The Educational Value of an Information-Rich Virtual Environment

Doug A. Bowman[1], Jean Wineman[2], and Larry F. Hodges[3]

[1, 3] College of Computing

[2] College of Architecture

[1, 2, 3] Graphics, Visualization, and Usability Center

Georgia Institute of Technology

Atlanta, Georgia

ABSTRACT

"Information-rich" virtual environments consist not only of three-dimensional graphics and other spatial data, but also include information of an abstract or symbolic nature that is related to the space. We present evidence that an environment of this type can stimulate learning and comprehension, because it provides a tight coupling between symbolic and experiential information. In our virtual zoo exhibit, students can explore an accurate model of the gorilla habitat at Zoo Atlanta, and also access information related to the design of the exhibit. This paper discusses the rationale behind the design of the application and the interaction techniques used to obtain information. We also present the results of an evaluation, showing that students who used the virtual environment had better test scores than those who only attended a lecture on the material. We show that the virtual experience allowed students to learn information directly, and also motivated them to better learn and understand material from a traditional lecture.

INTRODUCTION

The goal of virtual environment (VE) research is not to produce more realistic environments, faster 3D graphics, better sensory cues, or low latency. Rather, all of these are the means by which we hope to achieve the actual goal: useful applications that will benefit people. Although most are still in the research lab, a few categories of VE systems have shown great promise, including architectural walkthrough (Brooks et al, 1992), exposure therapy for phobias (Hodges et al, 1995), and training for hazardous duty (Tate, Sibert, & King, 1997). All of these share the characteristic that their success depends only on producing a satisfactory and believable experience to the user. That is, they must cause the user to suspend her disbelief, and to feel on some level that she is actually in the displayed environment.

It has long been suggested that education should be another key application area for VEs (Durlach & Mavor, 1995), and this follows from the argument that the experience should be the main ingredient of a successful VE. After all, "experience is the best teacher." However, the amount of published work in this area is small, and the number of systems which have been shown to be practical is even smaller. Why?

It seems that experience can only take a student part of the way to learning and understanding a subject. In most cases, it is necessary to have background knowledge, peripheral information, reflection, *and* experience before the subject can be comprehended by the student. Consider that high school students have experienced a phenomenon such as the refraction of light hundreds of times, but do not come to a complete understanding of it until they study optics in their physics classes.

Thus, experience is only one part of a practical education. In fact, it is dangerous to rely solely on experiences for learning, since incorrect mental models can often arise logically from experiential data (for example, one could draw the incorrect conclusion that acceleration due to gravity changes based on an object's mass by observing that a brick falls faster than a feather). Certainly, there are some concepts that can not be experienced directly in the world, such as the interactions between sub-atomic particles. In cases such as this, a virtual environment can provide an important first step in understanding, but other knowledge and teaching will also be necessary to produce complete comprehension.

This paper presents an educational virtual environment that provides both experiential and abstract information in a tightly coupled manner. In this way, students can avoid the pitfall of relying only on experience as a learning tool, but also have the opportunity to relate information that would normally be received in a lecture setting to an actual experience and a three-dimensional space. We call this an "information-rich" virtual environment (Bowman, Hodges, & Bolter, 1998, to appear).

Our system builds on the work of the virtual reality gorilla exhibit (Allison, Hodges, & Wineman, 1997), and is designed to teach college students about the design principles used in constructing an animal habitat within a zoo setting. Users can move about the habitat to see it from any point of view, and can also obtain information in text, audio, or image form relating to the design of various aspects of the exhibit. Before we discuss the specifics of the application, we will review some related work in educational and information-rich VEs. After describing the system, we will present an evaluation in which we tested the educational value of our system in the context of a college course on environmental design. We will conclude with a discussion of the results and further work that we hope to do in this area.

RELATED WORK

Information-Rich Virtual Environments

Many systems have been developed which use a three-dimensional space to present some form of information to the user. These include both immersive virtual reality systems and desktop 3D applications. There are basically two categories of such systems: scientific simulations and database visualizations.

Scientific simulations present a view of scientific data within a 3D environment, often with animated objects. Generally, they consist of abstract objects which are too small for the naked eye, such as atoms (Bergman et al, 1993), too large to be comprehended, such as the solar system (Song & Norman, 1993), or invisible, such as electromagnetic fields (Dede, Salzman, & Loftin, 1996) or fluid flow lines (Bryson & Levit, 1992). Users can examine these simulations from various positions, detect patterns that would not be obvious without the visualization, and make changes to conditions and immediately visualize the results.

Database visualizations take a complex and abstract dataset and organize it into an understandable visual representation, which can be navigated and accessed by the user (e.g. Benford, Snowdon, & Mariani, 1995; Fairchild, Poltrock, & Furnas, 1988; Fairchild, 1993; Risch et al, 1996; Robertson, Card, & Mackinlay, 1993). Here abstract properties of the data are mapped into perceptual qualities, such as size, shape, color, or motion, and relationships between pieces of data are represented spatially. The resulting 3D visualization can reveal patterns in the data due to spatial groupings which are not obvious from the original dataset.

Both of these types of information spaces present abstract or non-viewable information using a perceptual (geometric) form. Other forms of information, such as text or speech ("symbolic" information) are not usually present except as labels for the geometric objects. On the other hand, information-rich virtual environments *embed* symbolic information within a realistic 3D environment. For example, a virtual college campus may contain text describing various streets or buildings or spoken audio giving the characteristics of the athletic facilities. In this way, symbolic and perceptual information are integrated in a single environment (Bolter et al, 1995).

Our previous work in the area of information-rich VEs was conducted in the context of an application called the Virtual Venue (Bowman, Hodges, & Bolter, 1998, to appear). In this system, users could move about an accurate model of the Georgia Tech Aquatic Center, and obtain various types of information regarding the design and use of the venue and the sports of swimming and diving. A usability study revealed that the most effective types of information were those which were "tightly coupled" to the environment. That is, the information content was pertinent to or otherwise associated with the object or location in 3D space.

Experience with the Virtual Venue indicates that information-rich VEs can be an effective means of information retrieval, but our usability study did not allow us to compare it with other information gathering media, such as multimedia presentations, the World Wide Web, or printed text. Thus, in our current study, we have created a comparison between traditional lectures and classroom teaching augmented with the use of a virtual environment.

Virtual Reality and Education

Wickens (1992) gives an overview of some of the salient features of virtual reality and their relation to education. He argues that the closed-loop interaction style of VEs should increase learning and retention, because it requires effort on the part of the user to continuously choose their position, view orientation, and action, rather than being passively guided by the system. However, some of the other characteristics of VEs, such as three-dimensional and ego-referenced viewing, and "natural" interaction, may actually reduce a student's retention because he has not been required to put forth as much effort. Thus, he claims that the goals of user interface design (e.g. reduce mental workload for the user) and educational software design actually conflict in some ways. We would argue that a distinction needs to be made between cognitive load from task-related activities and system-related activities. That is, Wickens is correct that in educational environments, learning activities (task-related) should require effort and choice on the part of the user; however, system-related activities, such as selecting an object, changing display mode, or finding a menu item should require as little cognitive processing as possible. We do not want users to be distracted from learning because they cannot figure out how to use the interface.

Wickens also highlights the need for educational systems to teach the relationships between pieces of information, rather than just isolated facts. He says that "...the educational benefit of a VR experience should be enhanced to the extent that the learner is exposed to material from *both* a VR and a more abstract perspective, and learner attention is directed to the linkages or relatedness between these two perspectives" (Wickens, 1992). This is precisely the goal of an information-rich VE: users obtain both spatial and abstract information within the same context, and the relationships between the two are made evident by the location and type of embedded information.

There have been several reported VE systems that are intended for educational purposes, and it will be instructive to review some of them here. As we have already mentioned, there are scientific simulations and data visualizations that could be considered educational, since they "teach" the user information that might not have come to light without the 3D visualization. However, in the case of applications such as the virtual wind tunnel (Bryson & Levit, 1992), the users are already experienced in the field of computational fluid dynamics, and thus can understand the visualization as presented. Thus, such systems are not teaching concepts, but instead are demonstrating or applying concepts that may then reveal further specific information.

Two scientific simulation applications are intended for conceptual education, however. The ScienceSpace system (Dede, Salzman, & Loftin, 1996) and a VR physics simulator from Rice University (Brelsford, 1993) are both designed to teach important physics concepts using an immersive virtual environment. Both of these systems use a *constructivist* educational theory, meaning that students learn through personal interaction with the material (in this case manipulating physical elements and observing their behavior).

ScienceSpace has three virtual worlds which teach the concepts of Newtonian mechanics, electrostatics, and molecular structure and dynamics, respectively. These worlds are designed to promote learning through experience and experimentation. Thus, students can "become" a point mass to learn about collisions or move a charge through an electric field to see the magnitude and direction of the force on that charge. This allows students to experience phenomena that are not accessible in the physical world, and hopefully gain the ability to predict the results of a given situation. However, the authors acknowledge that incorrect mental models can be created based on experience alone, which could mislead the student. Thus, students are carefully guided through the learning process, with appropriate background information inserted when needed. In fact, the authors say that they are currently developing an automated “coaching” system which will embed feedback into the virtual environment, making it into an information-rich VE. No quantitative learning results have been reported for ScienceSpace.

The Rice physics system is not described in great detail, but is said to be a virtual representation of a physics laboratory, complete with pendulums and masses, and controls to change the force of gravity, location, air drag, friction, and so on. The author describes a study in which both junior-high and college students were divided into two groups. One used the VR system for one hour while the other attended a lecture over the same material. The groups were tested four weeks later, and the results showed that the VR group had increased their physics knowledge by a significantly greater amount. This clearly shows the promise of VR as an educational medium. However, it is difficult to tell from the article how students were given the necessary background information they needed in order to understand the results of their experiments in the virtual laboratory. The authors do say that students worked on specific written problems during their sessions, which may have themselves been sources of abstract information or formulas, and which definitely guided the experimentation of the users. In any case, it is doubtful that junior high students would discover the laws of motion by chance by spending one hour in a VR simulation.

The predecessor of our current design education application is the virtual reality (VR) gorilla exhibit (Allison, Hodges, & Wineman, 1997). The VR gorilla exhibit is also an educational application, designed to teach middle-school students about gorilla behaviors, vocalizations, and social interactions. Students begin in the visitors center, where they acclimate themselves to the virtual environment and the techniques for movement. By moving through the viewing window, they take on the persona of an adolescent gorilla, and the virtual gorillas in the exhibit react as they would to a young gorilla. The virtual gorillas have accurate movements, sounds, and behaviors, so that if the user makes a social faux pas, such as entering the personal space of the male silverback, he will see a realistic reaction (escalating annoyance leading to a “bluff charge” and roar).

The VR gorilla exhibit originally relied solely on experience to teach the students - it was hoped that they would draw their own inferences from using the system, since the gorillas’ interaction was not very complex (Wineman et al, 1997). Some students did this, but most needed assistance to understand the behavior of the virtual animals. This meant that a gorilla expert needed to stand next to the user, offering information and advice to aid the learning process. Eventually, we decided to automate most of this “expert information” by providing spoken information played through headphones at appropriate times, locations, or events (Allison et al, 1997). For example, when the user stares for too long at an older gorilla, an audio segment informs the user that the gorilla is becoming annoyed because of the staring, and tells the user what she should do to placate the other animal (look away and move away). We also implemented “mood indicators,” which are icons that indicate the current state of each of the virtual gorillas (content, annoyed, or angry).

These enhancements effectively make the VR gorilla exhibit into an information-rich VE. Even with the simplicity of the information to be presented, experience alone did not allow total comprehension. Both experience and abstract information were needed to give students a complete understanding.

VIRTUAL ZOO EXHIBIT FOR DESIGN EDUCATION

Our current application of virtual environments for design education uses the habitat model from the VR gorilla exhibit, which is an accurate model of the largest outdoor gorilla habitat at Zoo Atlanta (Figure 1). Elements including the visitors center, the moat, the terrain, and the trees, rocks, and logs are all modeled and positioned as they are in the physical exhibit. However, our focus has changed from teaching middle-school students about gorilla behavior to teaching college students about habitat design. The learning goal of the students who use this system is an understanding of the philosophy of environmental design and of the specific design decisions that were made for the Zoo Atlanta gorilla exhibit. Therefore, we have embedded new symbolic information content within the virtual exhibit, and have included new interaction techniques with which to access this information.

Like its predecessor, the design education application is based on the Simple Virtual Environment (SVE) library (Kessler et al, 1994), a software support library that takes care of the details of tracking, rendering, event-handling, and the like. The system runs on a Silicon Graphics Indigo2 Max Impact, and uses a Virtual Research VR4 head-mounted display (HMD) for visual output. Tracking is performed using a Polhemus Fastrak with three enabled receivers, including a special stylus with a button (see Figure 3a). One tracker is used for head position and orientation, while the other two allow us to implement a “pen & tablet” interaction metaphor, described below.

Figure 1. The Virtual Gorilla Exhibit

This application is quite different from the other educational VEs that were mentioned earlier. First, our system supports design education, while most previous efforts focused on math or science education. Design is a much less concrete subject, which may make it more difficult to teach. However, students are not required to understand complex formulas, and aesthetics play a major role. These characteristics make design education a natural fit for a VE. Second, our system is not based on constructivist learning, as most previous educational VEs have been. Rather, we have chosen to present design philosophies in an abstract form within a model of a space which follows those philosophies. For design education, students must be able to see examples of design concepts at work before they can begin to construct their own designs. Our system does incorporate some constructive elements, which allow students to modify the habitat's design, but these tools were not used until students had obtained information on design concepts (we will report on the use of the design tools in another article). Finally, as we have stated, our system explicitly integrates both symbolic and perceptual information in a single environment, so that learning and comprehension are enhanced.

Embedded Information

The design education application makes use of several embedded media types to accomplish its goal of presenting relevant information about habitat design within the context of the habitat itself. The most ubiquitous form is spoken audio. The virtual habitat contains 19 audio clips describing many aspects of the design. They range from general concepts regarding the philosophy of environmental design to quite specific pieces of information on features of the gorilla habitat itself. Some of the clips are taken directly from a recorded interview with one of the gorilla keepers from Zoo Atlanta. In general, we tried to use audio for most of the embedded information since it allows the user to view the environment at the same time as he is receiving information. For example, the student can look at the structure of the moat while listening to an annotation describing its design and construction.

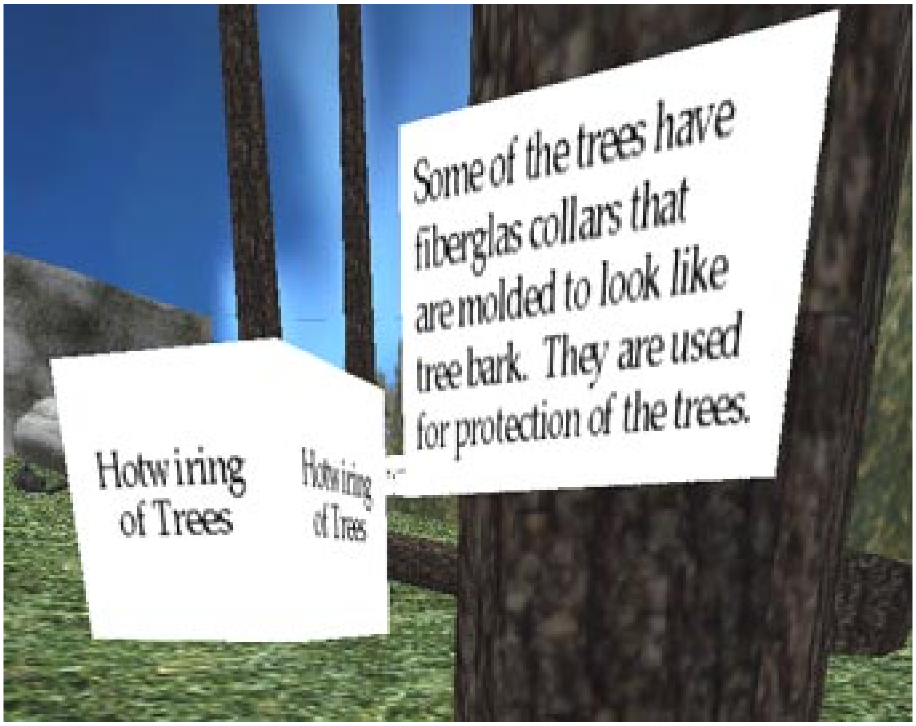

Figure 2. Audio (left) and Text Annotations in the Virtual Habitat

Some audio annotations are played automatically based on the current state of the system or the user's position within the environment. For example, some introductory material is played when the system is initialized, and a description of the design of the outdoor viewing areas is given when the user goes there. Most annotations, however, are represented by cubes in the environment (Figure 2) and are triggered explicitly by the user, as described in the section on interaction techniques. The user therefore explores and learns at her own pace and based on her own interests, rather than under the control of the system. All of the annotations were developed using a VE audio annotation toolkit, developed for the Virtual Venue system.

There are also five text annotations in the virtual exhibit, in the form of signs that are located on surfaces within the environment, such as on the walls of the visitors center or on a tree (Figure 2). Text was chosen over audio for information that was more complex and detailed, and therefore might require scanning back and forth and rereading. Audio is obviously less suitable for these purposes. We also inserted text annotations in areas which were already cluttered with audio clips. For example, we felt it would be confusing to have two audio annotations regarding the gorilla night building, so we instead used one audio clip and one text annotation.

Finally, we have embedded two images in the virtual habitat. These are used to enhance understanding of text and audio annotations, and to convey spatial relationships that would be difficult to describe in words. One of the images is a map of the entire gorilla exhibit at Zoo Atlanta, showing the habitat in plan view and its location relaative to the other three gorilla habitats, and illustrating the idea of a “zoogeographic” and “bioclimatic” zone, in which animals and plants from similar geographic and climatic regions are grouped together. The other image is a photograph of a gorilla playing near the window of the visitors center, which illustrates a point about usage of various areas of the habitat. The picture and an audio clip describing habitat usage are presented simultaneously to the user.

The embedded information was gleaned from a variety of sources. Interviews were conducted with both the principal architect in charge of the design of the gorilla habitat and one of the gorilla researchers at Zoo Atlanta. We also obtained some general information on the philosophy of the designers in several of their publications (Coe, 1985; Coe & Maple, 1987). Finally, we used maps and other information about the zoo that appears on the WWW home page for Zoo Atlanta[1].

Unlike the VR gorilla exhibit, the virtual gorillas in the design education system do not react to users. We felt that interaction with the gorillas would distract users and hinder the goal of teaching students about habitat design. However, there are gorillas in the environment, and they are used to help underscore some design concepts. The gorillas are arranged on a hillside in a similar way to the arrangement in which they have been observed using a

---

[1] http://www.zooatlanta.org

hillside in the wild. This fact informed the design of the exhibit's terrain, and is explained in an audio annotation. Also, the virtual gorillas help to give the student a sense of the scale of the design, which is quite important from an architectural point of view.

Interaction Techniques

An information-rich virtual environment cannot be effective unless the user can easily and efficiently access the information. For this reason, usable interaction techniques are a necessity and cannot be overlooked. We have chosen techniques for user navigation and object selection based on their simplicity, efficiency, and unobtrusiveness.

Usable navigation techniques should allow the user to move around the habitat freely and efficiently, while ensuring that the user does not become lost or disoriented in the 3D space. This is a difficult combination to achieve, since more freedom generally results in higher disorientation, and reducing disorientation depends on constraints, which diminish freedom of movement. Our previous studies on the subject of VE travel techniques (Bowman, Koller, & Hodges, 1997) aided us in combining two techniques which allow complete freedom of movement while also providing aids to reduce disorientation.

The first technique uses the stylus as a pointing device. Users point in the direction in which they wish to move, and hold down the stylus button to travel in that direction with a constant velocity. Users can see a representation of the stylus in the virtual world, so they can visualize the direction they are pointing. Our previous experiments have shown that this technique is accurate and efficient for most user positioning tasks, whether the user is traveling directly to an object or simply moving to a location to obtain a specific view of the world. The pointing technique decouples the user's head orientation from the direction of travel, allowing him to move in any direction regardless of the direction of his gaze. One important feature from a design perspective is the ability to fly upwards to get a bird's eye view of the entire habitat.

Some disorientation is prevented by keeping the user within the habitat using some simple collision detection routines. Users are not allowed to go below the ground or beyond the walls of the surrounding moat. However, flying in three dimensional space is still difficult for many people, and they may not be able to maintain spatial awareness of their surroundings, causing disorientation.

The second technique addresses some of these concerns. It is based on the "pen & tablet" metaphor (Angus & Sowizral, 1995) which we used in the Virtual Venue. In their non-dominant hand, the user holds a physical tablet (Figure 3a), and a visual representation of the tablet can be seen in the virtual environment. A map of the habitat appears on the tablet, and a red dot represents the user's current position (Figure 3b). The user can move by placing the stylus over the red dot, holding down the button, and dragging it to a new location on the map. When the button is released, the user is flown smoothly to the new position in the environment.

a. b.

Figure 3. a) Physical Devices Used in the "Pen & Tablet" Interaction Metaphor, b) User's View of the Virtual Tablet and Stylus

This technique has several advantages. First, the map displaying the user's position effectively combats disorientation. If the user feels lost, she can look at the map and find her position relative to some known landmarks. This is true whether the user has been using the pointing technique or the dragging technique. Second, by dragging to a specific location on the map, the user can move quickly to the area of interest, without having to navigate through the actual 3D environment. Third, since the user does not actually change position until he releases the stylus button, he can watch as he travels smoothly from his current location to the new one, and spatial awareness may be increased. Also, the pen & tablet metaphor itself has several advantages for many types of interaction, due to its unobtrusiveness (the tablet may be put aside if not needed), its inherent constraint (the physical surface of the tablet guides the stylus), and its use of two-handed interaction, with the dominant hand working relative to the non-dominant hand (Hinckley et al, 1997).

We considered a “view-up” map, which rotates so that the map is constantly aligned with the user’s point of view. However, Wickens (1992) and others have argued that such a display, while possibly enhancing navigation performance, may reduce retention of the layout of the 3D space. We have instead given users a fixed frame of reference within the ego-centric frame of reference, which forces them to expend some effort in forming mental links between the two types of views, and should cause increased retention of the spatial data.

The user can combine the two navigation techniques in any way. In our experience, most users utilize the pointing technique for exploration and obtaining interesting views of the habitat, and the dragging technique to quickly move to a new area when the information gathering task demands it. The map and the constraints on movement help the users to maintain awareness of their spatial location.

Our application also required a technique for object selection, since we wished to allow users to control the playback of audio annotations. Audio clips were represented by white cubes within the environment itself, on which were printed a title phrase so that users could know the theme of the annotation without playing it (this also prevented users from playing the same annotation over and over again). The annotations were also represented by icons on the map, so that users would be able to see the location of audio information and move to these locations more quickly.

To begin playback, the user selects the cube, and she can also stop it during playback by selecting the cube again. There are several requirements for the selection technique. First, it must be cognitively simple to use, since we want students to focus on the content of the annotations, not on the interaction. Second, it should be useful at a distance, since users need to be able to look around the environment while the annotation is playing. If they are forced to move in close proximity to the cube, it might block their view. Finally, the selection technique should integrate nicely with the navigation techniques we have chosen.

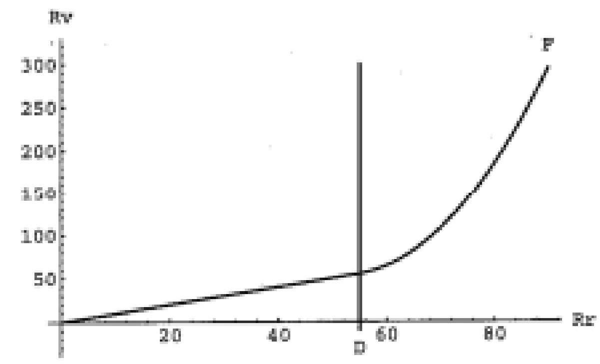

Figure 4. Non-linear Function for Virtual Arm Length Using the Go-Go Technique (reproduced from Poupyrev et al, 1996)

In a previous study on selection and manipulation techniques for immersive VEs, we had found that the ray-casting technique, in which the user points a virtual light ray at an object to select it, was ideal for object selection (although it was not as useful for manipulation). Unfortunately, this technique requires a button to activate the light ray, and our stylus button was already being used for the pointing technique for navigation. We could leave the light ray active at all times, but this might obscure views of the environment. Instead, we chose to use the "Go-Go" technique (Poupyrev et al, 1996) for object selection. This technique allows the user to stretch his virtual arm well beyond the length of his physical arm, using the mapping function shown in Figure 4. When the physical hand is beyond a certain distance from the user's body, the virtual arm begins to grow at a non-linear rate. Our study showed this technique to be nearly as efficient for object selection as ray-casting, although it may not be as accurate for small objects. Since our annotation cubes were fairly large, the Go-Go technique would allow easy object selection from a distance with no change to our navigation techniques.

EVALUATION AND TESTING

In order to test the efficacy of our virtual environment system for design education, we designed and implemented an evaluation within the context of a class on "The Psychology of Environmental Design," taught jointly by the College of Architecture and the department of Psychology. The class already contained a major section on the design of zoo exhibits, so our system fit neatly into the content of the class.

The evaluation was designed to test two hypotheses:

1. Students who augment the normal class presentations by using the virtual zoo exhibit will have greater understanding and increased retention of the material, and thus will perform better in an evaluation.
2. Students who use the virtual zoo exhibit will be more motivated to learn when the same material is presented in class and will be able to form more mental associations, and thus will perform better in an evaluation.

Note that in neither of these cases do we surmise that the virtual zoo exhibit should replace traditional classroom teaching; rather, we feel that it is best used as a supplement to the normal procedure of the class. It would be extremely difficult to create an information-rich VE that would gracefully contain the complexity of the material that could be presented in an hour-long lecture. Moreover, the attention span and patience of VE users is generally low. Therefore, we feel that the VE is best used to introduce material, create associations between abstract and spatial information, and to motivate further learning.

Method

The class of 24 students was divided into three groups: two groups of nine students and one group of six. Equal groups of eight students each were not possible, since the instructor wished to divide the class based on project teams, each of which had three members. The groups are summarized in table 1. Students were randomly assigned to project teams, and project teams were randomly assigned to groups.

The *control group* (nine students) had no change to the normal progress of the course. That is, they simply attended lectures, one of which covered material on exhibit design in general and the design of the gorilla habitat in particular.

The *information group* (nine students) attended class lectures, and also used the VE system to explore the virtual habitat and to gather embedded information. This was done in two phases: first, students explored the habitat with information disabled, in order to understand the layout of the exhibit; second, students gathered information within the VE using the techniques described above. Students were not given a time limit, but it was suggested that they spend 5-10 minutes in phase 1 and 10-15 minutes in phase 2.

The *habitat group* (six students) attended class lectures, and also used the VE system to navigate about the virtual gorilla habitat. They were able to explore the visitors building, the hillside, the moats, and the rocks, and could also fly into the air to get a bird's eye view of the environment (the same opportunities as the information group in phase 1). However, this group could not access any of the embedded information. The students were given no time limit, but it was suggested that they spend less than 20 minutes. This group was used as a check to ensure that any performance differences were due to the coupling between the information and the virtual environment, and not just because the novelty of the VE experience motivated higher learning during lectures.

Before their VE sessions, each of the students in the information and habitat groups received both written and verbal instructions on the use of the system as appropriate. They also signed an informed consent form and completed a background questionnaire that inquired about their age, gender, handedness, college major, computer usage and experience, and VE experience. Students completed their use of the VE before the class lecture took place. They were told that they were simply trying a new computer system that might be used in later classes, and were naive regarding the purposes of the experimenters.

In the class period after the lecture on this material (5 days later), a test was given to all students in the class. The test covered material relating to the philosophy of zoo exhibit design and specific information about the design of the gorilla habitat at Zoo Atlanta. At the conclusion of the test, students were told about the nature and purposes of our evaluation.

	Information	Habitat	Control
initial # students	9	6	9
final # students	8	3	5
VE usage	complete	habitat only	none

Table 1. Summary of experimental groups

Test subsets	# questions
V: VE-only	5
L: Lecture-only	12
B: Both	9

Table 2. Evaluation test summary

The test consisted of 26 questions, 24 fill in the blank and 2 multiple choice. Also, there were three subsets of questions relative to the information presented: 5 of the questions could only be answered from material presented in the VE, 12 of the questions could only be answered from material presented in the lecture, and 9 of the questions concerned material that was given in both the VE and the lecture (Table 2). Analysis of scores on these three subsets could help to show which of our two hypotheses, if any, was correct. If the information group scored higher on the questions relating only to the VE or to both the VE and lecture, the first hypothesis would be supported (students learned and understood more material because they encountered that information in the VE). If the information group scored higher on the questions that came only from the lecture, then our second hypothesis would be supported (students were more motivated to learn from the lecture because of the VE experience).

Unfortunately, the size of the three groups was reduced due to imperfect class attendance. Some students did not attend the lecture class, and others were absent on the day of the test. This left the control group with 5 students, the habitat group with 3 students, and the information group with 8 students.

Results

The results of the evaluation support both hypotheses given above, and are summarized in Figure 5. Due to the small sample size and a few outlying scores, the differences were not statistically significant. However, they do reveal a positive trend, and we are confident that in a larger test our claims could be statistically proven.

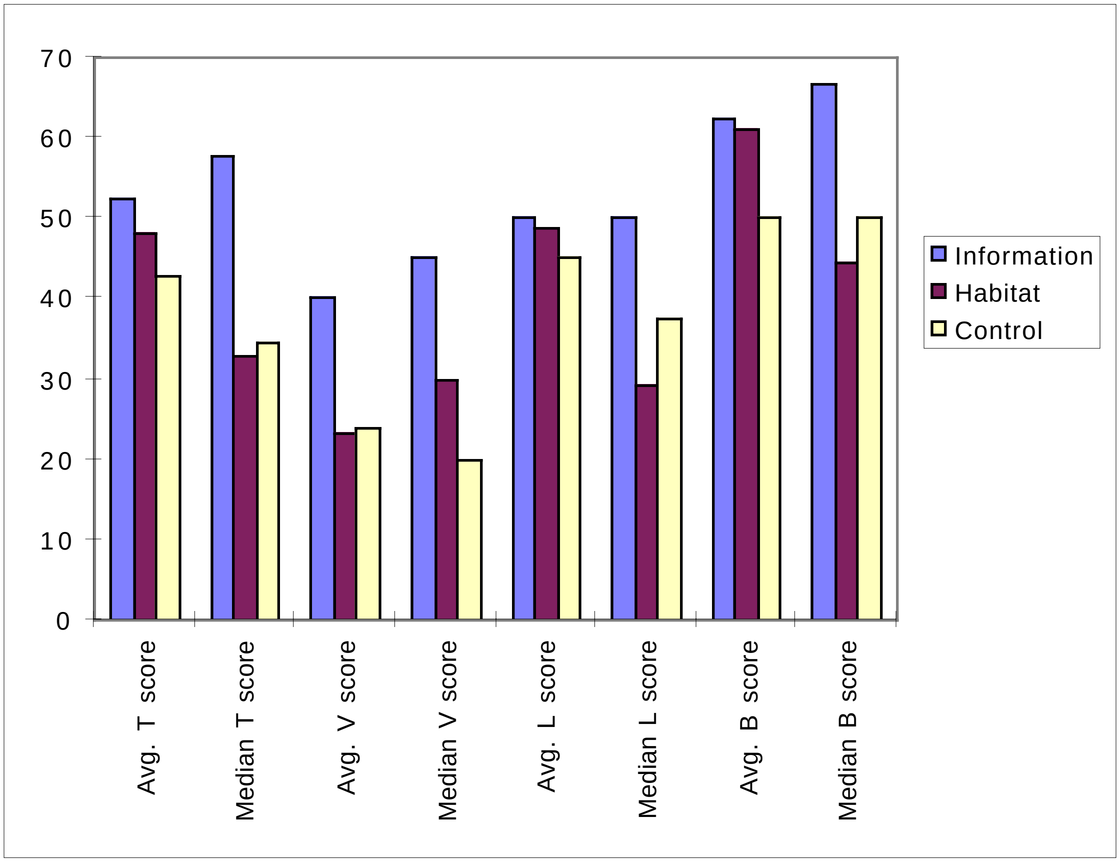

Figure 5. Summary of Test Score Results (T: Total score, V: Questions relating to information found only in VE, L: Questions relating to information found only in lecture, B: Questions relating to information found in both the VE and lecture)

The figure shows both the average and median scores for the entire test (T score) and for the three subsets of questions that we mentioned previously. The information group had the highest average and median score for each of the subsets and the complete test.

We also collected and analyzed some peripheral data not directly related to performance on the test. First, we wondered whether the time between the use of the VE system and test would have an effect on performance. This time lag ranged from 1 to 8 days, but we found no correlation between lag and test score. Second, we collected data on the number of pieces of information of each type (audio, text, and image) viewed by each member of the information group. These students viewed from 62 to 85 percent of the total information in the virtual exhibit. However, the information visited did not correlate with test score or scores on any of the three subsets of questions.

This was also true of the amount of time spent in the virtual exhibit, in both the exploring and information-gathering phases.

Finally, we analyzed the results of a normal in-class test the week following our evaluation, and found that the information group had a slightly higher score than the other two groups. One might conclude that the differences in test scores from our evaluation, then, were due to the fact that the information group as a whole was more intelligent! We calculated the ratio of the experimental test scores to the normal class test scores to check this conclusion. This ratio, on average, was higher for the information group (0.76), then for the habitat (0.66) and control (0.66) groups, confirming that the information-rich VE did indeed enhance the learning experience.

Analysis and Discussion

Taking a closer look at the summary statistics from our evalaation, some interesting trends and observations arise. The trends are stronger when we look at the median scores of the groups, because this removes some of the effect of outliers on the average score. However, the trends are visible when average is considered as well.

First, we note that both of our hypotheses were supported. Recall that the first hypothesis was that the VE system paired with a lecture on the same material would provide greater learning and understanding than a lecture alone. Here, it is important that we look at overall performance, as well as the performance of students on those questions which were answered in both the virtual exhibit and the lecture (B questions). Students in the control group had received the information necessary to answer these questions, but the information group students scored higher on these questions. This supports our assertion that the VE produced some absolute educational benefits.

One might say that these benefits occur only because the information group received this information twice (more time spent on task), and that the method of presentation was not the important factor. Even if this is true, it does not negate the hypothesis, which was that an information-rich VE is an effective method of education when paired with traditional classroom teaching. On an intuitive level, it is clear that the VE is not simply a second method of presentation equivalent to another lecture. It is better at exploiting associations between spatial and abstract information, and adds a strong experiential component to the educational process. The lecture excels at explaining concepts in detail and providing a strong theoretical foundation. We have shown that together these techniques are more effective than the traditional technique by itself.

Our second hypothesis was also supported. One of the most interesting facts in Figure 5 is that the information group scored higher than the other groups on questions which were answered only in the lecture (L score). The difference in averages is not very large, basically due to one high score in both the habitat and control groups, but the difference in the median grade is impressive. This strongly suggests that students who used the information-rich VE were able to draw on that experience during the lecture, creating more and richer mental associations that helped them on the test.

For example, a series of three questions (all in the lecture-only subset) asked students about the tendency of gorillas who had previously been in research labs to remain near the holding building in the back of the exhibit. Students were asked to describe this tendency, why it was a problem from a design point of view, and what changes had been made by the keepers to combat the problem. A student from the information group, listening to this information in the lecture, could visualize the position of the holding building at the back of the exhibit and immediately realize that there was no sight line between the visitor viewing points and this position. When the changes were described (throwing food from the top of the visitors center instead of near the holding building), this student could also visualize the center and what effect the change would have (gorillas would venture down the hill toward the visitors to get food). Thus, the information group had the opportunity to create mental associations that contained both spatial and symbolic information in a tightly coupled manner. To a limited extent, this is also true of the habitat group, but they were not given the names of the two buildings nor the locations of the visitor viewpoints, so it would be more difficult for them to create the same associations.

We also see from the results that our original assumption - that the information-rich VE alone is not very effective as a teacher - was generally correct. On the questions that could only be answered from the information in the VE, the information group had higher scores, but only about 15 points higher than students who had not received the information at all! The information group answered only about 2 of 5 questions correctly on average from this subset. This is lower than their overall average, and is the lowest of the averages of the three subsets of questions.

However, it is interesting to note that one of the students in the information group who did not attend the lecture still scored 30 percent on the test, which placed him higher than one student in the habitat group and two students in the control group who attended the lecture. This student obviously learned some material solely through his experience using the information-rich VE.

CONCLUSIONS AND FUTURE WORK

In this paper, we have presented evidence that information-rich virtual environments can produce increases in learning and motivation when combined with normal classroom teaching. Students not only performed better in an evaluation on an absolute basis, but also showed better retention of related material given in lecture form. This validates our claim that a VE which includes both spatial and abstract information allows learners to better understand the relationships between the two types of data. Experience combined with background knowledge is a richer form of learning that we should take advantage of. This type of learning can be achieved in other ways, such as laboratories and field trips, but these cannot offer the range of experience that can be produced in a virtual environment, and VEs also allow students to enter environments which are inaccessible because of their scale (a collection of molecules), their distance or cost (the coral reefs off New Zealand), or the danger involved (the inside of the gorilla habitat).

Moreover, students' assessment of the VE system was overwhelmingly positive. Students were asked about their comprehension of the space and the usefulness of the embedded information following their VE session. All of them were able to understand the layout and design of the virtual zoo exhibit through their virtual experience, and they also found it easy to access the embedded information. The students had been to the actual exhibit before using the virtual one, and most commented that they had a better sense of the space after VE usage, since they could travel to viewpoints that are not possible in the actual habitat. Some of the positive response is undoubtedly due to the novelty of virtual reality, but students' test performance revealed the practical effectiveness of the system as well.

The virtual zoo exhibit has another component which allows users to make modifications to the design of the habitat by moving trees, changing the terrain, repositioning visitor viewpoints, and so on. The class that participated in our evaluation was later part of a usability study involving these design tools. Each project team created a unique design in which they applied their knowledge of the philosophy of habitat design. Thus, information gathering and constructive learning are integrated into a single system, and students can compare their designs and design rationales to the originals. We will report on this study in another article, but some preliminary information is available via the Internet[2].

---

[2] http://www.cc.gatech.edu/gvu/people/Phd/Doug.Bowman/arch4751/

In the future, we plan to continue our study of information-rich virtual environments and their application to both education and general information gathering tasks. New information types and embedding techniques will be needed to create a tighter coupling between information and environment. We will also be studying ways to integrate experiential learning, such as the virtual physics experiments discussed earlier, into information-rich VEs. Finally, we will continue our research into effective and efficient interaction techniques and user interfaces for immersive VEs.

ACKNOWLEDGMENTS

The authors wish to thank the students in Dr. Wineman's Psychology of Environmental Design course for their participation in our study. We also acknowledge the help and support of other members of the Virtual Reality Gorilla team: Don Allison, Brian Wills, Kyle Burks, Kristen Lukas, and Lori Perkins, and the support of Zoo Atlanta and its Director and CEO, Dr. Terry Maple.

REFERENCES

Allison, D., Hodges, L., & Wineman, J. (1997). Gorillas in the Bits. In *Proceedings of the Virtual Reality Annual International Symposium*, 69-76.

Allison, D., Wills, B., Bowman, D., Wineman, J., & Hodges, L. (1997). The Virtual Reality Gorilla Exhibit. *IEEE Computer Graphics & Applications*, *17*(6), 30-38.

Angus, I. & Sowizral, H. (1995). Embedding the 2D Interaction Metaphor in a Real 3D Virtual Environment. In *Proceedings of SPIE, Stereoscopic Displays and Virtual Reality Systems*, *2409*, 282-293.

Benford, S., Snowdon, D., & Mariani, J. (1995). Populated Information Terrains: First Steps. In R. Earnshaw, J. Vince, & H. Jones (Eds.), *Virtual Reality Applications* (pp. 27-39). Academic Press.

Bergman, L., Richardson, J., Richardson, D., & Brooks, F. (1993). VIEW-An Exploratory Molecular Visualization System with User-Definable Interaction Sequences. In *Computer Graphics* (Proceedings of SIGGRAPH), 117-126.

Bolter, J., Hodges, L., Meyer. T., & Nichols, A. (1995). Integrating Perceptual and Symbolic Information in VR. *IEEE Computer Graphics and Applications*, *15*(4), 8-11.

Bowman, D., Koller, D., & Hodges, L. (1997). Travel in Immersive Virtual Environments: An Evaluation of Viewpoint Motion Control Techniques. In *Proceedings of the Virtual Reality Annual International Symposium*, 45-52.

Bowman, D. & Hodges, L. (1997) An Evaluation of Techniques for Grabbing and Manipulating Remote Objects in Immersive Virtual Environments. In *Proceedings of the Symposium on Interactive 3D Graphics*, 27-30.

Bowman, D., Hodges, L., & Bolter, J. (1998, to appear). The Virtual Venue: User-Computer Interaction in an Information-Rich Virtual Environment. To appear in *Presence: Teleoperators and Virtual Environments*.

Brelsford, J. (1993). Physics Education in a Virtual Environment. In *Proceedings of the Human Factors and Ergonomics Society 37th Annual Meeting*, 1286-1290.

Brooks, F. et al. (1992). Final Technical Report: Walkthrough Project. Report to National Science Foundation.

Bryson, S. & Levit, C. (1992). The Virtual Wind Tunnel. *IEEE Computer Graphics and Applications*, *12*(4), 25-34.

Coe, J. (1985). Design and Perception: Making the Zoo Experience Real. *Zoo Biology*, *4*, 197-208.

Coe, J. & Maple, T. (1987). In Search of Eden - A Brief History of Great Ape Exhibits. In *AAZPA Anuual Proceedings.*

Dede, C., Salzman, M., & Loftin, R. (1996). ScienceSpace: Virtual Realities for Learning Complex and Abstract Scientific Concepts. In *Proceedings of the Virtual Reality Annual International Symposium*, 246-252.

Durlach, N. & Mavor, A., eds. (1995). Virtual Reality: Scientific and Technological Challenges. National Academy Press.

Fairchild, K., Poltrock, S., & Furnas, G. (1988). Semnet: Three-Dimensional Graphic Representations of Large Knowledge Bases. In *Cognitive Science and its Applications for Human-Computer Interaction*. Lawrence Erlbaum Associates.

Fairchild, K. (1993). Information Management Using Virtual Reality-Based Visualizations. In A. Wexelblat (Ed.), *Virtual Reality Applications and Explorations* (pp. 45-74). Academic Press Professional.

Hinckley, K., Pausch, R., Proffitt, D., Patten, J., & Kassell, N. (1997). Cooperative Bimanual Action. In *Proceedings of CHI*, 27-34.

Hodges, L., Rothbaum, B., Kooper, R., Opdyke, D., Meyer, T., North, M., de Graff, J., & Williford, J. (1995). Virtual Environments for Treating the Fear of Heights. *IEEE Computer*, *28*(7), 27-34.

Kessler, D., Kooper, R., Verlinden, J, & Hodges, L. (1994). The Simple Virtual Environment Library Version 1.4 User's Guide. Graphics, Visualization, and Usability Center Technical Report GIT-GVU-94-34.

Poupyrev, I., Billinghurst, M., Weghorst, S., & Ichikawa, T. (1996). The Go-Go Interaction Technique: Non-linear Mapping for Direct Manipulation in VR. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, 79-80.

Risch, J., May, R., Thomas, J., & Dowson, S. (1996). Interactive Information Visualization for Exploratory Intelligence Data Analysis. In *Proceedings of the Virtual Reality Annual International Symposium*, 230-238.

Robertson, G., Card, S., & Mackinlay, J. (1993). Information Visualization Using 3D Interactive Animation. *Communications of the ACM*, *36*(4), 57-71.

Song, D. & Norman, M. (1993). Cosmic explorer: A Virtual Reality Environment for Exploring Cosmic Data. In *Proceedings of the IEEE Symposium on Research Frontiers in Virtual Reality*.

Tate, D., Sibert, L., & King, T. (1997). Using Virtual Environments to Train Firefighters. *IEEE Computer Graphics & Applications*, *17*(6), 23-29.

Wickens, C. (1992). Virtual Reality and Education. In *Proceedings of IEEE International Conference on Systems, Man, and Cybernetics*, 842-847.

Wineman, J., Hodges, L., Allison, D., & Wills, B. (1997). Virtual Gorilla Project: Gorillas in the Bits. *Proceedings of the Ed-Media/Ed-Telecom 97 World Conference*.