD, although we have made efforts to avoid this advantage to HUD. We did not use the option of “always facing user” from our taxonomy for WWD because most of the users from our pilot study indicated that they liked the labels to be fixed as it helped with orientating themselves.

	Navigation techniques		Text display techniques	
	F (1, 60)	p	F (1, 60)	p
Task 1	12.13	0.0009*	11.92	0.001*
Task 2	1.289	0.26	2.23	0.14
Task 3	11.54	0.001*	4.41	0.04*
Task 4	11.02	0.002*	4.35	0.04*

Table 3. F and p Values of Two-Factor ANOVA with Repeated Measure of Interaction and Text Display (Slots marked with * indicate statistical significance)

	Mean value of the completion time (s)			
	Navigation techniques		Text display techniques	
	Go-Go	HOMER	HUD	WWD
Task 1	56	138	55.6	139
Task 2	33.4	42.8	34.53	41.68
Task 3	48.9	66.44	43.5	71.9
Task 4	74	126.6	58.5	142.2

Table 4. Mean Value of the Completion Time of Navigation and Text Display for Tasks

Second, the text on a HUD is more readable. The size of the font is relatively small but clear. For WWD, we have to render a bigger font in order to have a clear view (Figure 3, 4). Also, the text might not be readable due to

the object being far away from the user. In this case, users had to travel long distances to view the text labels, which for the most part were blocked by the objects nearby.

Third, the occlusion problem is more apparent in WWD compared to HUD, and subjects had to travel longer in order to view the text information. In WWD, the text is occluded if the associated object is behind another object. However, this is not the case for HUD, which always displays the text on the screen. HUD does have the disadvantage with regard to finding the association between the text and the perceptual information. This is because the lines linking the labels and the abstract information are also drawn on the image plane and are not occluded. For example, in Figure 4, the right-most tree is *Siberian larch*, but the two labels of trees that are behind this tree are also shown. Therefore, to correctly identify the name of this tree, the user should change viewport to see the correspondence. However, most subjects did not find this as a problem for getting the right link, although there were two errors among all 128 questions.

In order to see the combinations of interaction and text layout techniques, we compared GW, GU, HW, and HU for the four tasks conditions. Figure 6 shows the mean values of completion time per question type. A similar trend has been observed for different task types.

It can be seen that the HW combination had the worst

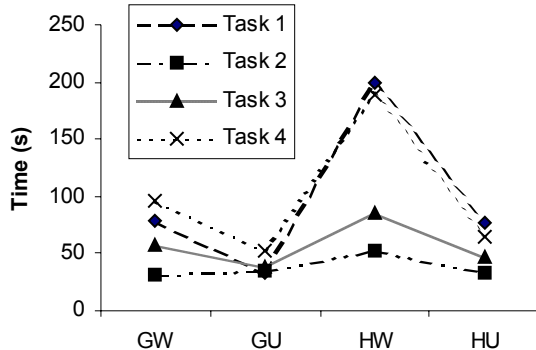


Figure 6. Combination of Text Display and Navigation Techniques vs. Tasks Completion Time

performance, while GU had the best. A single-factor ANOVA of GU and HW on each task was performed. We found that the GU combination is significantly better than the HW combination except for task set 2.

The spike that appears on the completion time of task 1 and 4 seems to be abnormal under the HW combination. We believe it is due to the compounding difficulty level in both the display and the travel techniques. For the tasks using GU combination, the performance was the best because of the inherent ease of navigation and ease of information access from the text displayed. Also, we observed that all subjects rotated around an object to

search for information instead of following a logical pattern although they were told of the two fixed directional labels for each object.

Task set 3 and Task set 4 are the same tasks, except the search required reading different amounts of information with different distributions, but the spike does not appear in task 3. This is because of the change in the amount of information to be searched compared to task 4, and we found the combination of interaction and display techniques had significant effect on task completion time ($p=0.04$). As we mentioned before, task set 3 required the subjects to identify a shed in the environment compared to finding a tree for task sets 4, and then locating additional perceptual and abstract information. But because the number of sheds in the environment was only four compared to the 28 trees, and the sheds being highly distinguishable from the trees (fewer numbers and larger size), subjects found it easy to complete these tasks.

Subjects also rated the BW as the worst (average response 4.3) and GU the best (average response 2.3) combination based on a perceived level of difficulty scale of 1 to 7, where 1 was most easy and 7 most difficult.

4.2 Differences in tasks

After careful analysis of the four sets of tasks and the statistical results, we believe the reason for significant differences in task completion times might be also because of the *difficult* naïve search tasks compared to the primed search. We refer to the tasks as *difficult* because the users had to search for one piece of information from a few tens of information (in our environment we had a total of 32 labels). Task set 1 is a single-step naïve search in which subjects were asked to find the color of a tree. The trees to be searched in some of the tasks were well hidden and the subjects had to look through considerable amount of information and travel over a large distance in order to find the answer.

Compared to task set 1, task set 4 added one more extension to the naïve search, i.e., to find perceptual information in the world that is close by and then to read the abstract information (here, text labels) corresponding to it. This step can be considered as a primed search because subjects knew where the perceptual and symbolic data were. Also, subjects traveled very short distances or just turned their bodies to read the information during the primed search, because the objects were arranged close to each other and no other unrelated or distracting information was displayed at the same time.

Task set 2 consists of naïve searches but was *not difficult* overall. Subjects started from a point located on one side of the environment and after looking at the objects that are directly visible from the starting point, they generally made an effort to go forward. During this

first exploratory move, most subjects found the location of the tree with the required color. Once the particular object (here a tree) was located in space, subjects just tried to get close to the object to read the abstract data displayed on it.

The lack of significant difference in performance times for task set 2 was because of the manner in which this environment was arranged. The disadvantage of the WWD (subjects can only view information from certain angles) was mitigated because all text labels were initially displayed facing the starting point of the subjects. Therefore the subjects read the information directly without the need for traveling around objects. And for HUD, most subjects still changed their views a little in order to verify that the information belongs to the object they are looking at and not to the partially visible objects that are behind the object of their focus.

5. Impact on IRVE Design

IRVEs show great promise in applications like education [6], immersive modeling [3], surgical simulation [2], etc. Overall, our main contributions to this area include: A systematic approach to study IRVE text layout and navigation techniques. More specifically, the take away lessons are:

- HUD is a better display technique for naïve search tasks related to abstract information in densely packed environments. This is because the user can directly access the abstract information without the need to locate the actual position of the perceptual object in a crowded world. The impact of this finding on education could be profound. For example in medical human anatomy teaching application, where there is heavy clutter due to the high number of objects, if the user is looking for an object called ‘pancreas’ without knowing where it actually is, the HUD makes it possible for the user to scan the labels of the objects joined to the perceptual data. After identifying the label, it would be possible to find the object by following the line joining the label to the object (Abstract to perceptual information task in our experiment).
- Go-Go interaction technique is better suited for navigation in environments that require easy and more flexible movements. An example of such an environment could be architectural walkthroughs. Here the users would like to have easy and simple navigation through the building structures. One of our conclusions regarding the two interaction techniques is that, the Go-Go technique performs better because there is no explicit target selection. This conclusion helps in making the choice of interaction techniques easier. HOMER would be difficult in this application because of the inherently few selectable objects present in architecture walkthrough

applications. Therefore, for applications like this, Go-Go would most certainly be better than HOMER.

6. Conclusions and Future Work

Understanding the text layout techniques, navigation techniques and their relationships is crucial to improve the usability of IRVE applications. We provided a taxonomy of the text layout techniques and conducted a testbed evaluation to study two text displays, navigation techniques and their relationships. The quantitative and qualitative results of our study indicate that for highly dense environments, HUDs are more effective. Go-Go technique was reported to be more comfortable and easy to use because of its simplicity and directness. According to our observations, the difficulty factor in HOMER appears to be because of the additional object selection component of the technique. Applying these observations to other manipulation based navigation techniques, the following inferences can be made: techniques that use ray-casting and require the user to select an object to navigate would probably have less usability; techniques that are direct and have more degrees of freedom are likely to perform better. This might suggest new navigation techniques that need to be designed for IRVE.

The best combination of navigation technique vs. text display techniques (out of the four studied) appear to be Go-Go and HUD. There are two possibilities why the techniques performed better than others: the fact that the user started from abstract information and then proceeded to locate the related perceptual information or because of the difficult naïve search tasks.

Our future work could be to perform a comprehensive study by considering all navigation techniques, and then to come up with new navigation techniques. This could provide a theoretical basis for judging these two display techniques over a wider variety of potential applications. Such work could easily result in useful guidelines for future application of text based IRVEs. Also, extending the testbed for other types of abstract information like voice, images, etc appears to be the next step in IRVE research.

References

- [1] Bell, B., Feiner, S., & Hollerer, T., “View Management for Virtual and Augmented Reality,” *ACM Symposium on Interface Software and Technology*, 101-110, 2001.
- [2] Billingham, M., Savage, J., Oppenheimer, P., & Edmond, C., “The Expert Surgical Assistant: An Intelligent Virtual Environment with Multimodal Input,” *Proceedings of Medicine Meets Virtual Reality IV*, 590-607, 1995.
- [3] Bolter, J., Hodges, L.F., Meyer, T., & Nichols, A., “Integrating Perceptual and Symbolic Information in VR,” *IEEE Computer Graphics and Applications*, July 1995.

- [4] Bowman, D., "Conceptual Design Space: Beyond Walk-through to Design, D. Bertaol (Ed), *Designing Digital Space*, New York: John Wiley & Sons, 1996.
- [5] Bowman, D. & Hodges, L.F., "An Evaluation of Techniques for Grabbing and Manipulating Remote Objects in Immersive Virtual Environments," *Proceedings of the ACM Symposium on Interactive 3D Graphics*, 35-38, 1997.
- [6] Bowman, D., Hodges, L.F., & Bolter, J., "The Virtual Venue: User-Computer Interaction in an Information-Rich Virtual Environment," *Presence: Teleoperators and Virtual Environments*, 7(5), 478-493, 1998.
- [7] Bowman, D.A., Wineman, J., Hodges, L.F., & Allison, D., "The Educational Value of an Information-Rich Virtual Environment," *Presence: Teleoperators and Virtual Environments*, 8(3), 317-331, 1999.
- [8] Bowman, D., Johnson, D., & Hodges, L.F., "Testbed Evaluation of Virtual Environment Interaction Techniques," *Presence: Teleoperators and Virtual Environments*, 10(1), 75-95, 2001.
- [9] Bowman, D., North, C., Chen, J., Polys, N.F., Pyla, P.S., & Yilmaz, U, "Information-Rich Virtual Environments: Theory, Tools, and Research Agenda," *To appear in the Proceedings of the ACM Virtual Reality System Technology*, 2003.
- [10] Card, S., Mackinlay, J., & Shneiderman, B. (Eds), "*Information Visualization: Using Vision to Think*," Morgan Kaufmann, 1999.
- [11] Darken, R.P. & Cevik, H., "Map Usage in Virtual Environments: Orientation Issues," *IEEE VR '99*, 133-140, 1999.
- [12] Hollerer, T., Feiner, S., Terauchi, T., Rashid, G., & Hallaway, D., "Exploring MARS: Developing Indoor and Outdoor User Interfaces to a Mobile Augmented Reality System," *Computers and Graphics*, 23(6), 779-785, 1999.
- [13] Kessler, G.D., Bowman, D., & Hodges, L.F., "The Simple Virtual Environment Library: An Extensible Framework for Building VE Applications," *Presence: Teleoperators and Virtual Environments*, 9(2), 187-208, 2000.
- [14] Lampton, D., Knerr, B., Goldberg, S., Bliss, J., Moshell, J., & Blau, B., "The Virtual Environment Performance Assessment Battery (VEPAB): Development and Evaluation," *Presence: Teleoperators and Virtual Environments*, 3(2), 145-157, 1994.
- [15] Mine, M., "Virtual Environment Interaction Techniques," *Technical Report TR95-018, UNC Chapel Hill CS Dept*, 1995.
- [16] Poupyrev, I., Weghorst, S., Billinghurst, M., & Ichikawa, T., "The Go-To Interaction Techniques: Non-linear Mapping for Direct Manipulation in VR," *Proceedings of the ACM Symposium on User Interface Software and Technology*, 79-80, 1996.
- [17] Poupyrev, I., Weghorst, S., Billinghurst, M., and Ichikawa, T., "A Framework and Testbed for Studying Manipulation Techniques for Immersive VR," *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, 21-28, 1997.
- [18] Survey of OpenGL Font Technology, <http://www.opengl.org/developers/code/features/fontsurvey/> (Available Aug 2003).
- [19] Anatomy Landmarks, <http://ovrt.nist.gov/projects/vrml/h-anim/landmarkInfo.html> (Available Aug 2003).