A Comparison of a Cave and a Fish Tank VR System for Counting Biological Features in a Volume

First Author Name (Blank if Blind Review)
Affiliation (Blank if Blind Review)
Address (Blank if Blind Review)
e-mail address (Blank if Blind Review)
Optional phone number (Blank if Blind Review)

Second Author Name (Blank if Blind Review)
Affiliation (Blank if Blind Review)
Address (Blank if Blind Review)
e-mail address (Blank if Blind Review)
Optional phone number (Blank if Blind Review)

ABSTRACT
Many scientists analyze 3D biological data by counting cells and other structures. Our hypothesis is that an immersive environment is better for such tasks in terms of user speed and preference than a non-immersive environment. To help test our hypothesis, we designed and implemented an interactive visualization tool for a Cave and a Fish Tank virtual environment and conducted a user study in which six users performed a counting task using our new tool in both environments. Our results showed that most users preferred the Cave and achieved better results in it.

Author Keywords
Human-centered environment, 3D user interface, volume rendering, fish tank, virtual reality, Cave, biology, confocal laser microscopy.

ACM Classification Keywords
H.1.2 [Models and Principles]: User/Machine Systems; H.5.2 [Information Interfaces and Presentation]: User Interfaces.

INTRODUCTION
To support a collaboration with biologists, we have designed and implemented a system for studying digitized volumes of multi-channel data. Our system runs in a Cave as well as a Fish Tank VR environment (see Figures 1 and 2), but we did not know which would be more effective. To help answer this question, we designed and ran a user study aimed at characterizing the advantages of each working environment.

The quantitation of cells or cellular components is critical in the study of many biological processes. The density of cells or certain components within a particular volume is often compared between control and experimental samples. Immunohistochemical techniques that make use of antibodies tagged (or labeled) with a fluorochrome or other molecule

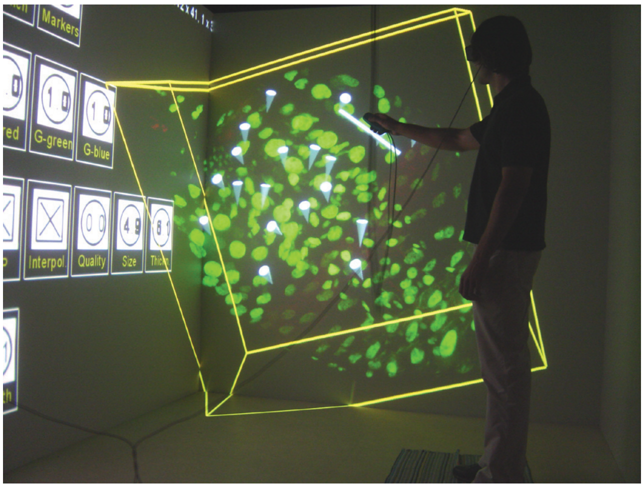

Figure 1. Viewing a cell data set in the Cave. The user places a white cone-shaped marker on a cell by reaching out to position it and pressing a button on the wand. On the left wall are several visualization controls.

Figure 2. Viewing a cell data set at the Fish Tank system. The user places a marker. We are using two monitors to provide a continuous wider field of view—the control widgets are on the right, and the data set is on the left.

that fluoresces when activated with particular wavelengths of light, allow biologists to specifically highlight (or recognize) the structures of interest within a tissue preparation. The digitized data is collected with a laser scanning confocal microscope (e.g., Leica TCS SP2 AOBS [13]) which generates a static volume data set.

Determining the density of cells in each data set requires two tasks: 1) identifying an equivalent sub-volume within each data set, and 2) accurately counting the cells in the sub-volume.

The first task requires domain knowledge to best identify and optimally register an equivalent sub-volume across samples. This might be done by using several anatomical features as landmarks for marking the boundaries in which to count cells. The second task requires uniquely identifying and tallying the cells within the volume. Here the main challenges are isolating individual cells, and due to the large number of cells, avoiding double-counting.

For this study, subjects only did the counting task. The counting task is more easily learned by most subjects, but knowing how to isolate the correct volume to count cells in is much harder. Therefore, we provided a pre-selected sub-volume for subjects to count cells in. (Nevertheless, we feel our systems that enable the user to be immersed in the 3D space would be valuable for an experienced biologist to use for identifying volumes to count cells in.) In the study, we provided a user interface that lets subjects adjust the scale of the data set, change their viewpoint, and place a marker in 3D which was useful for marking which cells they had already identified. The system automatically kept count of the number of markers that had been placed.

RELATED WORK

To our knowledge, there have not previously been any volume rendering systems which had been adapted to both CAVE [4] and Fish Tank [17] virtual reality environments.

An early volume visualization software for a CAVE, which is in many respects similar to ours, is the Crumbs system by Brady et al. [2]. It uses the same texture based volume rendering technique that we are using, but supports less functionality for volume rendering. An evaluation of the Crumbs system has been presented by Swartz et al. [14]. The Crumbs project has been discontinued, and to our knowledge the software has never been adapted to other virtual environments than CAVEs.

Another related project is the MediDesk, a volume rendering software developed by Wohlfahrter et al. [18]. It uses a tracked pen and a plexiglass prop for interactions, and runs on an immersive table display. It does not, however, run in CAVEs or Fish Tank environments.

General purpose visualization toolkits which allow volume rendering are VTK [16], Amira [1], COVISE [3], Ensight [7], Open Inventor [9, 10], and OpenGL Volumizer [15]. They can be used as a basis to develop software for CAVEs

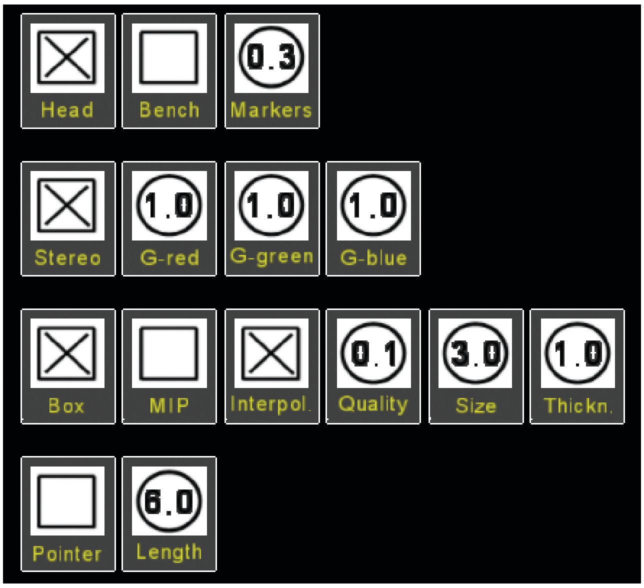

Figure 3. This user interface is shown on the left wall of the Cave and on the right monitor of the Fish Tank.

and Fish Tank VR, but they either do not support the development of custom VR user interfaces, are not easily adapted to Cave and Fish Tank, or are not extensible because their source code is proprietary.

A related user study which examined task performance in Cave and Fish Tank has been done by Demiralp et al. [6]. They compared Cave and Fish Tank in a task where users identified a feature on the surface of a potato-shaped three-dimensional object. This task required only little interaction with the system, namely rotation of the object and clicking one of two buttons, depending on the type of feature identified.

For similar reasons as Swan II et al. [8], we selected a real-world task as the basis of our formal user study. Swan II et al. did not work with volume data sets, nor did they use a Fish Tank environment for their user studies.

SYSTEM OVERVIEW

We implemented our software system in a four-sided Cave and for a dual-monitor Fish Tank virtual reality environment. Both environments use Polhemus trackers and Wanda input devices. The Wanda is a hand held input device (wand) with three buttons and a push-ball that reports relative force in two dimensions. We track the user's head and the dominant hand. Both systems provide active stereo with shutter glasses, and both are equipped with Nvidia Quadro FX 3000 graphics cards. The Cave uses the 3000G version of these cards to synchronize the four images. The resolution is 1024×768 pixels per screen in both environments. The Cave is driven by four PCs with dual Intel Xeon CPUs at 2.8 GHz and 1 GB RAM. The Fish Tank is driven by a PC with an AMD Athlon XP 1700+ CPU and 512 MB RAM.

Our software is based on an in-house library to support the rendering cluster, input devices, and volume rendering based on textured polygons [5]. We use the OpenSceneGraph API [11] for the graphical elements of the user interface.

USER INTERFACE

The central part of the UI is a set of rectangular widgets on the left wall of the Cave, or on the right screen of the Fish Tank (see Figure 3). These widgets are permanently displayed at fixed positions. We chose to display them in the plane of the screens so they are always projected at the same position, independent from where the head-tracked user is looking. Due to the unavoidable tracking lag, this gives them a much more stable feel than free-floating widgets, and even for non-tracked users they are always in the same place and therefore easily readable. This is especially valuable in the Cave, since it is meant to be used by small groups. The widgets usually do not block the user's view because typically our users do not need all four Cave walls to view the type of biological data sets they work with.

We implemented three types of widgets to trigger actions, change a state, or to set a scalar value:

- **Action buttons** trigger an action when the user moves the pointer over them and presses the left button on the wand.
- **Check boxes** change their state when they are clicked on. An "X" or a blank space indicate the state.
- **Dials** change their value when the user clicks on them and holds the button, while rotating the hand to the left for lower values, or right for higher values. Dials are the equivalent of sliders in desktop applications, but in virtual environments dials have been found easier to use (see [12]).

FUNCTIONALITY

We explained the following components of our software to the users to enable them to best take advantage of the provided functionality.

Moving the data set: The data set can be moved and rotated by pointing at it and clicking the middle button on the wand. While the button is depressed, the data set sticks on the pointer like on a fork. Repeated grabbing and moving allows to place it anywhere in the virtual world. Our wands also have a trackball which rotates the data set about its center when pushed in a direction. Unfortunately, the trackball on the Fish Tank's wand did not work during our study.

Reconstruction quality: Volume rendering suffers from low frame rates when a high reconstruction quality is desired, i.e., no information from the data set is skipped. In our task, the cells are sufficiently big that a lower rendering quality is possible without skipping entire cells. The users could change the reconstruction quality with a virtual dial on the Cave wall, which in turn changes the frame rate.

Volume size: Two dials allow setting the perceived size of the data, and separately the size along the z-axis (thickness).

Data intensity: The data set used for the tasks consisted of two color channels, red and green. To perform the task of counting cells, the green channel sufficed, the red channel did not even add any useful information. The user can change the intensity of each channel interactively, and before the tasks we set the red channel to be transparent to make sure the users were not distracted by the it. We set the intensity of green to the same default position for all users, and only a few users changed this value during the tasks.

Markers: Cone-shaped markers can be placed in the data set to mark a cell so that it is not counted twice. The computer keeps track of the number of markers placed, so that the users themselves do not need to count. The size of the markers can be changed with a virtual dial.

Pointer length: The markers showed up at the end of a virtual stick, which seemed to come out of the tip of the wand. The length of the stick can be changed with a dial. Shorter lengths allow more intuitive and direct marker placement, especially at low frame rates. A longer pointer, however, allows the user to move his hand around less and move the marker by merely turning the hand, instead of moving it. The latter proved to be particularly useful in the Fish Tank environment, where the range of motion is naturally limited by the sitting posture and the smaller space between the user and the screen, as opposed to the Cave.

EVALUATION

Six students participated in our formal user study: we recruited four computer science and one political science undergraduate student, all of whom had no or very little experience with virtual environments. We also invited one expert user, a Ph.D. candidate from the biology department who has used the system in the Cave before, but who had not previously seen the Fish Tank.

Note that in this paper, we will always refer to our users as "he" to conceal their gender.

The user study took place on three days. Each of the participants spent a total of about 1.5 hours in our lab.

The users first filled out a consent form and a short questionnaire, asking for statistical information like age, gender, and previous experience with virtual environments. Then we orally explained them the task: they were about to see two data sets from a confocal laser microscope depicting cells of the larva of fruit fly (Drosophila). The users should count the cells by counting the visible nuclei, which had been stained with fluorescent proteins before the microscope scanned them. Because each cell has a nucleus, their numbers are equivalent. To help the users perform this task, they should place virtual markers, the number of which the computer keeps track of. The markers help prevent counting cells multiple times, and to know which ones had already been counted. Figure 4 shows one user's data set when he was finished.

The task was divided into four sections. We started with a practice run with the practice data set (see Figure 5a) in one

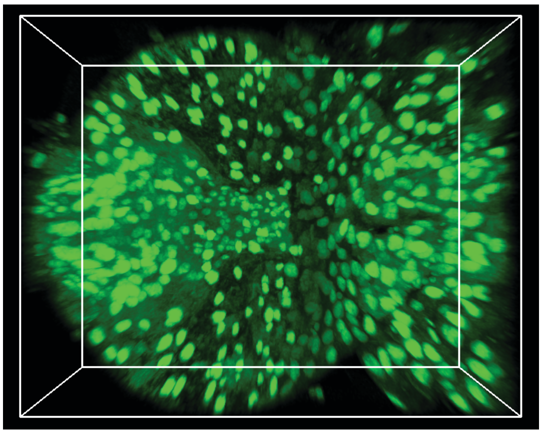

(a)

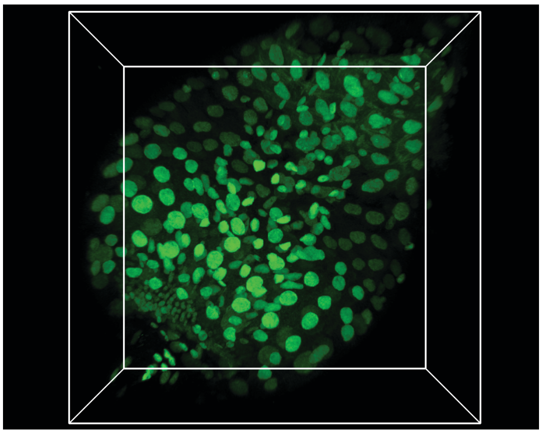

(b)

Figure 5. The two data sets we used in the user study. Both are scans of confocal laser microscopes, showing parts of internal organs of Drosophila larvae. (a) is the data set we used in the training, (b) is the one we used in the trial.

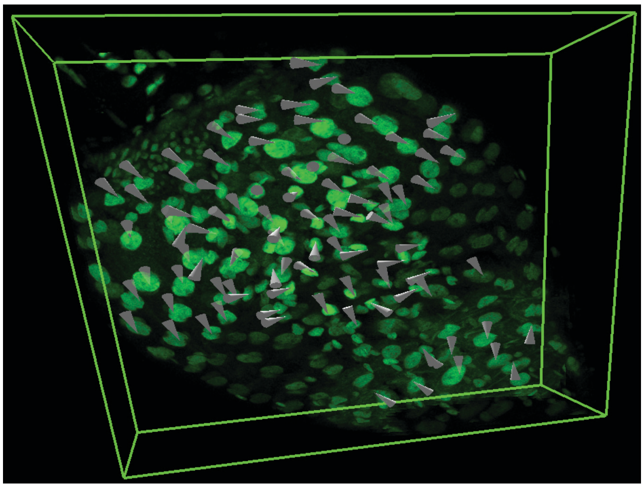

Figure 4. A subject marked 93 cells in a trial data set in the allotted ten minutes.

of the virtual environments (environment #1), and a trial run with the trial data set (see Figure 5b), again in environment #1. Then we walked the subject over to the building in which the other virtual environment (environment #2) was located and did a practice run in environment #2, with the same practice data set as before. Finally, we ran a trial with the trial data set in environment #2. For each subject, we alternated the order of environments: if one started in the Cave, the next one would start at the Fish Tank, and vice versa. Each of the four parts was timed to about 10 minutes. The practice runs were not as strictly timed as the trial runs to allow for differences in how quickly the participants felt comfortable using the systems. In most cases, they did not exceed the 10 minutes.

The subjects were encouraged to ask questions about how to use the system at any point. If we noticed that someone did not take full advantage of the system's capabilities, for instance if he struggled with cells that were too small to mark, we would make suggestions for improvement. We would not, however, take over the controls during the trial runs, or enter the Cave.

After the final run, the subjects were given a two-page questionnaire with ten questions to rank Cave and Fish Tank, along with five free questions. In the ranking questions we asked them to rank the two environments according to several criteria. For each question, there were seven boxes in a row to write in the first letter of the environments ("C" for Cave, or "F" for Fish Tank). Under the boxes we indicated the direction of "more" or "less" of whatever the criterion was. So the users effectively ranked each environment on a scale from 1 to 7.

In the free questions, we asked the users for their thoughts on a variety of other issues. The study ended with a short debriefing, in which we summarized the goals of the study.

In the following section, we will list the questions that we asked, and we will summarize the answers our subjects gave us. We will also present an analysis of the task performance, and our observations while the users were performing the tasks.

Results

Our user study produced three types of results. The first is the task performance, which answers directly the overall question of which environment is better suited to solve the given problem of counting cells. The second type are the results we obtained from the questionnaire. The third type are observations that we made while running the tests. We will present all these results in the following sections.

Task Performance

Figure 6 shows that the users placed on average more markers in the Cave than at the Fish Tank. Four out of five of the users placed between 1.3 and 2.2 times the number of markers in the Cave as compared to the Fish Tank. One

Table 1. Post-questionnaire ranking criteria and results: average and standard deviation (in parentheses)

#	Criterion	Cave	Fish Tank
1	Ease of learning how to use the system	5.7 (1.2)	3.8 (1.2)
2	Help provided by the system to perform the task	5.3 (1.6)	4.0 (1.4)
3	Ease of use once you had learned how to use the systems	6.7 (0.5)	3.8 (1.3)
4	The speed you experienced to perform the tasks	5.8 (1.2)	4.0 (1.7)
5	The perceived quality of the displayed data	5.3 (1.4)	4.7 (1.4)
6	The effectiveness of the user interface	6.2 (0.8)	3.8 (1.8)
7	How comfortable were you using the systems	6.7 (0.5)	3.3 (2.9)
8	How confident were you when using the systems	7.0 (0.0)	4.2 (2.6)
9	How much did you enjoy working with the systems	6.7 (0.5)	4.0 (1.8)
10	Your preference if you were a professional	6.3 (1.2)	3.5 (2.1)

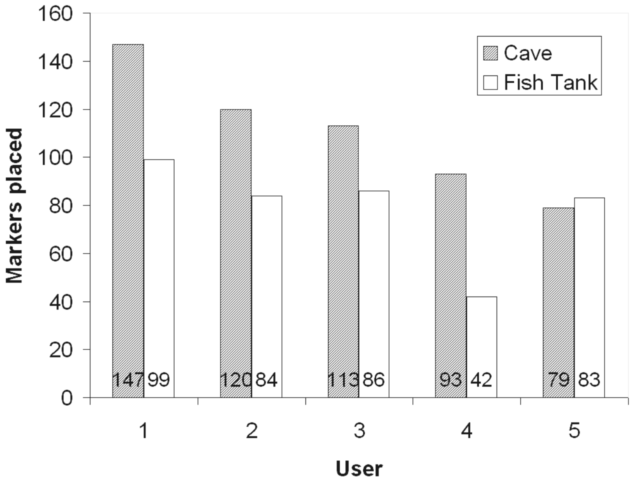

Figure 6. Overview of the number of markers the users placed in the trial runs. The average number of markers placed in the Cave is 110.4 (standard deviation: 26.1), at the Fish Tank it is 78.8 (standard deviation: 21.6).

user got dizzy in the Cave and did not finish the task, so we could not use his results in this statistics.

Questionnaire

The users filled out two questionnaires: one before the task performances ("pre-questionnaire"), and one after ("post-questionnaire").

In the pre-questionnaire, we asked the users about their age (average: 22.5 years, standard deviation: 5.7), sex (3 male, 3 female), and their dominant hand (all right handed). We asked for their concentration or job title: one was a biology graduate student, one a political science undergraduate student, the others were computer science undergraduate students. We also asked for vision deficits: all but one of the users answered they had deficits, but they were wearing contact lenses or glasses that corrected them.

We also asked for prior experience with 3D graphics programs: three users had minimal to no experience, three had previously used such programs in class or at home. To our question for prior experience with virtual reality systems, two users answered they had used virtual environments once, one had used the Cave several times, the remaining ones did not have previous such experience.

In the post-questionnaire, we asked the users to rank Cave and Fish Tank according to ten criteria, and we asked five free questions. Table 1 lists all the ranking criteria and shows average values and standard deviations of the results.

The Cave scored consistently higher that the Fish Tank. Only rarely did users rank the Fish Tank higher than the Cave: two users at criteria 4 and 5, one at criterion 7, and one at criterion 10.

After the ranking questions, we asked five free questions. In the following paragraphs we will present the most interesting answers that we got.

Question 1: "What do you think should be done to improve the effectiveness of any of the systems?"

One user suggested "some kind of scroller" to adjust the length of the pointer, because he changed it ften. He also suggested to "take into account the gap between the screens" at the Fish Tank: we were using a virtual desktop that extended over both monitors, which resulted in a slightly offset picture on the right monitor. Two users suggested, independently from one another, to include the non-dominant hand in the interaction. One suggested to "put the navigation in it", another suggested to have it "control rotation, length, width of data, etc, as you would change modes in Photoshop".

Question 2: "Is there anything you especially liked or disliked with any of the systems?"

One user liked about the Cave "the opportunity to put my hand through the image", and another one liked the "ability to immerse yourself in the object". Two users disliked about the Fish Tank that the "angle of my hand was awkward" or "in less natural positions than in the Cave". One user "disliked marking data points that were behind others" in the Fish Tank. Another one disliked about the Fish Tank "not being able to see the entire pointer—all you could see was on the actual screen".

Question 3: “Do you think we are doing a fair comparison of the systems? If no, please explain why not.”

All but one of the users thought the comparison was fair. One user thought that “if the Fish Tank had better sensors it would score slightly higher”.

Question 4: “Did you feel disoriented or discomfort at any time during the experiment? If so, please briefly describe.”

Two users reported that their arms and wrists got tired in the Fish Tank. One user got dizzy enough in the Cave experiment that he had to abort it after about 7 minutes. One user “felt frustrated and annoyed at my inability to maneuver as I wanted” at the Fish Tank. One user reported that at the Fish Tank he was “positioned uncomfortably”.

Question 5: “If you have any additional comments concerning your experience in this study, please feel free to provide them here.”

To this question we mainly got positive comments on how much the users liked the Cave. One user suggested to add a “long distance shot” of the data set to the existing visualization, because “it is hard to get a good view of the screen when you are up close”.

Observations

We were surprised at how quickly all candidates learned how to use the system, no matter if they started at the Fish Tank or in the Cave. Ten minutes practice time turned out to be enough in all cases, in fact in the second environment, some of the participants told us they were ready to start with the trial data set before the practice time was up.

There were considerable differences in how the subjects worked with the data sets in both environments. One of them, who did not have previous experience in virtual environments, moved around much more than the others. At the Fish Tank, he moved the entire arm with the wand, and he even got up from his chair at some point to see if he could work better when standing (he figured he could not). Note that this was even before he was in the Cave. During the practice run in the Cave he walked around most of all participants, and moved the data set to many different places in the Cave. When asked why he did this, he said he wanted to see what it looks like from a different perspective.

In the Cave, after having set up all viewing parameters, most of the other subjects ended up standing in the middle of the Cave while performing the trial task, not moving their bodies much and focusing on marking the cells. Differences were in the size at which they viewed the data set, the size they set the markers to, and the length of the pointer, i.e., the distance of the hand from the newly to be placed markers.

At the Fish Tank, the users were naturally much more stationary than in the Cave, as they were sitting in front of the two monitors. They could rest their arms on the chair’s armrests, but many of them did not take advantage of this for most of the time because that posture did not allow them to reach everywhere they wanted. Interestingly, all users preferred to place the data set behind the plane of the screens, even after having been suggested to move it closer to reach it better. Therefore, they had to set the distance from wand to marker to a higher value than otherwise, which required more effort to precisely place the markers.

It was interesting to see that most users selected a higher rendering quality at the Fish Tank than in the Cave. Even though this resulted in lower frame rates at the Fish Tank, they preferred their settings, despite us making them aware of their options. In the Cave, the users set the quality level so that they got about 8 frames per second (fps), at the Fish Tank they were happy with about 4 fps.

The users consistently chose longer pointers at the Fish Tank than in the Cave. In the Cave, most users preferred the marker to be as close to the hand as possible, without blocking it with the hand. At the Fish Tank, all users placed the data set behind the screen, so that the distance between hand and marker was larger than in the Cave.

DISCUSSION

Rendering quality. We were surprised at most users giving rendering quality higher preference at the Fish Tank than in the Cave. We hypothesize that due to the Fish Tank’s crisper display (more pixels per inch and clearer pixel boundaries), the artifacts are more noticeable at the Fish Tank than in the Cave and thus the human vision system cannot as easily ignore them and create a continuous image. This might be similar to a TV image which looks fine on a large but low resolution TV screen, but looks pixelated on a smaller but higher resolution notebook screen. In the Cave, the CRT projector technology slightly blurs the image.

Confidence. Another interesting result is that the users consistently scored the Cave higher in confidence than the Fish Tank. We hypothesize that this is due to the more natural interface in which they could use their whole bodies to move around in the virtual world and felt higher presence. Many users made the data sets big enough in the Cave that they could put their hands inside the boundaries of the data set. At the Fish Tank, all users chose to put the data set behind the screen, which did not allow them to put their hands in the data set. This, too, might have added to higher perceived confidence in the Cave. Another reason could have been the lower frame rate the users chose at the Fish Tank: the therefore higher latency, along with the longer distance between hand and marker that they chose, allowed a less direct placement of markers and required more visual feedback to accurately place markers.

Trackball. As stated earlier, the trackball on the wand which rotates the data set around its center, did not work at the Fish Tank, but the subjects used it in the Cave. We thought about the impact this might have had on the study. Eventually, we do not think the impact was very significant, because the users rarely rotated the data set at all once they had positioned it. At the Fish Tank they could have done rotation with the fork metaphor, and some users did so. After all, the data sets had ten times more detail within the slices than along the z-axis. Nevertheless, some users told us the trackball would have been nice if it had worked, even though they never actually tried to rotate the data set. It remains speculation whether they would have done this more often otherwise.

Cave preference. It was surprising to us that the Cave scored considerably higher in most rankings than the Fish Tank. We had not anticipated this after the study by Demiralp et al. [6]. What might be the reasons for this? We believe that it is due to the fundamental differences in the tasks. In [6], the data set is solid and only has features to look for on its surface. By rotating it the user can see all he needs to move on, and he does not need to interact in any other way than rotating it. So only three degrees of freedom of the input device are actually used, whereas our task required control over all six of them. Our data set has features (cells) all over the place, many of them can only be seen at specific viewing angles because they are blocked by other cells. This requires much more careful control over the viewpoint. Additionally, placing markers near these small features requires very fine movements of the wand, which were not required in [6]. On the other hand, it is possible that the users were overwhelmed by the wow effect of the Cave, which most of them had never seen before. We tried to make the subjects aware of the goals of our study to get results independent from this possible influence, but we will need to run more studies with differently experienced people to factor it out entirely.

Arm fatigue. We did not expect the users to complain more about arm fatigue when using the Fish Tank than in the Cave. However, one users said that he “could work in the cave all day” but at the Fish Tank only for about 15 minutes because of arm fatigue. We hypothesize that our method of pointing and clicking at features worked better in the Cave than at the Fish Tank because the hand was part of the virtual world, at the exact location where it was in the real world. In order to see where the hand is, the user would have to hold it up in front of him because this is where the screens are, but that is obviously an uncomfortable position in the long run. In the future, we will experiment with real hand positions that are offset from the virtual hands so that the users can leave them in a more comfortable position. We expect that this will require a similar learning effort as what people experience who use a desktop mouse for the first time: it requires some getting used to before one can intuitively move the pointer on the screen while the mouse is located offset from it.

CONCLUSION AND FUTURE WORK

Our study showed that subjects could mark cells more quickly in the Cave than at the Fish Tank system, and all but one subject preferred the Cave interface. In the Cave subjects scaled the data set up so individual cells were about 4-6 inches in diameter. However, at the Fish Tank they scaled the data such that cells were under one inch in diameter. Our hypothesis is that bigger features are easier to mark, but you only want to make them big if you have a large field of view and a large interaction space. Both conditions are features of the Cave, but not of the Fish Tank.

In the future, we will specialize the user interfaces more to the environments to try to optimize user performance. Both environments would have benefited from higher frame rates, so faster volume rendering techniques are needed.

ACKNOWLEDGEMENTS

This work was supported by DOE award number DE-FG02-03ER25582 and LLNL Research Subcontract No. B527302. We wish to thank Professor Kristi Wharton and Lorena Soares for their collaboration on this work. We would also like to thank our undergraduate research assistants Devon Penney and Alexander Rice, who helped implement the software system over the summer. This work used UNC’s VRPN library, which is supported by the NIH National Research Resource in Molecular Graphics and Microscopy at the University of North Carolina at Chapel Hill.

REFERENCES

1. Amira. *Visualization system, Indeed - Visual Concepts GmbH, Berlin, Germany*. URL: http://www.amiravis.com, 2003.

2. R. Brady, J. Pixton, G. Baxter, P. Moran, C. S. Potter, B. Carragher, and A. Belmont. *Crumbs: A Virtual Environment Tracking Tool for Biological Imaging*. IEEE Symposium on Frontiers in Biomedical Visualization Proceedings, pp. 18–25, 1995.

3. COVISE. *Visualization system for scientific and engineering data, VirCinity GmbH, Stuttgart, Germany*. URL: http://www.vircinity.com, 2004.

4. C. Cruz-Neira, D.J. Sandin, and T.A. DeFanti. *Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE*. ACM SIGGRAPH 93 Proceedings, pp. 135–142, 1993.

5. T.J. Cullip and U. Neumann. *Accelerating Volume Reconstruction With 3D Texture Hardware*. Technical Report TR93-027, University of North Carolina, Chapel Hill, 1993.

6. C. Demiralp, D.H. Laidlaw, C. Jackson, D. Keefe, and S. Zhang. *Subjective Usefulness of CAVE and Fish Tank VR Display Systems for a Scientific Visualization Application*. IEEE Visualization ’03 Poster Abstract, Seattle, WA, 2003.

7. EnSight. *Visualization system, Computational Engineering International (CEI), Apex, NC*. URL: http://www.ensight.com, 2004.

8. J.E. Swan II, J.L. Gabbard, D. Hix, R.S. Schulman, and K.P. Kim. *A Comparative Study of User Performance in a Map-Based Virtual Environment*. IEEE Virtual Reality ’03 Proceedings, Los Angeles, CA, pp. 259–266, 2003.

9. SGI Open Inventor. *Open source visualization API, Silicon Graphics, Inc., Mountain View, CA*. URL: http://oss.sgi.com/projects/inventor/, 2003.

10. TGS Open Inventor. *Object-Oriented 3D Graphics Developers Toolkit*. URL: http://www.tgs.com/pro_div/oiv_main.htm, 2004.

11. OpenSceneGraph. *Scenegraph based graphics library*. URL: http://www.openscenegraph.org, 2004.

12. J.P. Schulze, U. Wössner, S.P. Walz, and U. Lang. *Volume Rendering in a Virtual Environment*. Proceedings of the Fifth Immersive Projection Technology Workshop (IPTW'01) and Eurographics Virtual Environments (EGVE'01), Springer Verlag, pp. 187–198, 2001.

13. Leica TCS SP2. *Leica Microsystems home page*. URL: http://www.leica-microsystems.com, 2004.

14. K. Swartz, U. Thakkar, D. Hix, and R. Brady. *Evaluating the Usability of Crumbs: a Case Study of VE Usability Engineering*. Proceedings of the 3rd International Immersive Projection Technologies Workshop, May '99, Springer-Verlag, pp. 243–252, 1999.

15. OpenGL Volumizer. *Volume visualization API, Silicon Graphics, Inc., Mountain View, CA*. URL: http://www.sgi.com/software/volumizer/, 2003.

16. VTK. *Open source graphics API for 3D graphics, image processing, and visualization*. URL: http://public.kitware.com/VTK/, 2003.

17. C. Ware, K. Arthur, and K.S. Booth. *Fish Tank Virtual Reality*. Proceedings of the Conference on Human Factors in Computing Systems, Addison-Wesley Longman Publishing Co., Inc., pp. 37–42, 1993.

18. W. Wohlfahrter, L.M. Encarnacao, and D. Schmalstieg. *Interactive Volume Exploration on the StudyDesk*. Proceedings of the Fourth International Immersive Projection Technology Workshop, Ames, Iowa, 2000.

Contributions and Benefits Statement:

Our user study of Cave and Fish Tank environments showed that most users preferred the Cave and achieved better results in a counting task with it.