

Using Virtual Environments in the Teaching of Computer Graphics

Doug A. Bowman, Balaprasuna
Chennupati, Matthew Gracey
Department of Computer
Science, Virginia Tech, USA
{bowman, prasuna, mgracey}@vt.edu

Marcio S. Pinho
Instituto de Informática, UFRGS
and Faculdade de Informática,
PUCRS, Brazil
pinho@inf.pucrs.br

Kristin J. Wheeler
Department of Mathematics and
Computer Science, Grambling
State University, USA
krwheel2@vt.edu

Abstract

Education has long been touted as an appropriate application area for immersive virtual environments (VEs), but few immersive applications have actually been used in the classroom, and even fewer have been compared empirically with other teaching methods. This paper presents VENTS, a novel immersive VE application intended to teach the concept of the three-dimensional (3D) normalizing transformation in an undergraduate computer graphics class. VENTS was developed based on key principles for the use of VEs in education, systematically evaluated for usability, and refined based on the results of this evaluation. Students in a university course used VENTS, and their learning was compared to that of other students who either attended a lecture on the topic or used a 3D desktop application covering the same material. The results of pre- and post-tests showed a larger percent increase in test score for the VE group than the desktop or lecture groups, although these differences are not statistically significant.

1. Introduction

For many years, researchers in the area of immersive virtual environments (VEs) have pointed to *education* as a key application area. In his seminal 1965 article “The Ultimate Display,” Ivan Sutherland stated that if a three-dimensional (3D) graphics display could be controlled by head and body movements (an immersive VE), then the display could be “a looking glass into a mathematical wonderland” [17]. In other words, Sutherland realized that visualization of data in an immersive environment could enhance understanding and learning. In the National Research Council’s report on virtual reality [6], the use of VEs in education was also considered a crucial area for future research. The report envisioned the use of VEs to teach high school students on a worldwide network. It stated that VE technology could increase the student’s range of experience, provide macro contexts for situated cognition, provide micro worlds in which students could exercise their skills and knowledge, and expand the student’s peer group for collaborative learning [6].

Despite these predictions, the use of immersive VEs in educational settings has been limited to research prototypes and small-scale deployment in schools. Some of the impediments to larger-scale integration into the curriculum include:

- **Usability and human factors:** Many interactive VE applications are difficult to use because of 3D interaction issues [8] and human factors issues related to the displays and devices used.
- **Lack of evidence of improved learning or retention:** It has proven difficult to obtain statistically significant results that demonstrate the effectiveness of VEs in improving learning or retention over current teaching methods.
- **Cost:** Immersive VE systems have been too expensive to purchase for classroom use.

In our research, we have been developing educational applications of immersive VEs that address all of these issues. First, we carefully design the user interface (UI) to our applications based on the state of the art in 3D interaction research, and evaluate our applications to ensure their usability. Second, we perform empirical studies of the effectiveness of these applications in the classroom. Finally, we put together low-cost, portable immersive VE systems designed for classroom use.

In addition, our experience in this area has led us to postulate four principles for designing educational VEs:

1. **Interactivity:** Educational VEs should provide students many opportunities to interact with and explore the environment, rather than simply offering a static visualization.
2. **Complementarity:** VEs should be used to complement, not replace, other methods of teaching and learning in the classroom.
3. **Information-Richness:** Educational VEs should contain many different types of information (geometric, textual, auditory, concrete, abstract, etc.) and provide links between the various types to help students build associations and to support different learning styles.
4. **Augmentation:** Educational VEs will be most effective when they portray information that is not normally visible or accessible in the physical world.

In this paper, we present the design and evaluation of an immersive VE application intended to teach a difficult concept in a computer graphics course. The *3D normalizing transformation* [7] is a viewing transformation that maps an arbitrary 3D perspective view into a canonical 3D view in order to simplify the processes of clipping and projection. Understanding this concept requires students to visualize 3D structures moving through 3D space, and students often have difficulty absorbing or comprehending this information. Our application allows students to step through the phases of the transformation at their own pace, within an immersive VE, viewing the transformation from any location. We call this application the Virtual Environment Normalizing Transformation System (VENTS).

In the following section, we review related literature on the use of immersive VEs in education. We continue with a detailed description of the normalizing transformation and typical methods of teaching this content. Next, we describe VENTS and a usability evaluation that helped us to refine the implementation. Finally, we describe our empirical assessment of student learning comparing VENTS to other teaching methods, and present issues to be addressed by future work.

2. Related work

Wickens [18] 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 his 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. We argue that a distinction needs to be made between cognitive load from task-related activities and system-related activities. In educational environments, learning activities (task-related) should require effort and choice; however, system-related activities, such as selecting an object or finding a menu item should require little cognitive processing. One does not want users to be distracted from learning because they cannot figure out how to use the interface.

The best example of an immersive VE system used for educational purposes is the ScienceSpace project [5, 15]. This project developed three virtual worlds designed to teach students concepts in physics, such as gravitation and electromagnetism. Controlled experiments using this

system with students have shown a great deal of potential. Our own work on the Virtual Habitat system [2, 3] also showed a trend towards increased learning in an immersive, information-rich VE. It is still not clear, however, whether immersive VEs have educational advantages over non-immersive environments, such as VRML worlds presented on a standard monitor.

Few immersive VE applications have been deployed in schools, but a notable exception is the work of Johnson and his colleagues [10, 14]. This research has focused mainly on elementary school education and collaborative learning, with a great deal of success.

3. Educational content

VENTS is designed to teach students of computer graphics the concepts and steps involved in the 3D normalizing transformation.

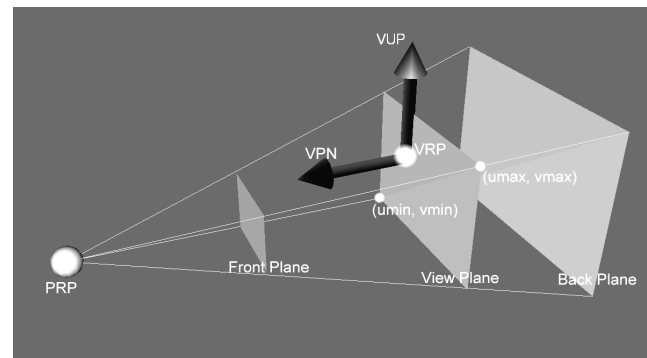


Figure 1. Virtual camera model used to define a 3D perspective view

3.1 Normalizing transformation

A 3D perspective view can be described by defining various properties of a virtual camera within a world coordinate system. Foley et al. [7] use the following values to define these properties (see figure 1):

- **View Reference Point (VRP):** A point in world coordinates defining the origin of the camera coordinate system.
- **View Plane Normal (VPN):** A vector in world coordinates normal to the viewplane. The VPN forms one axis of the camera coordinate system.
- **View Up Vector (VUP):** A vector in world coordinates that specifies the orientation of the camera (i.e. that defines the other two axes of the camera coordinate system).
- **Projection Reference Point (PRP):** A point in camera coordinates that defines the position of the virtual camera.
- **Window extents (u_{min} , u_{max} , v_{min} , v_{max}):** Minimum and maximum values in camera coordinates that

define the boundaries of the viewing window on the viewplane.

- **Front and back clipping planes (F, B):** The planes $n = F$ and $n = -B$ in camera coordinates that define the front and back sides of the view volume.

The normalizing transformation is used to map an arbitrary 3D viewing situation defined in this manner into a canonical 3D viewing situation, in order to simplify the processes of clipping and projection. Five sub-transformations, or steps, are involved in the normalizing transformation for perspective views (paraphrased from [7]):

1. Translate such that the VRP is at the origin of the world coordinate system.
2. Rotate such that the u , v , and n axes of the camera coordinate system are aligned with the x , y , and z axes of the world coordinate system, respectively.
3. Translate such that the PRP is at the origin.
4. Shear such that the centerline of the view volume becomes the z axis.
5. Scale such that the sides of the view volume become the planes:

$$x = z, x = -z, y = z, y = -z, z = -z_{min}, \text{ and } z = -l.$$

Each step has an associated 4x4 transformation matrix, and these matrices are multiplied together to produce the matrix describing the entire transformation.

Typically, students have difficulty understanding this process because of its complexity and the necessity of imagining these abstract 3D elements moving through 3D space. Specifically, many students do not understand the overall configuration of the viewing volume; many have difficulty comprehending that the transformation is actually applied to the vertices of the objects in the scene; and many cannot see why the shear in step 4 is necessary.

3.2 Teaching approaches

Because of the complex and abstract nature of the normalizing transformation, a number of approaches have been used in teaching this concept to students of computer graphics. Probably the most widely used approach is to display static 2D images (themselves perspective projections!) illustrating the state of the view volume after each step of the transformation. This approach has the problem of student misinterpretation of the true 3D structure portrayed in the 2D image.

To overcome this problem, Foley et al. [7] suggest building the coordinate systems, vectors, points, and objects using Tinker Toys or some other physical 3D representation. This can be useful in understanding the 3D nature of the problem, and in illustrating some of the steps of the transformation, such as the rotation step. However,

it would be very difficult to convey a shear or scale transformation effectively using this type of physical prop.

This suggests that an interactive 3D graphics application might be effective in helping students understand the normalizing transformation, since all the steps of the transformation could be shown correctly, in 3D. An excellent example of this approach is the Exploratories project [16], which includes 3D Java applets for many concepts, including perspective viewing parameters and the normalizing transformation. A possible shortcoming of this approach, however, is the indirect manipulation of the view using a mouse rather than natural head and body motions.

4. Immersive application

We designed and implemented VENTS as an immersive tutorial to be used in learning about the 3D normalizing transformation. Our development and usage of VENTS follows the four principles outlined in the introduction: the system is highly interactive, is designed to be used as a supplement to classroom teaching, contains geometric, textual, and mathematical information, and portrays an abstract concept not normally visible.

4.1 Hardware and software systems

VENTS runs on a Windows 2000-based personal computer. The user wears a head-mounted display (HMD) For our experiment, we used the Virtual Research V8, which offers 640x480 resolution and a 60-degree diagonal field of view. The user's head is tracked by an Intersense IS-900 VET tracking system, and the same system is used to track a stylus and tablet held by the user (see figure 2). The stylus has two physical buttons. Note that this is not the low-cost, portable system for classroom use. That system consists of a Daeyang i-Visor HMD and an Intersense Intertrax2 3DOF head tracker, plus a 12-button keyboard for input. VENTS was implemented within the Simple Virtual Environment (SVE) library [11].



Figure 2. Physical devices used in VENTS

4.2 Environment

The user is placed within a large virtual cube. Each face of the cube has a different texture map, so that users can easily orient themselves. Objects in the environment include a world coordinate system, a small cube, and a view frustum (figure 3a). The world coordinate system consists of three white axes that have a fixed position. The cube is a simple object that represents the “scene” to be transformed by the normalizing transformation.

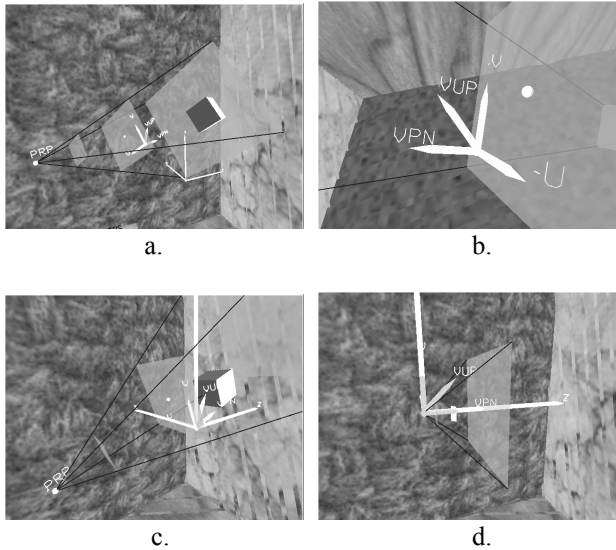


Figure 3. Views of the VENTS environment: a) initial configuration, b) close-up view of camera coordinate system, c) view of the environment during the rotation step, d) view of the environment following the scale step

The view frustum is a pyramidal volume showing the current viewing parameters. The frustum includes the PRP, VUP, VPN, camera coordinate system, view window, and front and back clipping planes. The PRP is represented as a white sphere, and the VUP, VPN, and camera coordinate system as green vectors (figure 3b). The labels on the PRP, VUP, VPN, and camera coordinate system rotate to face the user so they can be read from any location. The view window and front and back clipping planes are slightly transparent, so that the rest of the scene can be viewed through these planes.

VENTS displays user-controllable animations of the viewing transformation steps. Figures 3c and 3d show examples of the state of the system during the rotation step and at the end of the entire transformation, following the scale step.

4.3 User Interface

Users travel through the environment using a pointing technique [12], by pointing the stylus in the desired direction of travel and holding down one of the stylus buttons. This technique has been shown to perform well in situations when most travel is relative to other objects in the environment [4].

The physical stylus and tablet are used to create a two-dimensional (2D) UI within the 3D environment, using the “pen & tablet” metaphor [1, 2]. The user sees a virtual pen and a virtual tablet at the same locations as the physical devices, and a virtual interface is shown on the surface of the virtual tablet (figure 4). This interaction style has numerous well-known advantages, including users’ prior knowledge of 2D interfaces, the constraint provided by the physical surface of the tablet, the use of proprioception [13] and two hands [9] to improve performance, and the ability to put the interface away when it is not needed.

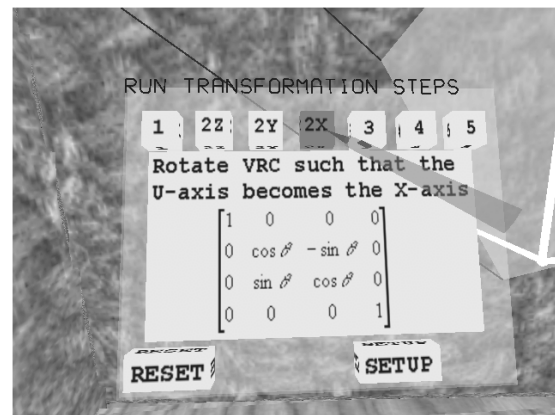


Figure 4. 2D interface on the surface of the virtual tablet; run mode is shown

The tablet interface in VENTS has a row of virtual buttons across the top of the tablet, a large information-display area in the center, and a reset button and mode button at the bottom. The mode button is used to toggle between “run” and “setup” modes.

The run mode allows the user to animate each of the steps of the normalizing transformation. When the user touches one of the buttons with the stylus, a description of that step and the associated transformation matrix are displayed in the information area (figure 4). When the user “clicks” the virtual button (by clicking one of the stylus buttons while the stylus is touching the virtual button), an animation of the step is shown in the environment.

Setup mode allows the user to change the initial viewing parameters (figure 5). The user can set the position of the PRP, define the direction of projection

(DOP – the vector from the PRP to the center of the viewing window), change the height or width of the viewing window, or change the distance from the PRP to the viewing window. To set the PRP position, the user clicks the PRP button, and then clicks the stylus button in 3D space at the desired PRP position. To change the DOP, the user clicks the DOP button, and then rotates the stylus to match the desired DOP vector. The other options simply increase or decrease the appropriate parameter by a fixed amount.

The reset button on the tablet restores all parameters to their initial values and restores the scene to the state before step 1 of the transformation.

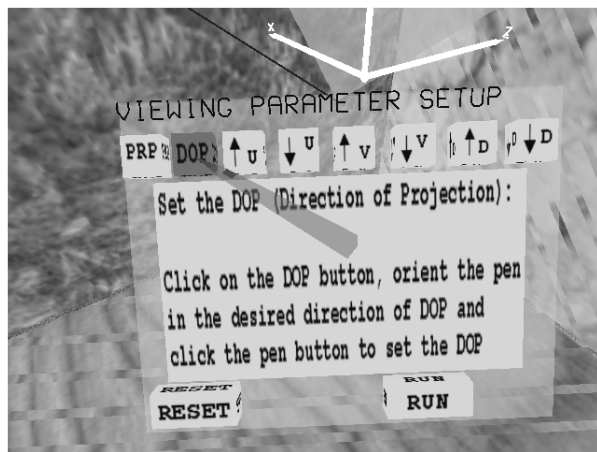


Figure 5. Setup mode interface

5. Usability evaluation

We ran a usability evaluation on our initial implementation of VENTS to determine any major usability problems and refine our UI.

5.1 Evaluation design

Our usability goals for VENTS were for the system to be understandable and usable by novices, since our target users were students who would likely use the system only once. Therefore, we designed a usability evaluation that tested users' initial performance (time to complete task and number of errors) on a variety of common tasks within the system. Users were given no initial instruction, except in how to hold the physical devices and press the physical stylus buttons. After this initial phase, we allowed users to explore the system for a few minutes, and then measured performance again on tasks similar to the first set.

Ten users, recruited from computer science classes, participated in the study. Each user completed a demographic questionnaire, completed the tasks within

VENTS, and completed a questionnaire regarding their perceived level of satisfaction, ease of use, etc.

5.2 Results

Performance on the initial set of tasks was not as good as we hoped. Task times were slightly higher than expected, and five of the ten users made at least one error on these tasks. However, after the exploration phase, all users completed the second set of tasks quickly and with no errors. The questionnaire results showed that users enjoyed using the system and that the system's features were rated highly for their ease of use (an overall average response of 4.5 on a 5-point scale).

5.3 System modifications

The usability study suggested a number of ways that we could refine our implementation in order to improve users' initial performance with the system. The biggest problem was misunderstanding of the run/setup mode button. We improved the textual information that is displayed when the user touches this button, and clearly labeled the top of the tablet interface to indicate which mode is currently being used. We also changed the word "stylus" to the more common word "pen" in all of the help text and added the reset button in case users became confused in setup mode.

6. Summative evaluation of learning

In order to assess the benefits provided by VENTS, we designed and ran a summative evaluation comparing VENTS to two other teaching approaches: an interactive desktop application and a traditional lecture.

6.1 Method

The experiment was performed within the context of an undergraduate course in computer graphics at Virginia Tech. During a regular class meeting, all students attended a lecture explaining the concept of the 3D normalizing transformation and the details of the five steps. This lecture involved verbal instruction by the professor, PowerPoint slides with text and static 2D images, and additional images drawn on a chalkboard.

Participation in the experiment was voluntary. Students volunteering for the experiment were assigned to one of four groups: a lecture group (L), a desktop group (D), a "fixed VE" group (VF), or a "changeable VE" group (VC).

The lecture condition consisted of an additional lecture (besides the earlier one attended by all students). This 15-minute lecture covered the steps of the normalizing

transformation using a specific example (applying the transformations to the three vertices of a triangle). The only materials used were verbal descriptions and PowerPoint slides.

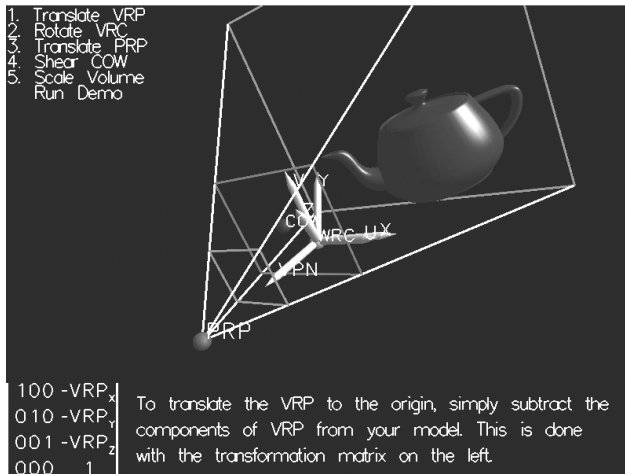


Figure 6. Desktop tutorial on the 3D normalizing transformation

In the desktop condition, students used an interactive 3D graphics application (shown in figure 6) on a typical desktop PC. It contains the same information as VENTS: a graphical representation of a view frustum, coordinate systems, and an object to be rendered, along with labels for important points and vectors. The bottom of the screen contains an information area describing the current transformation step and displaying the associated matrix. In the upper left corner of the screen is a menu with which the user can select a specific transformation step, or choose to see all the steps consecutively (“demo”). When a step or the demo is selected, the system displays an animation showing the effects of the step on the view frustum and object. Finally, the user can obtain different views of the scene by clicking the left mouse button and dragging to rotate the scene, or by clicking the right mouse button and dragging to zoom in or out. The desktop application does not allow the user to change the viewing parameters. During the experiment this application was run in full-screen mode.

The fixed VE group used VENTS, but was not allowed to change the initial viewing parameters using setup mode. The changeable VE group used VENTS with no restrictions on feature use.

6.1.1 Subjects. 25 students participated in the study. Because of various constraints, we had to restrict the number of students in the VE and desktop groups. One constraint was that the experiment took place in the context of an actual course, during the lecture period for the course. Only two lecture periods (75 minutes each) could be spared for the experiment. One period was used

for groups L and D, and the other for groups VF and VC. Three PCs were available for group D, and we were able to schedule two sets of subjects for a total of six. Only one VE system was available, and each student spent 10 minutes using the VE. The schedule allowed us to assign only three students to group VF and three to group VC. This left 13 students in group L. With the exception of three computer science graduate students, all subjects were undergraduate computer science or computer engineering majors.

6.1.2 Experimental design. The experiment used a between-subjects design with instruction method (L, D, VF, or VC) the independent variable. The dependent variable was the result of a written test of knowledge about the normalizing transformation. We administered a pre-test to all groups to assess their level of comprehension and retention of the material presented in the earlier in-class lecture. After the instructional component all groups were given a post-test. All groups received the same pre- and post-tests. Both tests consisted of four questions – two multiple-choice and two open-ended questions requiring a written or sketched answer. One question (“What is the purpose of the normalizing transformation?”) was the same on both tests. The other three questions on the post-test asked about the same concepts as the other three questions on the pre-test, but using different wording and/or a different question style. Tests were scored on a scale from zero to four, with each question worth one point, and ½-point partial credit available on the open-ended questions.

6.1.3 Procedure. Subjects in all groups first completed an informed consent form and the pre-test. Subjects in groups D, VF, and VC also received written instructions on the use of the desktop or VE system, respectively. Next, the subjects either listened to the lecture (15 minutes) or used the desktop or VE system (10 minutes). Although time on task is not equal for all groups, we considered this acceptable since the hypothesis was that the VE (even with less time on task) would improve learning. Finally, all groups completed the post-test.

6.2 Results

We considered two different ways of measuring the learning gains from the pre-test to the post-test: absolute difference and percent increase. The absolute difference was found by subtracting the pre-test score from the post-test score. The percent increase is the absolute difference divided by the pre-test score. Figure 7 shows the average results for these two metrics for each of the four experimental groups, plus a fifth group combining the results from both VE groups.

As the figure shows, the desktop group exhibited the largest absolute difference, followed by the fixed VE group. The fixed VE group had the largest average percent increase. Because of the unavoidably small sample sizes (especially for the VE groups) and high variability, these differences are not statistically significant based on analyses of variance (ANOVA).

The figure also shows that group VC had an average absolute difference of zero. Because of this, the combined group of all VE users (group VB) had the smallest average gain in test scores. However, group VB still had a higher percent increase in score than did group L. We discuss the performance of group VC further in the following section.

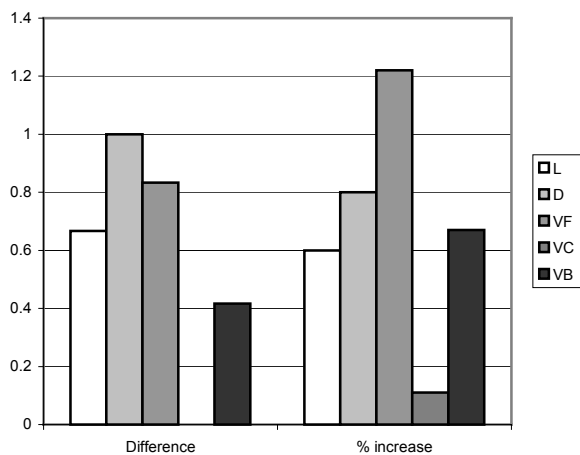


Figure 7. Average difference and percent increase in pre- and post-test scores for the experimental groups (note: VB is a combination of both VE groups)

6.3 Discussion

The results of the experiment were encouraging, although statistical significance was not achieved. We showed that even with a smaller amount of time on task, users of VENTS could learn a great deal more about the normalizing transformation than they knew before using the system. The desktop application was also useful in helping students understand the material, leading us to the tentative conclusion that for learning this particular concept, interactive 3D representations provide a real benefit.

Several important issues were raised by the experiment. First, why was the performance of the VC group so poor? Based on our observations of these users, two factors seem to have contributed to this result. First, these users spent more time working with the setup interface than in viewing and understanding the animated transformation steps. This group was explicitly instructed

to use the setup mode, and this may have caused them to focus on it more than they would have under normal circumstances. Second, some of the interactions were still difficult to understand. In particular, when setting the direction of projection, the view volume rotated about the VRP, rather than the PRP as users expected. We still feel that the ability to change the parameters could lead to greater understanding, since students should be able to generalize what they see into a more complete and correct mental model, but students will need more time and an even more transparent interface in order to do this.

We also noticed some interesting behaviors by members of groups D, VF, and VC. Desktop users tended to exhaust the possibilities of the application in two or three minutes, and few of these users modified their viewpoint significantly during the transformation steps. Most VE users, on the other hand, seemed to be engaged with the system for the entire ten-minute session, perhaps due to the fact that it was easier to navigate and obtain different views of the same animation. Engagement itself is an indicator that the VE provides an environment conducive to student learning.

Subjects in group D tended to stop slightly more often in order to read the textual descriptions of the transformation steps than those in groups VF and VC. The text was more readily available on the desktop, although it also obscured some of the view. In the VE, users had to raise the tablet in order to read the text. Some users even became so comfortable with the pen & tablet interface and the location of the virtual buttons that they selected the buttons “blindly,” with the tablet out of view.

We noticed that subjects in all three interactive groups (D, VF, and VC) had visible “moments of understanding,” typified by nodding the head or making an utterance that clearly implied, “Oh, now I see.”

Another important issue involves the effective use of VENTS in a classroom context. One of the key benefits of the desktop application is that it can be downloaded and used by all students on their own computers. However, using VENTS requires specialized hardware and software. How can the benefits of the immersive environment be provided to all students? Clearly, there is not time available during normal classroom hours to give each student 10-15 minutes with the system. Students could come to use the system outside of normal class time, but it would be difficult to convince them that the benefits outweighed the costs. Our solution involves bringing the system to the classroom using the low-cost, portable setup mentioned earlier. Further, we plan to have only a few students use the immersive system directly, but also to project the immersed student’s view on a large screen. The interesting question here is whether the loss of interactivity for the majority of students will decrease the learning benefits of the system.

Finally, as we noted in the introduction, obtaining statistical evidence for the educational effectiveness of VEs has proven elusive. Aside from the fact that we need larger sample sizes (which we will achieve by running the same study in subsequent semesters), our study also suggests that more and better learning metrics are needed. Written tests are easily quantifiable, but other methods such as structured interviews, sketches, or manipulation of physical models might also help us to understand more deeply the learning gains by the student.

7. Conclusions and future work

We have presented VENTS, an immersive VE tutorial on the 3D normalizing transformation. Based on our four principles for educational VEs, VENTS provides students with an effective new way to visualize and understand this difficult 3D concept. Although the lack of statistical support for our claim that VENTS improves learning is disappointing, our experiment still shows that this approach holds significant promise.

Our current work includes further refinements to the VENTS UI based on observations during our experiment, and preparations for larger-scale use and assessment of VENTS in our computer graphics class. In addition, we are developing other educational VE applications in the areas of computer networking and building structures. We will continue to perform empirical assessments of learning using VENTS and these other applications.

Acknowledgements

This work was supported in part by National Science Foundation grant DUE-0127326. We gratefully acknowledge the work of Nathan Egge on the desktop application and the time and effort of the students in CS 4204, computer graphics. Special thanks go to Qinwei Zhu for helping to administer the experiment.

References

1. Angus, I. and Sowizral, H. Embedding the 2D Interaction Metaphor in a Real 3D Virtual Environment. In *Proceedings of SPIE, Stereoscopic Displays and Virtual Reality Systems*, 1995, pp. 282-293.
2. Bowman, D., et al., *Designing Animal Habitats Within an Immersive VE*. IEEE Computer Graphics & Applications, 1998. **18**(5): p. 9-13.
3. Bowman, D., et al., *The Educational Value of an Information-Rich Virtual Environment*. Presence: Teleoperators and Virtual Environments, 1999. **8**(3): p. 317-331.
4. Bowman, D.A., Koller, D., and Hodges, L.F. Travel in Immersive Virtual Environments: an Evaluation of Viewpoint Motion Control Techniques. In *Proceedings of the Virtual Reality Annual International Symposium*, 1997, pp. 45-52.
5. Dede, C., Salzman, M., and Loftin, R. ScienceSpace: Virtual Realities for Learning Complex and Abstract Scientific Concepts. In *Proceedings of the Virtual Reality Annual International Symposium*, 1996, pp. 246-252.
6. Durlach, N. and Mavor, A., eds. *Virtual Reality: Scientific and Technological Challenges*. 1995, National Academy Press.
7. Foley, J., et al., *Computer Graphics: Principles and Practice*. 1990, Reading, MA: Addison-Wesley.
8. Herndon, K., van Dam, A., and Gleicher, M., *The Challenges of 3D Interaction*. SIGCHI Bulletin, 1994. **26**(4): p. 36-43.
9. Hinckley, K., et al. Cooperative Bimanual Action. In *Proceedings of CHI: Human Factors in Computing Systems*, 1997, pp. 27-34.
10. Johnson, A., et al. Exploring Multiple Representations in Elementary School Science Education. In *Proceedings of IEEE Virtual Reality*, Yokohama, Japan: IEEE Computer Society Press, 2001, pp. 201-208.
11. Kessler, G., Bowman, D., and Hodges, L., *The Simple Virtual Environment Library: An Extensible Framework for Building VE Applications*. Presence: Teleoperators and Virtual Environments, 2000. **9**(2): p. 187-208.
12. Mine, M., *Virtual Environment Interaction Techniques*. 1995, UNC Chapel Hill CS Dept.
13. Mine, M., Brooks, F., and Sequin, C. Moving Objects in Space: Exploiting Proprioception in Virtual Environment Interaction. In *Proceedings of ACM SIGGRAPH*, 1997, pp. 19-26.
14. Roussou, M., et al., *Learning and Building Together in a Virtual World*. Presence: Teleoperators and Virtual Environments, 1999. **8**(3): p. 247-263.
15. Salzman, M., et al., *A Model for Understanding How Virtual Reality Aids Complex Conceptual Learning*. Presence: Teleoperators and Virtual Environments, 1999. **8**(3): p. 293-216.
16. Simpson, R., Spalter, A., and van Dam, A. Exploratories: An Educational Strategy for the 21st Century. In *Proceedings of SIGGRAPH '99 Educational Schoolhouse*, 1999.
17. Sutherland, I. The Ultimate Display. In *Proceedings of the IFIP Congress*, 1965, pp. 505-508.
18. Wickens, C. Virtual Reality and Education. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics*, 1992, pp. 842-847.