Proceedings of *CS237 '14*

Interdisciplinary Scientific Visualization, Brown University

Prof. David H. Laidlaw TA: Steven R. Gomez

December 2014

Program

1. **Evaluating the Effects of a 3D-Immersive Environment on Learning Strategies for the Visual Control of Navigation**
   Diego Morales, Nicholas Catoni, Timothy Parsons
2. **CAVEBAT: Visualizing Bat SONAR in 3D**
   Elijah ben Izzy, Evan Birenbaum, Nate Burnell
3. **CAVEBAT: An Interactive 3D Visualization of Biosonar Sound Beams**
   Evan Birenbaum, Nathan Burnell, Elijah ben Izzy, Laura Kloepper, James Simmons
4. **Visualizing Biosonar Beams in 3D with Point Clouds**
   Nathan Burnell, Elijah ben Izzy, Evan Birenbaum
5. **Evaluating the Effectiveness of a 3D-Immersive Environment on Learning Strategies for the Visual Control of Navigation**
   Nicholas Catoni, Timothy Parsons, Diego Morales, William Warren, Youssef Barhomi
6. **Evaluating Attributes of Virtual Reality Displays for Interacting with Isosurface Renderings**
   Johannes Novotny, Wesley Miller, David H. Laidlaw, Wallace S. Lages, Doug A. Bowman, John J. Socha, Bireswar Laha
7. **Encoding Qualitative Uncertainty in Visualizations for Archaeological Research**
   Richard Nguyen, Scott Houde, Sheila Bonde
8. **Visualizing Qualitative Uncertainty in Archaeological Research**
   Scott Houde, Richard Nguyen
9. **Evaluating Effectiveness of Human Navigational Experimentation in an Immersive CAVE Environment**
   Timothy Parsons, Nicholas Catoni, Diego Morales
10. **Evaluating Attributes of Virtual Reality Displays for Interacting with Isosurface Renderings**
    Wesley R Miller, Johannes Novotny

Evaluating the Effects of a 3D-Immersive Environment on Learning Strategies for the Visual Control of Navigation

Author
Diego Morales
diego_morales@brown.edu

Co-Author
Nicholas Catoni
nicholas_catoni@brown.edu

Co-Author
Timothy Parsons
timothy_parsons@brown.edu

Brown University

ABSTRACT

We present our study of navigational decision-making in immersive virtual reality and the results we found. Most users are not accustomed to the experience of virtual reality, and we found in our experiment that a significant number of participants reported feeling motion sickness to the point that they could not continue and had to quit around 48 trials. After adjusting the height of the simulation camera, reducing the scale of the scene in relation to the size of the user, and giving the user slightly more control over rotation speed, we then found that users reported feeling far less sick than before and were able to finish 96 trials. Furthermore, among the users that finished, we observed that many reported strong feelings of boredom that were affecting their pathing decisions.

Keywords: Navigation, psychophysics, sickness, monotony.

1 INTRODUCTION

Human beings are exceptionally good at navigating our environments. We possess an innate sense for not only avoiding obstacles but also reaching goals. In fact, most of the navigational decisions that we make are barely conscious. Humans can view a scene, take in all of its visual cues for navigation, and then make rapid, dynamically changing navigation decisions based on how the scene changes. Furthermore, we can also do much of this simultaneously as we are actually making decisions and getting feedback from those decisions.

Previous research by Fajen and Warren has been able to create a fairly accurate model of human navigation in the presence of obstacles and goals[1]. Their model can create smooth, realistic paths through obstacle positions towards a goal state using various angles related to the position of the user (Figure 1). Fajen and Warren used sets of differential equations to minimize user distance from the obstacles while still making sure to safely avoid them. They collected their data in Brown University's Virtual Environment Laboratory(VENLab) using a head-mounted display.

Understanding the ways humans make basic navigational decisions given specific visual cues is extremely important, but the health and saefty of participants in user studies is more important. Given that creating walkable real world environments to be used for simple user trials is a large effort, utilizing virtual reality technology to simulate real world situations is beneficial in terms of cost and effort. However, while humans have no problem navigating the real world, navigating virtual worlds introduces several problems. Researchers must especially consider not only how virtual reality affects users' decisions compared to the real world but also how virutal reality affects the participants themselves.

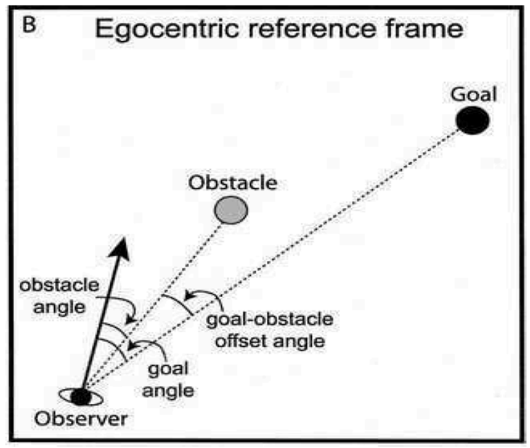

Figure 1: The egocentric reference frame including angles from the current direction to the goal and obstacles [1].

2 EXPERIMENT

Our experiment (Figure 2a) consisted of creating an environment similar to the one used by Warren in his experiments (Figure 2b), attempting to match textures and simplicity. This was done in order to compare the walking results from the VENLab with flying results in the CAVE. We used the Blender game engine to create a virtual reality simulation which runs the user through 96 trials. In each of these trials, the user is spawned some distance form the blue goal post, and after moving forward one meter, using a joystick to control flying movement, we display seven obstacles. The user must reach the goal state while avoiding the obstacles. We used Brown University's CAVE to run our program and conduct user trials.

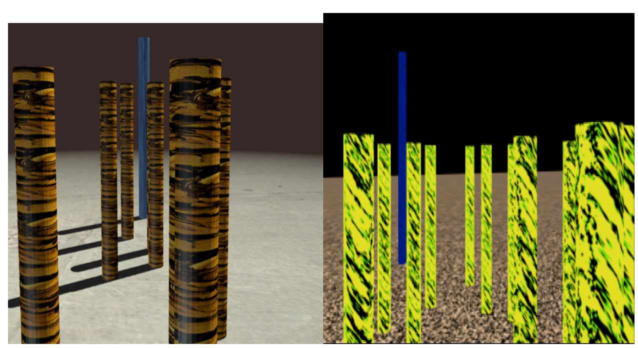

Figure 2: (a) LEFT: Our CAVE experimental environment.

(b) RIGHT: Warren's original VENLab experimental environment.

2.1 First Round of Users

We started by running six users though two sets of 48 trials. The users were each asked to stand in the center of the CAVE and hold the joystick with whatever hand they preferred. They then performed their trials after performing a few practice trials to become accustomed to the virtual environment. Of the six users that participated, five out of the six were unable to finish the entire experiment. All five reported feelings of heavy motion sickness and nausea that forced them to stop, and many were sure after their first set that a second set of 48 trials would be too much for them.

We collected information from every user about the motion sickness, even the one who did not experience it, and we gathered observations about possible causes from Brown CAVE experts. Most users stated that the environment just didn't feel natural; the most frequent comments we received about the possible sources of the motion sickness were: flying/floating movement in space, the perceived size of the obstacles and goal posts relative to the user, the speed of rotational movement controlled by the joystick, and obstacle shadows mixing with the obstacle textures. We then attempted to fix these issues and proceeded to run more user studies.

2.2 Second Round of Users

We were able to get seven new participants and one old participant for our updated simulation. Also for this round of studies, we chose to instead have three sets of 32 trials in order to give the users more breaks and decrease the likelihood of anyone getting sick. Of the seven new users, all of them finished all of their trials. However, they all still reported some sense of unnaturalness that built up to slight discomfort in the environment and a light level of tolerable motion sickness.

The one user that was able to come back again from the first round of studies reported that the simulation felt "a lot better." She was unable to start the second set of 48 in the first round, but she was able to finish all 96 with the updated simulation. While this was a significant improvement, she also mentioned that she still felt some tolerable level of motion sickness.

This round of 8 users still reported light levels of motion sickness despite being able to finish. The most frequent comments we received about the possible sources of motion sickness in the new simulation were: just doesn't feel natural, obstacles immediately and suddenly appearing close to the user at the beginning of every trial, rotational movement in the simulation compared to more natural sideways movement, and observing non-stereo images in the uncovered periphery of the CAVE goggles.

3 Discussion

Compared to Warren's models, our data for users' navigational decision making was confounded by two primary factors: motion sickness and monotony. Motion sickness was not only preventing users from finishing, it was also affecting their decisions by preventing them from taking in enough visual cues to make good navigational decisions. Monotony of the trials, on the other hand, led to bored users choosing to take easier and more interesting, while still inefficient and indirect, paths to the goal despite being told to find the most efficient path in the shortest time.

3.1 Motion Sickness

Of the users that reported feeling heavy motion sickness, most mentioned that their navigational decision making on the last 10-20 trials they performed, when their motion sickness was worst, was not as informed as the rest they had done first. This was because as their motion sickness increased from trial to trial, the susceptible users felt it more and more necessary to stop looking around at the environment and at obstacles, in order to remediate their increasing sickness. These users felt as though the same visual cues they utilized to make intelligent navigational decisions were also giving them increasing levels of sickness, so eventually they decided to keep their eyes fixed at the goal. Many reported that this was the only way they were able to finish a single set.

After revising the simulation, the most received comment was that the environment didn't feel natural. Given that this generality could apply to various problems, we concluded that there were numerous causes of motion sickness that varied person to person. While we were able to drastically improve our users' ability to finish the trials, more work must be done to completely eliminate this sense of motion sickness as it could still be affecting users' navigational decisions and thus the data used to compare to Warren's model.

3.1.1 Prior Virtual Reality Experience

We observed an unexpected pattern after asking all users if they had prior virtual reality experience. Six users that reported having had prior VR experience also reported feeling sick, and two users that reported not having had prior VR experience also reported not feeling sick. This suggests that prior VR experience could be unrelated to feelings of motion sickness on the CAVE, although more users will need to be tested to prove this.

3.2 Monotony

Of the users that had little problem with motion sickness during their trials, they reported that boredom was affecting their navigational decision making despite being told to find the shortest and most efficient path to the goal post. These users either went out of their way to explore the environment (Figure 3a), or these users chose not to attempt to find the quickest and most efficient path instead choosing to continually take an easier but less efficient and slower path, such as around the obstacles instead of through them (Figure 3b).

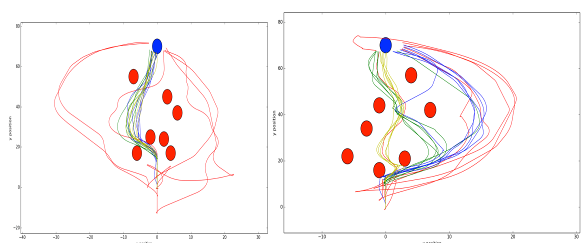

Figure 3: (a) LEFT: Red user wondering around environment. (b) RIGHT: Red and blue user taking easy but inefficient paths around all of the obstacles.

4 Conclusion

Studying human navigational decision making is important, and doing so in VR is practical. However, in order to get scientifically accurate, unbiased results, one must consider all of the possible ways VR can change the decisions users make. Further research can attempt to further refine the environment in hopes to reach an environment users deem "natural."

References

[1] J. Fajen, B. R., & Warren, W. H. (2003). Behavioral dynamics of steering, obstable avoidance, and route selection. *Journal of Experimental Psychology: Human Perception and Performance*, *29*(2), 343–362. doi:10.1037/0096-1523.29.2.343

CAVEBAT: Visualizing Bat SONAR in 3D

Elijah ben Izzy* Evan Birenbaum† Nate Burnell‡

Brown University

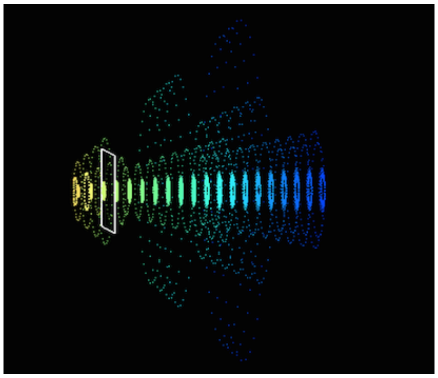

Figure 1: CAVEBAT rendering of bat biosonar beam profile

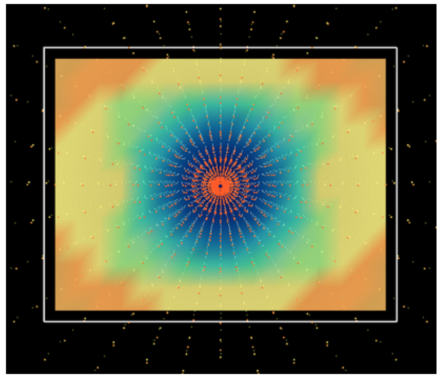

Figure 2: CAVEBAT rendering of biosonar beam looking directly at the bat.

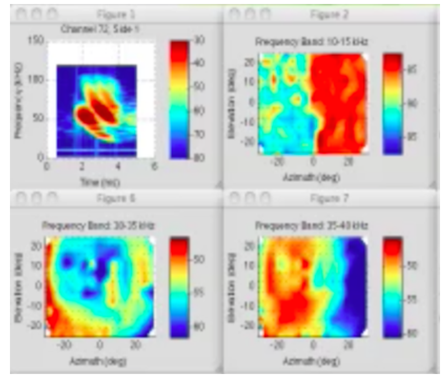

Figure 3: Previous visualizations: heatmaps of biosonar at different frequencies

Abstract

We present and evaluate CAVEBAT (Configurable Accoustic Environment for Biosonar Analysis and Testing), a novel tool for visualizing biosonar in 3D. Through the process of echolocation, Bats emit complex sound pulses to locate prey and learn about the environment around them, allowing them to navigate through dark environments. While a variety of mathematical models describe this process, it is still poorly understood and difficult to explain. Our tool enables researchers to better visualize and communicate the properties of bat biosonar. We ran user studies evaluating the effectiveness of CAVEBAT as an educational and comunication tool for both experts in the field and potential users with less experience. Both studies confirmed its usefulness, and taught us what makes 3D visualizations of sound effective.

Keywords: Visualization, SONAR, bats, echolocation

1 Introduction

1.1 Background

Bats emit complex patterns of high-pitched sound to track prey and determine their location, with regards to other bats as well as the environment around them. They emit two distinct harmonics (pitch ranges) simultaneously, which allow them to better understand the echoes they receive. As the sound is emitted, the frequency of each harmonic decreases over time. The big brown bat (Eptesicus fuscus) emits a first harmonic ranging from 55khz down to 20khz, and a second harmonic ranging from 110khz to 40khz. This *downsweep* of frequencies enables bats to figure out the exact portion of the pulse that is returning to them. [1, 6]

Bats' emission of biosonar also varies by the angle of the sound with regards to the target. In general, off-axis amplitudes tend to be lower than on-axis amplitudes. The extent of this variance depends both on the frequency of the sound and its angle of propagation, forming complex lobes of high-amplitude sound components. This property is best described by the *Piston Model* for sound propagation. [3]

*e-mail: elijah$_{b}en_{i}zzy@brown.edu$

After the sound is emitted, it attenuates (loses amplitude), depending on a variety of conditions. As a sound beam radiates outward, it loses energy proportional to the square of the radius. Furthermore, the beam loses energy to the air around it. Higher frequencies tend to dissipate more rapidly than lower frequencies, although the rate of attenuation is heavily dependent on atmospheric conditions such as tempaerature and humidity. [2, 5]

1.2 Research Goals

We aimed to display the above properties of biosonar emission and propogation in a configurable environment that allows users to fully explore the mechanisms with which bats internalize the environment around them. We collaborated with Brown University's Batlab, a group of scientists dedicated to understanding the process of echolocation in bats as well as other animals. We chose to evaluate our method of visualizing sound in 3D as a tool for educating people about the properties of sound, illuminating effective methods of sound and biosonar visualization.

2 Previous Work

While previous visualizations have managed to display biosonar to various levels of effectiveness, they have generally been constrained to 2-dimensional heatmaps or artistic rendering of the sound. The work discovering the previously mentioned effects propose limited visualizations, yet the focus is on the scientific and not the visualization components. False killer whale biosonar data has enabled a variety of visualizations, yet these mainly consist of adjacent heatmaps rendering the amplitudes at different frequency or line charts showing the relationship between frequency and amplitude composition. [4, 7]. The Batlab provided us with animated heatmap renderings of data they had collected in a series of experiments recording the sound emitted by specially trained bats.

3 Implementation

Our product is a desktop-based software tool that renders a biosonar beam over time in the environment experiments carried out by Brown's Batlab. In these experiments, a bat is placed a meter before a microphone array, and the bat's target is placed a meter after

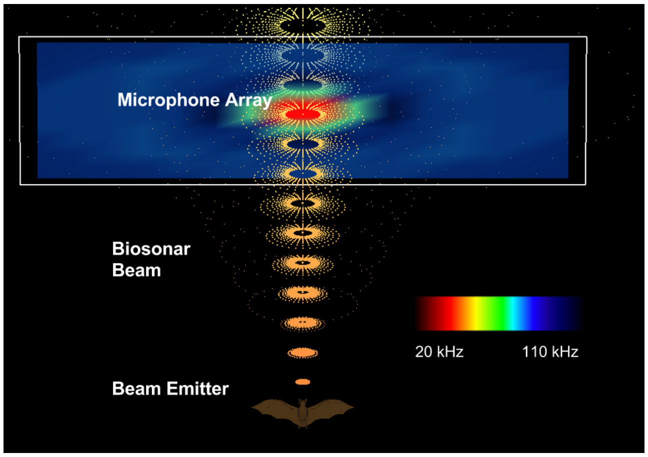

Figure 4: CAVEBAT display.

that. The bat then emits sonar to track the target, which is recorded by the microphones.

We chose to render the beam as a cloud of points. The color of each point represents the frequency of sound at the physical location of that point, ranging from blue (representing low frequency) to red (representing high frequency). The size of the point corresponds to the amplitude of the sound there. The microphone array gives the user a heatmap, which dynamically renders the amplitude of a cross-section of the sound beam. We chose to render more particles towards the interior of the beam to help give the user a clearer picture of its structure and avoid clutter.

The user is allowed to play or pause the simulation, filter between the two harmonics, record a video of the sound, and change the speed at which the simulation is run. The simulation runs in a continual loop.

4 Evaluation

To effectively evaluate this product, we chose to run qualitative surveys and interviews on a variety of potential users. For the purpose of making biosonar discoveries, our tool has a limited group of intended users. As such, our collaborators were the most valuable source of comparison between CAVEBAT and previous visualizations. However, as we also wished to measure our tool as an educational platform with which to explain concepts of biosonar visualization, we asked less experienced users a variety of open-ended exam-like questions to gauge what they learned from using our tool. They generally took about 15 minutes to complete the questions, although we gave no strict time limit.

4.1 Expert Feedback

Our collaborators gave positive feedback, praising its use a ”communications tool”, and less as a tool with which they could immediately gain insight. However, they were excited about the insights they would be able to gain by altering the environment parameters and experiment configuration. They said the tool ”gives a really good way of showing that only a subset of space is illuminated”, and that it represented ”very crucial” aspects of amplitude attenuation. They asked for the ability to include a screen-capture feature (which we have since implemented) so they could show videos in presentations. They confirmed that our renderings were similar to how they had previously envisioned biosonar propagation, yet said the sound beam was much ”crisper” than they had imagined.

4.2 User Study

Our users reacted positively overall to our tool. They quickly saw and commented on the *downsweep* of frequencies, as well as the general pattern of amplitude decrease over time. However, our visualization failed to teach them some of the intricacies of sound propogation, including the details of the piston model and the effects of atmospheric attenuation on specific frequencies.

From this study, we learned a variety of valuable visualization lessons. It seems that rendering sound in 3D is better for use as an educational tool than an insight tool, as experts already know the details of the sound models they are working with. As such, it is useful for quickly bringing people up to speed. From this study we conclude that volumetric rendering of sound as point-cloud data is an effective visualization method. However, our users generally though that representing amplitude by point size is potentially confusing, as there is little variance between different amplitude. People often confused the amplitude with the point distribution that we used. Our users confirmed that the heatmap was an effective rendering tool, yet we received mixed feedback on the contrast between the heatmap and the point cloud (the heatmap had an opposite scale to the point cloud). Users appreciated the general aesthetics of our rendering, and one went so far as to call it ”beautiful”.

5 Conclusion

We presented CAVEBAT, an interactive tool for visualizing bat biosonar in 3D. We rendered a variety of mathematical sound models using a point cloud to represent a beam of bat biosonar. We tested it on both experienced and inexperienced users, and found it to be an effective educational tool. Furthermore we learned which visualization techniques worked and didn’t work in time-varying volumetric rendering of sound.

Acknowledgements

I would like to thank the rest of my team, as well as my collaborators, James Simmons and Laura Kloepper. I would also like to thank David Laidlaw (the professor) and Steven Gomez (the TA) of *Interdisciplinary Scientific Visualization*.

References

[1] M. E. Bates, J. A. Simmons, and T. V. Zorikov. Bats use echo harmonic structure to distinguish their targets from background clutter. *Science*, 333(6042):627–630, 2011.

[2] K. Ghose and C. F. Moss. The sonar beam pattern of a flying bat as it tracks tethered insects. *The Journal of the Acoustical Society of America*, 114(2):1120–1131, 2003.

[3] D. J. Hartley and R. A. Suthers. The sound emission pattern of the echolocating bat, eptesicusfuscus. *The Journal of the Acoustical Society of America*, 85(3):1348–1351, 1989.

[4] S. D. Ibsen, P. E. Nachtigall, J. Krause-Nehring, L. Kloepper, M. Breese, S. Li, and S. Vlachos. Spatial orientation of different frequencies within the echolocation beam of a tursiops truncatus and pseudorca crassidens. *The Journal of the Acoustical Society of America*, 132(2):1213–1221, 2012.

[5] B. D. Lawrence and J. A. Simmons. Measurements of atmospheric attenuation at ultrasonic frequencies and the significance for echolocation by bats. *The Journal of the Acoustical Society of America*, 71(3):585–590, 1982.

[6] J. A. Simmons. Temporal binding of neural responses for focused attention in biosonar. *The Journal of experimental biology*, 217(16):2834–2843, 2014.

[7] M. M. Yuen, P. E. Nachtigall, M. Breese, and A. Y. Supin. Behavioral and auditory evoked potential audiograms of a false killer whale (pseudorca crassidens). *The Journal of the Acoustical Society of America*, 118(4):2688–2695, 2005.

CAVEBAT: An Interactive 3D Visualization of Biosonar Sound Beams

Evan Birenbaum* Nathan Burnell* Elijah ben Izzy* Laura Kloepper† James Simmons†

Brown University

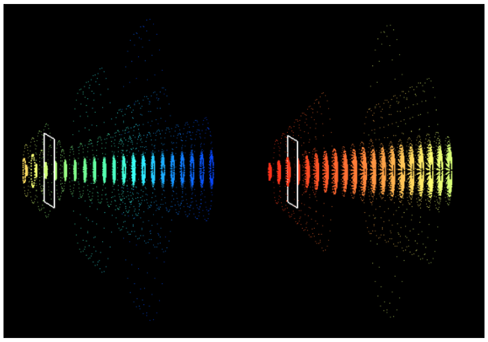

Figure 1: CAVEBAT renderings of two different harmonic frequency sweeps emitted in one biosonar beam. Redder particles indicate lower frequency regions of sound while bluer particles indicate higher frequency regions of sound.

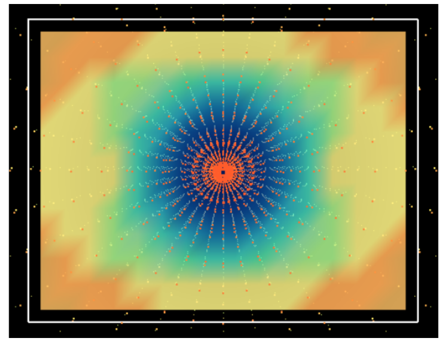

Figure 2: CAVEBAT heatmap displaying the relative amplitude of the biosonar beam as it passes through an array of microphones. Bluer indicates higher amplitude while redder indicates lower amplitude. Higher amplitude sound is proportional to higher energy.

ABSTRACT

We present CAVEBAT (**C**onfigurable **A**coustics **V**isualization **E**nvironment for **B**iosonar **A**nalysis and **T**esting), a novel tool for visualizing and interacting with biosonar sound beams in 3D. Biological sound beams, such as those used in bat echolocation, are complex, invisible, volumetric structures. While conventional methods of study examine beams using time-dependent cross sections, CAVEBAT enables biosonar researchers to view these structures as a 3-dimensional point cloud, using a variety of visual cues such as color and point size to indicate qualities of the sound at a particular point in time and space. We evaluate the usefulness of the tool as a research aid using feedback from professionals in the field of bat biosonar research, James Simmons and Laura Kloepper of Brown University's Simmons Lab. Additionally, we evaluate the educational value of the tool with a qualitative user study of individuals with no experience with the field.

Keywords: 3D visualization, biosonar, acoustics, human-computer interaction.

1 BACKGROUND

1.1 Biosonar

In order to navigate the environment, bats emit high-frequency sound waves into their surroundings and use the reflections of the sound to "see" the world, including obstacles, prey, and other bats. To avoid interference with other sounds, bats emit these waves in complex patterns with a variety of structural features. Our collaborators, James Simmons and Laura Kloepper of Brown University's Simmons Lab, focus their studies on understanding the echolocation of the big brown bat (**Eptesicus fuscus**).

*e-mail: ebirenba/nate/ebenizzy@cs.brown.edu

†e-mail: laura_kloepper/james_simmons@brown.edu

The shape of a biosonar emission is well represented by the *piston model* of sound propagation, which describes how sound is created by air escaping a closing piston. Under this model, the source of the sound is a single point and the relative amplitude of waves emanating from this point depend upon their angle from the axis of the piston. In general, off-axis waves tend to have lower amplitude than those closer to the piston axis, with frequency-dependent variations resulting in "lobes" of high-amplitude sound around a central beam of the highest amplitude sound. [3] This model forms the basis of the emission model implemented in CAVEBAT.

A single biosonar emission includes two distinct frequency ranges, called *harmonics*, which give the bat better perception of its surroundings. Each of these harmonics has a *downsweep* of frequencies in which the highest frequency sound is emitted first, sweeping downwards toward the lower frequency of the harmonic. The big brown bat emits a first harmonic that sweeps from $110,000kHz$ down to $40,000kHz$ and a second harmonic that sweeps from $55,000kHz$ down to $20,000kHz$. [1] After emission, sound waves *attenuate*, or lose amplitude, based on a variety of factors, including distance from the source, frequency, and atmospheric conditions such as temperature and pressure. [4] Each of the above factors is included in the propagation model implemented in CAVEBAT, resulting in a realistic simulation and visualization of a biosonar beam.

1.2 Related Work

Current visualizations of biosonar are limited by the number of dimensions they are able to display. For example, 2-dimensional slices of beams may be recorded with large arrays of microphones, and this data can be visualized as a heatmap displaying the relative amplitudes and frequencies of different parts of the slice. Although a number of these may be placed side-to-side for a primitive picture of the beam over time, it fails to truly capture the 3-dimensional, time-varying nature of a biosonar beam. [5]

Research in architectural acoustics indicates that simple point clouds are an effective method of visualizing sound. This strategy is appropriate for displaying the propagation of a sound through a room, but is not necessarily useful for understanding properties of the sound at a particular moment in time. [2] The visualization present in CAVEBAT combines these two concepts, providing a 2-dimensional heatmap slice of a 3-dimensional point cloud that uses qualities such as color and point size to indicate additional features of the sound.

2 Methods

2.1 Implementation

CAVEBAT features a visualization based on a standard experiment performed by collaborators Simmons and Kloepper. A bat is trained to "chirp" towards a target situated behind a large array of microphones. These microphones measure the sound pressure and frequency of the sound beam at different locations during a series of timesteps, which can then be synthesized into the 2-dimensional heatmap. Below is the visualization produced by CAVEBAT, including both the heatmap and the new point-cloud representation of the sound volume as described by the *piston model* of sound propagation.

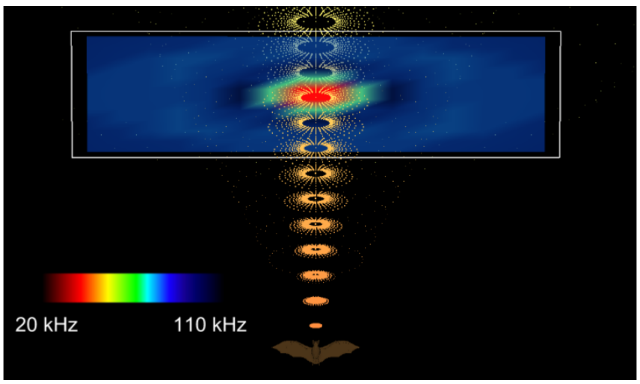

Figure 3: A rendering produced by CAVEBAT with an annotation indicating the frequency which corresponds to the color of a particular particle. The size of a particle indicates the relative amplitude, and the heatmap colors changes based on the amplitude of the beam passing through it, with redder indicating higher amplitude and bluer indicating lower amplitude.

Two primary goals motivate the design of CAVEBAT. First, the program is highly configurable, allowing users to change many aspects of the simulation, including variables such as the orientation and origin of the emitter and target, the speed of sound (to allow for modeling propagation through other mediums like water), and the harmonic frequencies. Second, the program is simple to use, providing the user with a basic interface which supports an orbital camera, pause/play/restart functionality, and the ability to filter the beam by frequency harmonic. In addition to the desktop application, CAVEBAT is implemented for Brown University's immersive CAVE environment using the VRG3D library, allowing users to step directly into the particle cloud representing the sound beam.

2.2 User Study

To evaluate CAVEBAT as an educational tool, we gave several individuals with little experience with sound propagation a survey which required them to use the visualization to answer a series of qualitative questions about the sound beam. Questions were focused on both the general structure of the sound beam and how certain aspects like frequencies and amplitudes changed over the course of propagation. Prior to beginning the survey, users were given a brief introduction to bat echolocation, as well as a tour of the features available in the CAVEBAT program.

3 Results

3.1 Professional Feedback

We solicited feedback from our collaborators on the usefulness of the tool in their studies. Simmons and Kloepper commented that CAVEBAT would be most helpful as a communication tool, allowing them to visualize and present their findings to other biosonar researchers and, perhaps more importantly, other scientists not as familiar with the field. In addition, they noted that the resulting visualization did match their own visions of the model. Both expressed interest in further development of the product, noting that it would be particularly interesting to compare model-based visualizations with data gathered from real experiments.

3.2 User Study

Survey results and feedback from the user study indicate that CAVEBAT does effectively function as an educational tool to explain basic properties of biosonar emission and propagation. Most users quickly identified the down-sweep in frequencies during an emission, and many recognized the general amplitude attenuation that occurs over the course of propagation. However, there was some confusion with regards to what the individual particles represented with respect to the sound as a whole, and most had difficulty understanding how changes in amplitude are dependent on the frequency of the sound at that point.

4 Discussion

The new visualization produced by CAVEBAT provides a more effective means of portraying biosonar sound propagation. Future development will allow researchers to compare different emission and propagation models with results seen in nature and provide them with a simple way to share these results. Additional user feedback will help adjust the visualization to better portray the appropriate features of the sound beam and develop it both as an educational tool and a research tool.

An open problem with the simulation is the determination of sample points. The sound beam is a continuous volume, and therefore representative points must be selected to render it as a point cloud. Currently, points are distributed so that a larger proportion fall closer to the axis of propagation, as this is where the "more interesting" sound falls. This reflects the fact that the higher amplitude sound waves are focused on the axis of propagation, but may not be the optimal way of sampling the sound particles.

Acknowledgements

The authors wish to thank James Simmons and Laura Kloepper for both their assistance in understanding the necessary sound propagation models and their professional feedback on the final product. In addition, the authors wish to thank David Laidlaw and Steven Gomez for their guidance throughout the CS237 research process.

References

[1] M. E. Bates, J. A. Simmons, and T. V. Zorikov. Bats use echo harmonic structure to distinguish their targets from background clutter. *Science*, 333(6042):627–630, 2011.

[2] M. e. a. Bertram. Phonon tracing for auralization and visualization of sound. *Visualization*, (5):151–158, 2005.

[3] D. J. Hartley and R. A. Suthers. The sound emission pattern of the echolocating bat, eptesicusfuscus. *The Journal of the Acoustical Society of America*, 85(3):1348–1351, 1989.

[4] B. D. Lawrence and J. A. Simmons. Measurements of atmospheric attenuation at ultrasonic frequencies and the significance for echolocation by bats. *The Journal of the Acoustical Society of America*, 71(3):585–590, 1982.

[5] J. A. Simmons. Temporal binding of neural responses for focused attention in biosonar. *The Journal of experimental biology*, 217(16):2834–2843, 2014.

Visualizing Biosonar Beams in 3D with Point Clouds

Nathan Burnell* Elijah ben Izzy† Evan Birenbaum‡

Brown University

ABSTRACT

We present a new simulation and visualization of biosonar beams. Biosonar researchers are confronted with the challenge of studying biosonar beams, which are complex, invisible, volumetric structures. Expert feedback, combined with a non-expert user study, suggests that our 3D point cloud representation is an effective way to convey the important properties of these complex volumes.

Keywords: Biosonar, point clouds, scientific visualization.

1 INTRODUCTION

The goal of our work is to enable researchers to study biosonar beams and convey their findings more effectively to experts and non-experts alike. Due to the complex nature of these beams, 2D representations are potentially misleading or fail to capture the complete picture. To that end, we present a new simulation and 3D visualization based on existing research and the physics of sound. Since 3D visualizations are uncommon in this field, we conduct a user study to determine the educational merit of our novel visualization of biosonar beams.

When an animal performs an echolocation task, it emits a beam of sound that can be crudely envisioned as a hemisphere of sonic energy expanding from the animal's emitter (mouth, nose, or forehead) over time. In the case of the big brown bat, the actual structure is more akin to a cone surrounded by a couple doughnuts wrapped in a sphere [7]. The size, structure, intensity, and frequency mixture depends on the species and the task (e.g., searching, navigating, tracking, etc.). These properties also vary with time and length of the call.

Biosonar researchers are faced with the problem of studying these complex structures. Their research informs and inspires new technologies, such as real-time sonar systems [2]. They convey their findings with limited representations of the complex data, which depict only a small portion of the beam at a given time [3, 4]. In order to help overcome this limitation, we present a new, domain-specific 3D visualization that offers another option for studying and teaching biosonar processes.

1.1 Related Work

Past research has identified the mechanisms of beam formation in various animals [7, 3]. We combine these findings with known physical properties of sound, resulting in a new, comprehensive simulation of biosonar formation and propagation. This simulation drives a 3D visualization of the sound beam.

Prior research in acoustics simulation and visualization focuses on the propagation of simple sound waves as collections of points [5]. These techniques are not immediately applicable to rendering volumetric sound beams. We extend this work to support complex sonic structures and simultaneous display of multiple attributes.

Visual exploration of 3D beam shapes could advance the bioacoustics field in unobvious ways, as is the case in other fields [6]. The experimental recordings in [1] are portrayed as top-down, grayscale videos, and although the authors mention the need for extra dimensionality in their microphone layout, the concept of 3D visualization is not realized.

2 METHODS

Our methods are broken down into two sections. First, we describe the simulation of a big brown bat's biosonar beam. Next, we describe the visualization technique. Methods described in this paper are implemented in a software package that runs on a desktop or in a virtual reality CAVE.

2.1 Simulation

The simulation is broken down into two parts. The first deals with the formation of the beam at the emitter (the animal). The second deals with the propagation and decay of the beam over the medium (i.e., air of some temperature and humidity).

The emitter of a big brown bat is its mouth. This is modeled as a small circle on a sphere (see Figure 1). This serves as the emitter surface and is parameterized by its diameter and depth (or focal length). In addition, the bat produces two simultaneous harmonics that sweep from higher to lower frequencies over the duration of the call [3, 7]. We approximate this behavior with a linear model (the behavior is approximately linear). The final component of the emission model dictates the relative amplitude of sound along the surface of the emitter. For this, we use the piston model of sound, which has been shown to be an accurate fit for bat calls [7].

The sound propagates outward from the surface and attenuates. This is in part due to spherical spreading (property of thermodynamics) and atmospheric absorption (depends on frequency, temperature, and humidity). The attenuation behavior in air is given by equation 1, below.

$$A_t = A_0 \cdot \frac{r_0}{d_t} \cdot L(f, d_t, h) \tag{1}$$

This determines the amplitude A of the sound component with frequency f at time t after traveling distance d_t from the initial location r_0. The spherical spreading loss is proportional to distance traveled. The atmospheric loss $L(f,d_t,h)$ is implemented with a publicly available lookup table [8]. The term h is the atmospheric relative humidity.

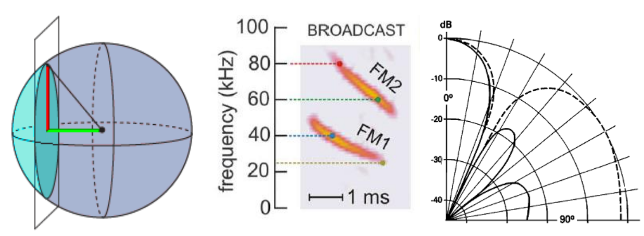

(a) (b) (c)

Figure 1: (a) The bat mouth as a small circle (radius in red) and focal length (in green). Sound is emitted from the light blue surface. **(b)** The frequency behavior of the call of the big brown bat over duration of the call [3]. **(c)** The piston model displaying amplitude loss as a function of off-axis angle [7].

* e-mail: nate@cs.brown.edu
† e-mail: elijah_ben_izzy@brown.edu
‡ e-mail: evan_birenbaum@brown.edu

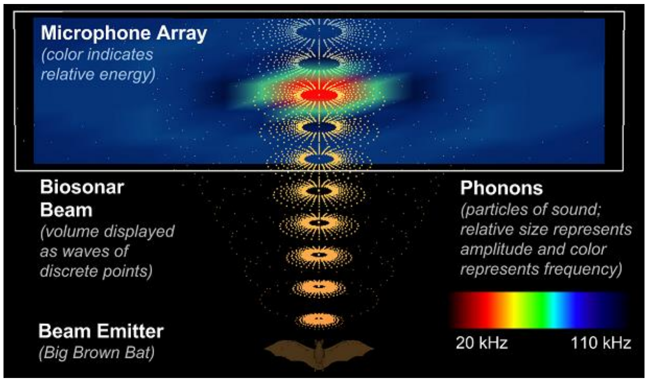

Figure 2: Annotated summary of our 3D visualization, as seen from a camera above the scene; to view several animations, please visit: *http://tinyurl.com/ouad6tm*

2.2 Visualization

Our interactive biosonar visualization is comprised of a few parts (see Figure 2). First, the biosonar beam itself is represented as a cloud of colored points that spread outward from the animal towards an imagined target (distance set by the user). Second, a planar microphone array intersects the beam, which changes color based on incidental sound. Third, the simulation can be paused, sped up or slowed down, similar to playing a video but with the inclusion of 3D camera controls. Additional controls allow the user to alter the visualization, such as by excluding a harmonic.

2.2.1 Biosonar Beam

The biosonar beam is displayed as sequential waves of discrete points, also called phonons (sound particles). Although the beam is a continuous volume, point clouds are a common volume rendering technique and are already commonly used in acoustics modeling [5]. Each phonon is assigned a color (determined by the mix of frequencies at that point) and a size (determined by the greatest amplitude of all frequency components). The distribution of points is determined by a user setting, allowing the user to concentrate the points on different areas of the beam while exploring its overall shape.

2.2.2 Microphone Array

The microphone array is a common instrument used in biosonar experiments. We include a virtual equivalent as both a way to enrich the visualization and provide a familiar element to our expert users. The animal is trained to call towards the 'wall' of microphones. Researchers record the results and display them as 2D intensity maps [4]. Likewise, we interpolate the intensity at each microphone by finding the 8 nearest phonons that enclose the microphone in a box. The intensity is depicted by color, ranging from red for high intensity and blue for low intensity.

3 Results

Since 3D visualization of biosonar beams is a seemingly unexplored problem area, we obtained qualitative feedback from our expert collaborators and a small set of non-experts. The experts were excited to see their data in 3D and greatly appreciated the microphone array, as it helped them understand the scene. Their primary concern was in the point distribution method and how it could potentially confuse observers into thinking the beam was acting as a series of sonic waves rather than a single beam.

In order to assess the quality of our visualization and simulation as an educational tool, we conducted a user study with 7 non-expert participants. Subjects are first asked about their experience with biosonar, acoustics, and echolocation. Then, they are given several minutes to interact with and view the 3D visualization of a big brown bat's biosonar beam. Immediately after, each user is given a short quiz designed to assess how well they inferred different properties of sound beams from the visualization.

The results of this user study are mixed. Although the majority of participants successfully identified the harmonic behavior and frequency down-sweep, most had difficulty determining the intensity of the beam and were confused between point size and density. Color was appreciated as an indicator of frequency, but there were confusions between color used for frequency (in the beam) and intensity (in the microphone array).

4 Discussion

Our results suggest that 3D point cloud visualization is a promising technique for rendering biosonar beams. While color proves to be an effective visual attribute, point size is not as meaningful to our expert and non-expert users. Instead, most users associated point density with intensity. This leaves us with an open issue: how do we render a point cloud such that users can see within the structure, yet accurately represent intensity with point density (which may be greater outside than inside the volume)? We suspect that a greater level of interaction (through configuration options) will alleviate this issue. For example, based on initial feedback, we changed the point distribution parameters and the expert users' primary concern is now remedied. If this were a parameter the users could change, they may be able to discover the distribution that works best for the current task.

5 Conclusion

Our comprehensive simulation and visualization of biosonar beams is effective in conveying some properties of complex sound beams. We find that color is an effective visual attribute, as is point density (if used correctly; in our case, we may have misused it). Overall, our expert users are excited to produce videos and screenshots to share with others in their field, suggesting that this visualization tool has filled a much needed hole for this research community.

References

[1] Ghose, Kaushik, and Cynthia F. Moss. "The sonar beam pattern of a flying bat as it tracks tethered insects." *The Journal of the Acoustical Society of America*114.2 (2003): 1120-1131.

[2] Schillebeeckx, Filips, et al. "Biomimetic sonar: Binaural 3D localization using artificial bat pinnae." *The International Journal of Robotics Research* (2010): 0278364910380474.

[3] Bates, Mary E., James A. Simmons, and Tengiz V. Zorikov. "Bats use echo harmonic structure to distinguish their targets from background clutter." *Science* 333.6042 (2011): 627-630.

[4] Kloepper, Laura N., et al. "Single-lobed frequency-dependent beam shape in an echolocating false killer whale (Pseudorca crassidens)." *The Journal of the Acoustical Society of America* 131.1 (2012): 577-581.

[5] Bertram, Martin, et al. "Phonon tracing for auralization and visualization of sound." *Visualization*, 2005. VIS 05. IEEE. IEEE, 2005.

[6] Van Dam, Andries, et al. "Immersive VR for scientific visualization: A progress report." *Computer Graphics and Applications*, IEEE 20.6 (2000): 26-52.

[7] Hartley, David J., and Roderick A. Suthers. "The sound emission pattern of the echolocating bat, Eptesicusfuscus." *The Journal of the Acoustical Society of America* 85.3 (1989): 1348-1351.

[8] Tables of Physical & Chemical Constants (16th edition 1995). 2.1.4 Hygrometry. Kaye & Laby Online. Version 1.0 (2005) www.kayelaby.npl.co.uk

Evaluating the Effectiveness of a 3D-Immersive Environment on Learning Strategies for the Visual Control of Navigation

Nicholas Catoni* Timothy Parsons* Diego Morales* William Warren† Youssef Barhomi†

Brown University

ABSTRACT

We present a novel experiment to determine the effectiveness of a 3D-immersive environment in relation to learning strategies for the visual control of navigation. Specifically, we implement Fajen and Warren's experiment, where users navigate around obstacles towards a goal, [2] in the Brown's CAVE - a stationary immersive display system that projects information on the left, right, and front walls of a cube as well as the floor - to compare with their results from Brown's VENLab - a head-mounted display (HMD) system, where users have space to physically walk around while performing a given task. Through user studies and questionnaires, we determine the success of translating this experiment into the CAVE and evaluate the multiple issues found during this process.

Keywords: Visual navigation, virtual reality, obstacle avoidance.

1 BACKGROUND

1.1 Motivation and Theory

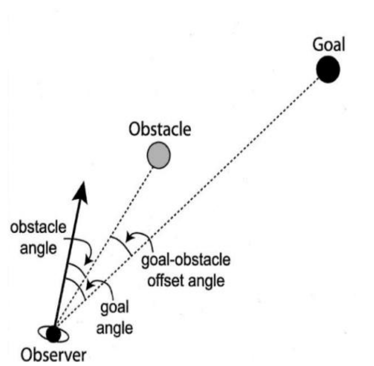

Figure 1: Fajen and Warren's egocentric reference frame [2].

We, as humans, do not have a great understanding of our visual system. Our collaborators - William Warren and Youssef Barhomi of Brown's Cognitive, Linguistic and Psychological Sciences department - focus their research on one area in particular: visual navigation. Specifically, how do humans navigate around obstacles towards a goal? When navigating a complex visual scene, humans constantly analyze their surroundings to make navigational decisions. Warren illustrates this phenomena in an agent-environment feedback loop. To summarize, the agent takes in visual information from the environment which it then turns into an action.

This action changes the agent's environment, causing the cycle to repeat [6]. Beyond this, Fajen and Warren built a model for human steering in regards to avoiding obstacles to reach a goal (Figure 1). In their experiment, they consider the angle between the observer and the obstacle, the angle between the observer and the goal, and the angular acceleration of the person. Using this data, they determine the minimal angle to take towards the goal to maintain a safe distance from the obstacle [2]. Similarly, Barhomi et al. vary how 'realistic' the environment feels to the user. In other words, where Fajen and Warren use basic poles (similar to our environment shown in Figure 2), Barhomi et al. build richer and textured environments with trees, grass, sky, etc. The idea is that at some point, the level of 'realism' will cause the data to vary from Fajen and Warren's model, and thus, help create a new model based on data gathered from an environment more representative of the world a human navigates through.

1.2 Related Work

The majority of visual navigation experiments have been performed with a stereo HMD [2, 3 5], such as Brown's VENLab, or on 3D-desktop displays [1, 4, 5]. However, there is a lack of literature on visual navigation in a 3D-immersive environment similar to Brown's CAVE, where a user is immersed in a cube-like structure with images projected on the walls, head-tracking glasses, and a joystick to navigate while standing in place. Similarly, there is not much comparison between a CAVE-like environment or HMDs or 3D-desktop displays. This experiment yields more insight into this area.

2 IMPLEMENTATION

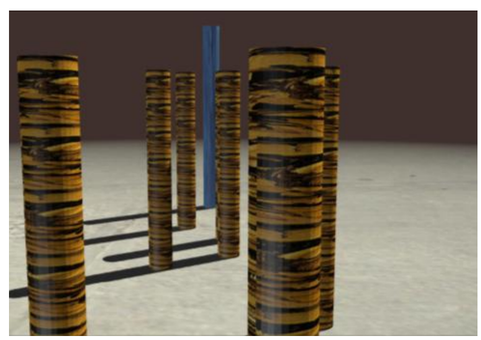

Figure 2: Screenshot of environment, implemented in Blender. Blue pole represents goal, brown-black poles represent obstacles.

The environment for our experiment - shown in Figure 2 - matches the specs of Fajen and Warren's experiment [2]. For a given trial, the user is instructed to reach the goal (blue pole) as quickly and as safely (meaning avoid the obstacles) as possible. They begin seven meters from the goal, with only the goal showing. Using the joystick, they move towards the goal. After the user moves one meter from the starting point, the obstacles (brown-black poles) instantaneously appear, forcing the user to navigate around them. The obstacles must appear in this way since Warren's model assumes as much [6]. Once the user navigates through the goal, they are reset to the starting point, and they can start again after a three-second delay. This process repeats for 96 total trials, in separate groups of 32 trials. We use eight unique arrays describing obstacle location, and for a given 32 trials, the user randomly navigates through each obstacle array forwards twice and backwards twice. For each trial, we track the users position in the environment.

3 Results

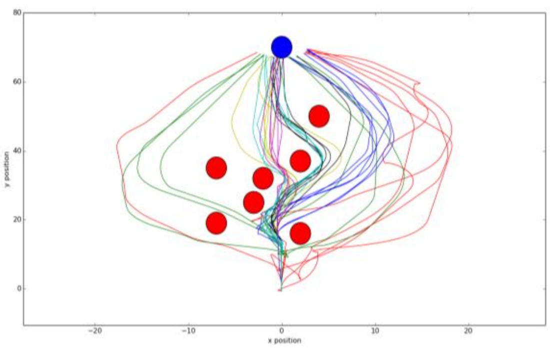

Figure 3: Sample plot of data from the seven users during the second iteration of the experiment.

The aforementioned implementation described in section 2 describes our final iteration. However, due to multiple users becoming ill, we were forced to go through two iterations of the experiment. Hypotheses as to why users became sick, as well as results stemming from the location data and user questionnaires, are discussed in the following sections.

3.1 First Iteration: User Sickness

With the first iteration of our experiment, five of six users became sick to the point of being unable to finish the experiment. This happened due to motion sickness and/or cybersickness. From both user comments as well as observations, we hypothesized the following reasons were attributing to the sickness: too many trials at once, a feeling of floating rather than walking through environment, obstacles seemed too large, virtual body size seemed incorrect, rotation speed too fast, camera height seemed to high, pole shadows blended into obstacle pole texture as well as background.

From this, we changed the trials from 2 groups of 48 trials to 3 groups of 32 trials with 2-3 minutes of rest between groups, lowered the camera height to an average height (1.7m), decreased the rotation speed, removed pole shadows, fixed the aspect ratios of the obstacle size and virtual body size, and generally, attempted to make the navigation feel more 'natural.'

3.2 Second Iteration: Users Finish

Following the changes, all seven users to perform our experiment finished the entire 96 trials. However, only two users reported no feeling of sickness throughout the experiment while five reported some level of discomfort due to the implementation of the environment. Furthermore, we asked one of the initial six users who felt very ill to retry the experiment following the changes. While she still felt some level of sickness (and did not finish due to time constraints), she referred to the second iteration as "A lot better" and could have finished the experiment if it had been possible.

Again, based on user feedback and our observations, we believe that some of the issues remaining are as follows: the obstacles popping up instantaneously is jarring, rotating around obstacles using the joystick feels unnatural when a sidestep would be more natural, virtual body size still feels incorrect, standing in place feels unnatural (users want to move), users sometimes notice non-stereo images in the periphery due to the CAVE head-tracking glasses coverage limitation. Essentially, the environment still just doesn't feel 'natural,' as users want to navigate as they normally would rather than being limited by the controls of a joystick while standing in place.

3.3 User Monotony

Aside from sickness affecting users, we hypothesize that the monotony of the experiment led to inefficient paths by at least one user (the red lines in Figure 3). During as well as after the experiment, this user noted that the navigation task was "boring" and asked questions in relation to how large (i.e. where are the edges) of the virtual world. On all 16 arrays, this users consistently chose paths that strayed far outside the other users' paths. Furthermore, other users stated how they actively chose a certain path to try and make the experiment more interesting.

3.4 Prior Virtual Reality Experience

Prior to the experiment, we hypothesized that users with more virtual reality (VR) experience would have less issues with the navigation task. However, out of eight users who noted they felt sick during the experiment, seven had previous interaction with VR. Furthermore, of the three users with no VR experience, two did not feel any sickness.

4 Conclusion

This experiment yielded insight into the comparison between two VR environments, namely the CAVE and the VENLab at Brown. From this, we concluded that the CAVE may not be the best environment for Fajen and Warren's visual navigation experiment [2] to a variety of reasons, such as issues with rotating instead of the ability to sidestep, virtual body size, CAVE-specific issues (inability to walk, limited visual field coverage of the glasses). The majority of these issues don't exist in an immersive system such as the VENLab since a user is wearing a HMD while walking around the obstacles. This setup easily ensures a camera height in relation to the user's height, there are no visual field coverage issues with an HMD, the user can sidestep instead of rotate around obstacles, and arguably most importantly, the user is actually walking rather than standing in place.

Furthermore, this experiment showed that the monotony of the navigation task may impact a given user's path, making it more inefficient. Also, that the VR experience of a user may not matter in terms of whether or not they get sick.

The main open problem remaining for is to further implement the experiment in the CAVE with all users feeling comfortable. We hypothesize that this could be accomplished by allowing the user to walk-in-place rather than stand as well as building a more 'realistic' environment, as described in section 1.1.

References

[1] Youssef Barhomi, Abigail Yanke, David Mely, Stephane Bonneaud, William Warren, and Thomas Serre. A data-driven approach to learning strategies for the 1visual control of navigation. Poster. 2014.

[2] Brett R. Fajen and William H. Warren. Behavioral dynamics of steering, obstacle avoidance, and route selection. *Journal of Experimental Psychology: Human Perception & Performance*. 2003.

[3] Beatriz Ray and Mariano Alcaniz. Research in Neuroscience and Virtual Reality. 2010.

[4] Bernhard E. Riecke, Henricus A. H. C. van Veen, and Heinrich H. Bulthoff. Visual Homing Is Possible Without Landmarks: A Path Integration Study in Virtual Reality. *Presence*. 2002.

[5] Beatriz Sousa Santos, Paulo Dias, Angela Pimentel, Jan-Willem Baggerman, Carlos Ferreira, Samuel Silva, and Joaquim Madeira. Head-mounted display versus desktop for 3D navigation in virtual reality: a user study. *Multimedia Tools and Applications*. 2009.

[6] William H. Warren. The dynamics of perception and action. *Psychological Review*. 2006.

Evaluating Attributes of Virtual Reality Displays for Interacting with Isosurface Renderings

Johannes Novotny* Wesley Miller * David H. Laidlaw * Wallace S. Lages† Doug A. Bowman †
John J. Socha † Bireswar Laha ‡

Brown University, Virginia Tech, Stony Brook University

ABSTRACT

We evaluated the influence of two fidelity attributes of CAVE VR environments on the performance of users in complex isosurface exploration tasks. Previous studies have indicated that a higher field of regard (FOR) and the use of stereoscopic displays significantly improves user performance in specific use cases. In our work we repeat one of these experiments in order to generalize the findings over multiple CAVE setups. The results of our user study confirm the positive effect of higher fidelity setups on user performance in general, but differ in various aspects from the baseline data.

Keywords: Virtual reality, CAVE, data analysis, human-computer interaction, evaluation.

1 INTRODUCTION

In this work we evaluate how two fidelity aspects of a CAVE virtual environment system affect the users in search and exploration tasks of 3D isosurface data. The analysis of complex three-dimensional data is a common task in many scientific domains including medicine, biology and engineering. While specialized 2D desktop applications are available for many of these tasks, recent advances in display hardware (e.g. 3D displays and head mounted displays) make it feasible to use Virtual Reality (VR) environments for data analysis and exploration. VR setups enable researchers to access data visualizations in an immersive way which may lead to more insight into the 3D structure compared to traditional desktop systems [3].

While VR systems often differ in their basic construction their fidelity attributes can be described through a set of distinct components. These components include the field of view, field of regard (FOR), stereoscopy (ST), resolution, and head-tracking (HT) capabilities of a system. Single and combined effects of fidelity components on user behavior have been found in several controlled user studies within single VR setups [4] and across multiple systems [1][3]. Due to the uncommonness of CAVE environments, multi-system studies often compare very different VR setup types (e.g. fish tank, head mounted display, CAVE)[3]. To address this limitation we partially repeat a recent evaluation of VR fidelity components in our own CAVE environment. The study by Laha et al. [2] reports that high fidelity settings for FOR, ST, and HT components of a CAVE system increase user performance in specific search and spatial judgment tasks on isosurface volume visualizations. Their results indicate that higher fidelity settings not only affect task correctness but also task completion time as well as subjective difficulty and confidence ratings of the users.

*e-mail: novotny/wmiller/dhl@cs.brown.edu
†e-mail: wlages/dbowman/jjsocha@vt.edu
‡e-mail: bireswar.laha@stonybrook.edu

In this work we directly extend the study of Laha et al. by validating their experiment for the FOR and ST components. Our results confirm that both components decrease the task completion times of users, but we report notable differences in the answer grade metric and are not able to confirm all of the previously published individual and combined effects.

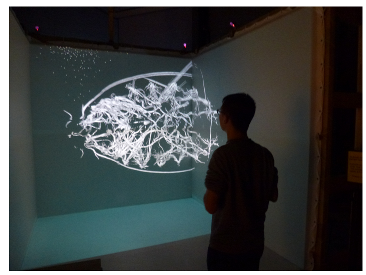

Figure 1: A user during a training task in the 270 stereo configuration

2 EXPERIMENT

In collaboration with the authors, we have conducted an experiment to evaluate the field of regard and stereo components of our CAVE system, that resembles the previous experiment as close as possible.

To avoid unnecessary variation we used two manually segmented micro-CT data sets used in the original experiment. The data sets contained the isosurface representations of beetle tracheal systems of two specimens of the *Pterostichus* (Fig. 1) and the *Platynus* genus.

For this experiment we used a four-screen CAVE setup with three rear-projected 8' by 8' walls and a top-projected 8' by 4.5' floor with screen resolutions of 1280×1280 and 1280×720. Stereo view was supported through active shutter glasses. An OptiTrack system was used to track positions and orientations of the users head and an Aimon PS Elite wand (9 buttons & joystick). To prevent visual differences we configured the software framework of the original experiment, VRUI and MeshViewer, to work with our CAVE setup. One of our collaborators personally visited our setup to verify that it resembles the original environment as close as possible.

2.1 Tasks and Procedure

In order to obtain comparable results we use the set of tasks designed by Laha et al. [2] with minor modifications to accommodate for the different setup. The tasks aim to cover a wide range of use cases in various scientific domains and are grouped into five categories: searching and counting (s/s,c), spatial judgment (sj), quantitative estimation (qe), shape description (sd), and pattern recognition (pr). Table 1 shows the number of tasks in each category. Each task requires the user to examine the given trachea data set for specific features, e.g. single tracheal tubes for size estimation (qe) or tracing (sj). For detailed task descriptions please refer to the appendix in the paper of Laha et al. [2].

We evaluated two settings for each fidelity component. For the field of regard, defined as the horizontal angle of screen space surrounding the user, we tested 90°(front and floor screen) and 270°(all four screens) settings. For the stereo setting we evaluated stereoscopic and monoscopic view, leading to four independent setting combinations.

For our user study we recruited 11 (7 male, 4 female) volunteers from the graduate student body. None of the participants reported prior experience with VR environments and isosurface exploration. Two participants were used as pilot users, the others were divided into the four setting groups for the main study. While all users passed the initial spatial ability test, one user had to be dismissed as an outlier based on self-reported confidence and difficulty scores.

The experiment procedure consisted of several steps. The introduction included the signing of the IRB approved informed consent form as well as a background questionnaire and a spatial ability test. After that each participant was introduced to the used data sets and the CAVE hardware. A set of five training tasks allowed users to familiarize themselves with the environment (Fig. 1). After a short break participants were asked to complete the 15 tasks of the main study. Graded answers, completion time as well as subjective scores for task difficulty and confidence in the answer were collected for each task. Finally participants completed the study by filling out a post-experiment questionnaire and were given the chance to ask questions and give feedback in a free-form interview.

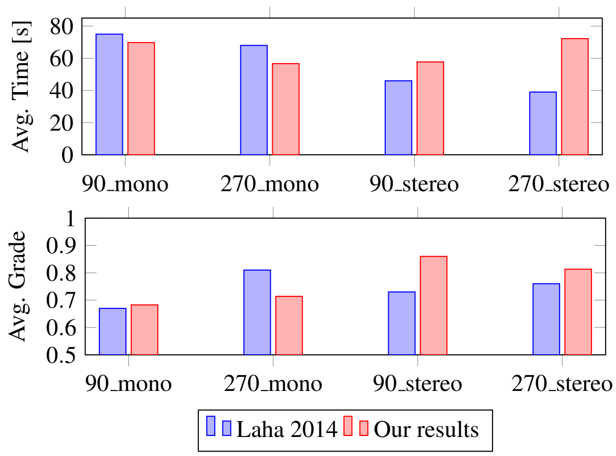

Figure 2: Comparison of average completion times and average grades within the four evaluated setting groups

2.2 Results and Discussion

Figure 2.1 shows a general overview of our study result for the average task completion time and answer grades in comparison to previous results. We report that a high FOR and reduce the time required to complete tasks and increases their grades among users in the monoscopic setup, matching our expectations. However, in the stereo settings we observe an opposing trend caused by a single outlier in the *270°stereo* group.

One goal of this study was to confirm statistically significant findings reported by Laha et al. in our own setup. Table 1 lists the previously observed individual and combined effects for individual tasks. Out of 17 effects we were able to strongly confirm six effects and found similar trends for four more. In the remaining cases we either found equal results or opposed trends between the tested settings. The two unconfirmed effects for the stereo setup in the grades metric (Grade-ST) result from equal results in the mono and stereo cases. We were able to find evidence for a majority of the stereo effects.

Table 1: Evaluation of significant effects and interactions found by Laha et al c[2]. A cross (X) denotes a significant main effects of a component in a given metric while connected circles (O) indicate significant component interactions. The color encodes which effects and interactions were confirmed by our own results, where green indicates strong confirmation, orange a weak similar trend and red no confirmation

	Grade		Time		Difficulty		Confidence	
	FOR	ST	FOR	ST	FOR	ST	FOR	ST
T1 (s,c)								
T4 (s,c)		X						
T5 (s)			O	O				
T8 (s)	O	O		X				
T9 (s)		X		X	O	O		
T2 (sj)			X					
T6 (sj)		X						
T10 (sj)				X				
T11 (sj)	X				O	O		
T13 (sj)								
T14 (sj)		X				X		
T3 (qe)								
T15 (qe)								
T7 (sd)				X			O	O
T12 (pr)								

In general, we expected to confirm a higher number of the previous effects with our own user study. One significant difference between the two experiments is that our study evaluated a very low number of participants per group, which causes our results to be more sensitive to outliers. Additionally, inherent differences of the used CAVE systems (e.g. size, resolution, contrast, tracking accuracy and latency) could have more influence on the user performance than we anticipated. The higher task completion times of our participants might have been caused by a different motivation for volunteering in the experiment. Our users volunteered out of interest in the CAVE environment alone with no additional incentive, while users of the previous study had to participate in user studies for class credit.

3 Conclusion

We presented the results of controlled user study evaluating two fidelity components of CAVE environments, based on a previous experiment. While we could not entirely confirm the preceding results, we still observed positive effects of high fidelity components on task performance in multiple cases. However, we will need to increase the number of study participants in order to obtain more significant results.

References

[1] C. Demiralp, C. D. Jackson, D. B. Karelitz, S. Zhang, and D. H. Laidlaw. Cave and fishtank virtual-reality displays: A qualitative and quantitative comparison. *IEEE Trans. Vis. Comput. Graphics*, 12(3):323–330, May 2006.

[2] B. Laha, D. Bowman, and J. Socha. Effects of vr system fidelity on analyzing isosurface visualization of volume datasets. *IEEE Trans. Vis. Comput. Graphics*, 20(4):513–522, April 2014.

[3] Prabhat, A. Forsberg, M. Katzourin, K. Wharton, and M. Slater. A comparative study of desktop, fishtank, and cave systems for the exploration of volume rendered confocal data sets. *IEEE Trans. Vis. Comput. Graphics*, 14(3):551–563, May 2008.

[4] E. Ragan, R. Kopper, P. Schuchardt, and D. Bowman. Studying the effects of stereo, head tracking, and field of regard on a small-scale spatial judgment task. *IEEE Trans. Vis. Comput. Graphics*, 19(5):886–896, May 2013.

Encoding Qualitative Uncertainty in Visualizations for Archaeological Research

Richard Nguyen* Scott Houde† Sheila Bonde‡

Brown University

ABSTRACT

Architectural historians and archaeologists when presented with visualizations are concerned with the presentation of images that suggest certainty in the presence of unknown or questionable data. Data in these fields draw from qualitative sources such as artistic depictions and literature, as opposed to vector fields or MRI data that have been the focus of much of uncertainty visualization research. We present a study of a set of modelling techniques to visualize qualitative uncertainty while informing excavation planning.

Keywords: qualitative uncertainty, visualization, archaeology, architectural history, Bourgfontaine

1 INTRODUCTION

Finding ways to display uncertainty is one of the great problems in visualization [1], but work in visualizing qualitative uncertainty is sparse. In particular, architectural historians and archaeologists are concerned with the presentation of images that suggest certainty in the presence of unknown or questionable data [3]. This research in uncertainty is distinguished from previous efforts in the nature of its data. Researchers in the humanities typically deal with data from qualitative sources such as artistic depictions and literature, as opposed to vector fields or MRI data [2].

Efforts have been made in modelling techniques to visualize this qualitative uncertainty through color and transparency and time-variant featurization [7, 6, 5]. However these efforts have typically been made post-excavation of their sites. In response, we present a study of a set of modelling techniques to visualize qualitative uncertainty while informing excavation planning.

2 UNCERTAINTY AT BOURGFONTAINE

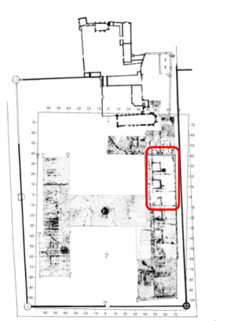

Figure 1: GPR data within architectural drawing of existing foundation at Bourgfontaine Charterhouse.

*e-mail: rnguyen@cs.brown.edu
†e-mail: scott_houde@brown.edu
‡e-mail: sheila_bonde@brown.edu

Located in northeastern France, the Bourgfontaine Charterhouse is one of many Carthusian monastaries. While the site is still in the process of excavation, we take particular interest in three cells in the cloister of the charterhouse. These three cells are planned for excavation in the near future and motivate our study.

Geologists have conducted experiments on the site with ground-penetrating radar (GPR) to produce slices of varying depths from 0.7 meters to 1.35 meters [4]. Shown in Figure 1, they have provided the most representative slice for potential foundations of the charterhouse.

Varying sources like this, existing foundations above-ground at the site, artistic depictions, and literature were combined into a 3D model created in SketchUp. Uncertainty associated with each of the sources was then encoded using three different techniques on the model: color, color and transparency, and texture.

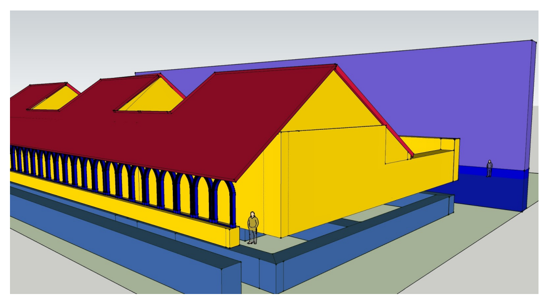

Figure 2: Model of cloister cells using color to encode uncertainty of sources and features.

3 EVALUATION

A user survey was conducted amongst art history and archaeology graduate students and faculty to evaluate the three techniques at visualizing qualitative uncertainty. Surveyors were presented with each technique on the model and allowed to interact with it by panning, zooming, and orienting the view. After interacting with an encoding technique, surveyors were prompted with four statements:

1. The 3D rendering is simple to interpret.
2. The difference in uncertainty among features is clear.
3. The model is clear in representing uncertainties.
4. I would find a model like this one to be useful in my work or research.

The surveyors responded to these statements with a Likert scale of answers "Strongly Disagree", "Disagree", "Neutral", "Agree", and "Strongly Agree". Following evaluation of each individual encoding technique, surveyors were asked to rank each in order of preference.

Table 1: Uncertainty Encoding Technique Ranking (n = 21)

Encoding	Sum	Avg.	Mode
Color	36	1.71	1
Color & Transparency	39	1.86	1
Texture	51	2.43	3

3.1 Results

Responses to the four Likert scale-evaluated statements are summarized in Figure 3 into four areas of concern: 1) Simplicity, 2) Difference, 3) Clarity, and 4) Utility.

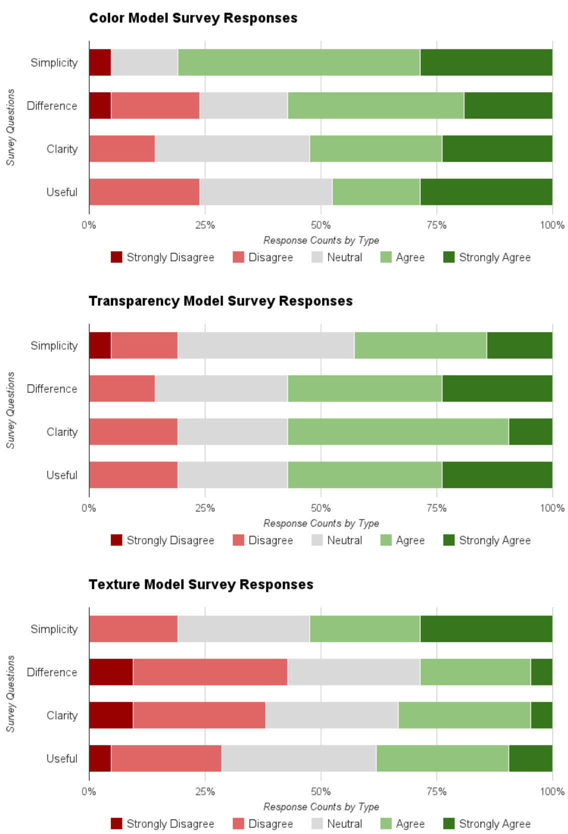

Figure 3: Survey results for each uncertainty encoding technique.

The results of the survey demonstrate a need for 3D models in archaeological research, with 52% of surveyors responding positively (”Strongly Agree” or ”Agree”) to Utility. If extended to include Neutral answers, this percentage comes to 79%. Anecdotally, the process of modelling the site has led to dig site boundary adjustments for the three cloister cells at Bourgfontaine. It is clear that future work in pre-excavation modelling has the potential to bear fruit for archaeologists.

In preference, Texture was clearly the least preferred out of the three encoding techniques. This was reflected in a substantial percentage of negative responses (”Strongly Disagree” and ”Disagree”) in the four Likert scale-evaluated statements given for Texture. 43% and 39% of surveyors respectively found that this encoding technique did not clearly visualize or distinguish between the uncertainties of sources and the features. This can be attrributed to the lack of contrast between textures and that of the bright, saturated primary colors of the Color and Color & Transparency encodings.

The Color and Color & Transparency encodings received similar feedback from surveyors, with close to half of all responses being positive for each statement. The distinction between the two encoding techniques seem to be in their perceived simplicity: 81% of surveyors found the Color encoding to be simple over 42.9% for the Color & Transparency encoding.

Although these two encoding methods were by and far the most preferred among surveyors, comments gleaned from the survey suggest the limited nature of using color to encode qualitative uncertainty. One surveyor noted that distinguishing between blue and purple hues in the encoding was difficult, while another pointed at the lack of support for color-blind individuals. These comments suggest that as more sources of uncertainty are encorporated, the contrastive benefit that colors have over texture may be diminished.

4 Conclusion

Though work and literature in visualizing qualitative uncertainty is sparse, a survey of graduate students and faculty in art history and architecture notes the utility of such work, especially pre-excavation. Surveyors preferred encodings with more contrast to distinguish levels of uncertainty between different features and sources. However an encoding technique relying solely on color is inherently limited in quantifying larger numbers of sources of uncertainty. In the future, work in quantifying qualitative uncertainty can benefit from exploration into other encoding techniques, such as displacement or distortion of uncertain features or utilizing photorealism.

Acknowledgements

The authors wish to thank David Laidlaw and Steve Gomez, instructor and teaching assistant respectively for CSCI2370 Interdisciplinary Scientific Visualization at Brown University.

References

[1] C. Johnson. Top scientific visualization research problems. *Computer graphics and applications, IEEE*, 24(4):13–17, 2004.

[2] C. R. Johnson and A. Sanderson. A next step: Visualizing errors and uncertainty. *Computer Graphics and Applications, IEEE*, 23(5):6–10, 2003.

[3] P. Miller and J. Richards. The good, the bad, and the downright misleading: archaeological adoption of computer visualisation. *BAR INTERNATIONAL SERIES*, 600:19–19, 1995.

[4] A. Saintenoy, F. Rejiba, E. Leger, S. Bonde, and C. Maines. Ground penetrating radar prospection at the charterhouse of bourgfontaine. In *Ground Penetrating Radar (GPR), 2014 15th International Conference on*, pages 83–86, June 2014.

[5] M. Sifniotis. *Representing archaeological uncertainty in cultural informatics*. PhD thesis, University of Sussex, 2012.

[6] M. Sifniotis, P. Watten, K. Mania, and M. White. Influencing factors on the visualisation of archaeological uncertainty. In *Proceedings of the 8th International conference on Virtual Reality, Archaeology and Intelligent Cultural Heritage*, pages 79–85. Eurographics Association, 2007.

[7] T. Zuk, S. Carpendale, and W. Glanzman. Visualizing temporal uncertainty in 3d virtual. In *Proceedings of the 6th International Symposium on*, pages 99–106.

Visualizing Qualitative Uncertainty in Archaeological Research

Scott Houde* Richard Nyugen†

Brown University

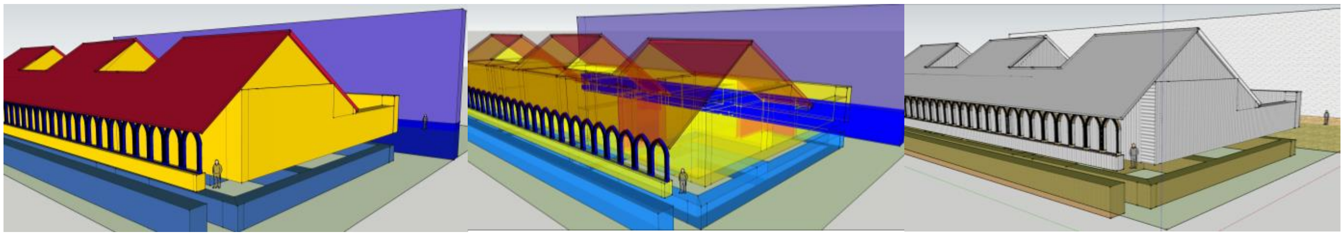

Figure 1: The three variants of the model generated in this project: from left to right, color, transparency, and texture. Online at http://tinyurl.com/og9vo5p

Abstract

This project explored three different ways to represent qualitative uncertainty in visualizations for archaeological research and will provide insights obtained from user feedback. Our 3D models brought together existing 2D information; site surveys, one slice of ground penetrating radar (GPR) data, comparative sites and artistic sources from a Carthusian monastery in northern France. Each source was assigned a qualitative value of certainty which was then assigned a color or texture for the model. Finally students and faculty from the fields of architectural history and archaeology interacted with the models and answered a short survey. Our methods and results from this survey are presented here.

Keywords: qualitative data, uncertainty visualizations, 3D models, architectural history, archaeology, Bourgfontaine.

1 Introduction

Finding better ways to represent uncertainty in 3D visualizations is an unresolved problem that many have considered. For our work we joined with an architectural historian and archaeologist to present qualitative uncertainty. We feel that this area has been under explored and that architectural historians and archaeologists could greatly benefit by utilizing uncertainty visualizations.

Our collaborator has done some architectural uncertainty previously but her focus was not on evaluating the visualizations. Her current research site is the charterhouse of Bourgfontaine, a Carthusian monastery north-east of Paris that was built in 1323-1325. The site survives with some structures intact but many in ruin. She has collaborated in the past with geologists at this site who have done extensive work with GPR and they have provided her with data outputs, including a 2D plan.

2 Related Work

Prior work on architectural uncertainty was done by our collaborator for the monastic church of Saint Jean-des-Vignes [1]. We built on that work by having more visual levels with three variants and conducting a user survey to determine preference. The GPR data slice was a result of work done by A. Saintenoy et al. [2] which was performed at Bourgfontaine in 2013 and 2014. We focused not on displaying the GPR cube data instead we extrapolated out actual foundations and walls above them from the provided slice.

Much work has been done on uncertainty in 3D visualizations; a call for a framework for doing so was done by Johnson and Sanderson in 2003 [3]. However, to our knowledge, no other study has focused on qualitative uncertainty which is what was done in this research.

3 Methods

Our work took various pieces of data and brought them together in one 3D model of the three cloister cells that are going to be excavated in the summer of 2015. For quantitative data, we have a site survey our collaborator performed of the remaining structures, with one slice of GPR data added in scale. These two pieces gave us a sense of the foundation of the three cloister cells we modeled. To construct the walls and roofs, we needed to use artistic sources. One particular image, a painting by Louis Licherie of the 17th century [4] was our primary source, although we also consulted images and plans of comparative existing sites.

We used free modeling software, SketchUp, since the intention was for our models to be used in the field by any user and we needed something fairly simple for all users to work with. We constructed a master model based on a scan of the site survey with GPR data. This was adjusted to be in scale so that accurate measurements could be obtained. GPR results were extruded down from the imaginary ground plane while the existing wall foundation, which is only ankle high, was both pulled up above ground and extruded down. Using the painting as reference, the fortification wall was added above the existing visible foundation and the cloister walls above the GPR results. Next cloister walls that should have existed but have no GPR foundation results were added. The arches along the cloister alley are interesting in that they are highly certain since most of the arches have been recovered from the site while the wall they sit on is less certain and the roof they supported is very uncertain. In total, we had seven classes of elements, each of which was given a certainty between 0 and 1 by our collaborator. These certainty scores are the qualitative element since they were only based on her experience and the source of the information. For example, we assigned a value of .75 to the GPR data which records existing though hidden foundations, while we gave the roof a value of .25. (We know there must have been a roof but the exact supports and pitch are unknown.)

After we constructed the model, we generated three variants: a color, a transparent and a textured version. In the color version, each color indicates a different level. The colors were chosen using a triad rule on a color wheel where the most common level was the primary color. (Green was avoided since our ground plane was green.) For the transparency version, the opacity of each class was set to its certainty level that had the same colors as the color variant. The texture variant used simple neutral colored textures that were chosen from SketchUp.

* e-mail: Scott_Houde@brown.edu

† e-mail: Richard_Nyugen@brown.edu

Once the model variants were built, we loaded the model up to Sketchfab (https://sketchfab.com/) which allowed us to share the models via a web browser. We constructed a survey which provided some background information on the site and our project and then presented responders with one model at a time along with four Likert scale (1-5) questions and space for comments. On each model page we had a short paragraph, a link to the online interactive model, a picture of the model and a legend explaining each color/texture. We also collected model rank (1 – 3), final comments, basic demographic data and comparative certainty level information for our collaborator.

4 Discussion

We received 21 responses to our survey and all were included in our analysis. The responses were evenly split between male and female (10 each plus 1 choose not to respond to any demographic questions) but the respondents were weighted towards the 18-25 age band, 52.4%, with 26-35, 33.3%, and 46-55, 9.5%, making up the balance. There was a slight difference between the 18-25 age group and remaining groups for model preference. The younger group had a small preference for the color model while the older groups were equally split between color and transparency when ranking the models. These differences are very slight and do not appear to be statistically significant. Small differences were again noted between the male and female respondents. Females were equally split between color and transparency while males very slightly preferred transparency. The variance was slightly higher for female than male but again these differences between groups do not appear to be statistically significant.

Looking at all the responses together the mode rank for Color and Transparency were both 1, but the average rank for Color was 1.71 versus 1.86 for Transparency. Texture had a mode rank of 3 and an average rank of 2.43 and was clearly least preferred.

For each model we asked four Likert scale questions. These questions can be reduced to one word each by which they are labeled as in Figures 2 and 3. The questions were: 1) The 3D rendering is simple to interpret (*simplicity*); 2) The difference in uncertainty among features is clear (*difference*); 3) The model is clear in representing uncertainties (*clarity*); 4) I would find a model like this one to be useful in my work or research (*useful*).

4.1 Texture Model

The texture variant was by far the least liked of the three. One respondent commented *"The textures are extremely annoying to look at, and make the model seem as if it was made up out of thin air."* Those responding also felt that some of the textures were too similar making differences hard to distinguish. We purposefully choose neutral simple textures but unfortunately some of these textures indicated specific building materials and some responding knew that, which based on their comments, compounded their confusion.

4.2 Color and Transparency Models

The color and transparency variants offer some more interesting analysis. Color was a slight favorite but if you look at the question responses in Figures 2 and 3 you can see that transparency was equal to or better than color in every question except simplicity. 81% agreed or strongly agreed that the color model was simple. Simple might not always be the best as transparency was as good or better in every other question. In particular transparency has a lower variance in the difference question. I think some insight into why is illustrated in this comment. *"This model* [transparency] *is more helpful than the last, but is much more effective in the 3D version than in the image above"*.

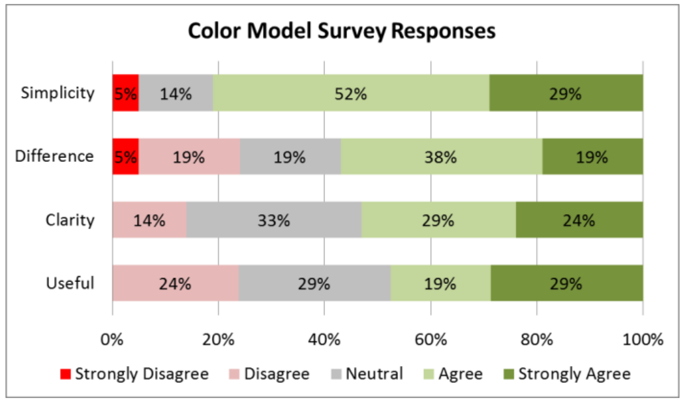

Figure 2: Survey responses to questions about the color model,

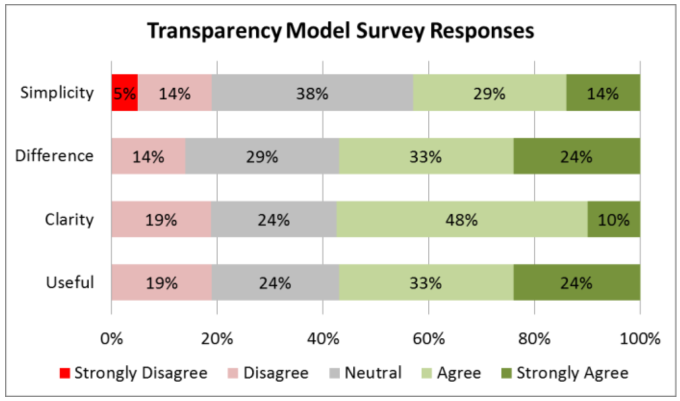

Figure 3: Survey responses to questions about the transparency model, N = 21

5 Conclusion

Our respondents showed a clear preference for the color and transparency variants. Color was overwhelmingly chosen as being simpler to interpret but we had several comments about how much better the transparency was as an interactive 3D model versus the transparency image. Interestingly only one respondent commented on how the color and transparency variants would be difficult for a color blind user to interpret. Despite their preference for these two variants several users found the colors too bright and garish. We strongly recommend a more subdued color palette than the one we used.

Some open questions are: Would it be possible to display 7, or more, classes and be color-blind friendly? Use of textures would seem one way to be safe for color blind users. How many levels would display the most information clearly? How best to get across the most information without overwhelming the reader

We are extremely gratified that the act of even building these models proved helpful for our collaborator as she had several insights upon seeing the 3D model take shape. One of these has resulted in her adjusting the locations of trenches she will excavate in the summer of 2015.

References

[1] Sheila Bonde, Clark Maines et al. "The Virtual Monastery: Re-Presenting Time, Human Movement, and Uncertainty at Saint-Jean-des-Vignes, Soissons." Visual Resources 25.4 (2009): 363-377.

[2] A. Saintenoy, et al. "Ground Penetrating Radar prospection at the charterhouse of Bourgfontaine." Ground Penetrating Radar (GPR), 2014 15th International Conference on. IEEE, 2014.

[3] Chris R. Johnson and A. Sanderson. "A next step: Visualizing errors and uncertainty." Computer Graphics and Applications, IEEE 23.5 (2003): 6-10.

[4] Louis Licherie. *Chartreuse de Bourgfontaine*. 1687.

Evaluating Effectiveness of Human Navigational Experimentation in an Immersive CAVE^{TM} Environment

Timothy Parsons* Nicholas Catoni† Diego Morales×

Brown University

ABSTRACT

We conducted a User Study to evaluate the effectiveness of using a CAVE^{TM} for experimentation in the field of human navigation. We re-implemented a psychophysics experiment done by Professor William Warren of Brown University using a Head-Mounted Display (HMD) and compared quantitative and qualitative results. The conclusion our group came to is that a CAVE^{TM} does not allow for the level of immersion a HMD provides for navigational tasks, and therefore is an inferior, and potentially inadequate, tool for evaluating human navigational responses.

Keywords: Navigation, Psychophysics, Immersive 3D Virtual Reality.

1 INTRODUCTION

Much is unknown about how the human brain functions, and the level of knowledge remains similarly poor when we look at the more specific branch of human visual perception as it pertains to navigation. In the area of cognitive science, specifically psychophysics, much has been done to study how humans navigate. Knowing how humans navigate could be very useful in other fields too, such as robotics. If we know how humans navigate, we can tell robots to navigate the same way, solving a problem that is currently very much unsolved and complex.

Finding the right environment to hold navigational experiments and get accurate and useful results is unfortunately difficult. Ideally, we would perform navigational experiments in the real-world, where we are assured that stimuli will elicit natural responses. However, gaining the level of control necessary for a useful navigational experiment in the real-world is hard. Virtual Reality offers us the possibility for much more control as well as possibly similar responses to stimuli. Prior to the last two decades, low-resolution 2D images were used to study low-level visual responses to stimuli [3]. Only in the last two decades has technology in the field of Virtual Reality (VR) emerged that allows for more realistic simulation of stimuli [3,4].

There are several examples of navigational experiments in VR, but the vast majority use navigational tasks to evaluate different aspects of VR. It is important to know which environments are eliciting natural responses for navigational experimentation, and therefore we are more interested in how different VR experiments affect the level of 'naturalness' of humans' navigational responses. To evaluate this, we want a baseline for low-level navigational responses. For this, we look to Warren et. al 2003. Professor Warren of Brown's CLPS department performed an experiment where users (wearing a HMD) had to navigate towards a goal and avoid obstacles. They are able to then test a model they developed against these results [1]. Our objective was to recreate this experiment in another environment – Brown's CAVE^{TM}, and directly compare the effectiveness of the CAVE^{TM} and the HMD.

* e-mail: taparson@cs.brown.edu
† e-mail: ncatoni@cs.brown.edu
× e-mail: dmorales@cs.brown.edu

2 THEORY

Our hypothesis is that a VR environment with higher fidelity combined with a mode of virtual navigation that more closely mimics how humans naturally navigate will give a more natural visual response. When we say fidelity, we mean quantitatively (resolution, Field of View, etc.), as well as qualitatively (a user's 'sense of presence'). The CAVE^{TM} wins in some categories of quantitative fidelity such as field of view and the ability to see one's limbs, but lacks the ability to move around large distances naturally as with a HMD in a large room. Because of this limitation, we chose to implement a flying-in-place mode of navigation for our experiment in the CAVE^{TM}. The user uses the joystick to rotate the world around him/her and to move forward. We hope to see a difference quantitatively and/or qualitatively between our experiment and Warren's so we can speak to the effects of system fidelity and 'naturalness' of motion - which we will call 'proprioceptive immersion – on human navigational response.

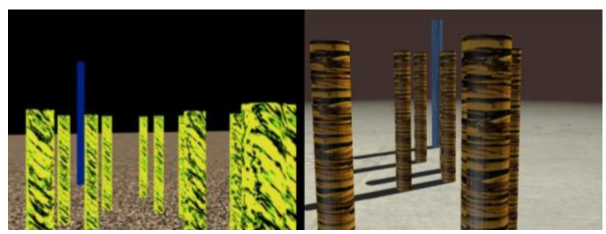

Figure 1: Left: Warren's experiment. Right: Our implementation.

2.1 Procedure

The procedure for the experiment from Warren et. al that we aimed to recreate is as follows: A plain atmosphere (grayish floor, black background) is presented to the user with a blue cylindrical goal in the distance. The user walks towards the goal – after travelling a meter, seven obstacles appear. The user must navigate around these obstacles and proceed to the goal. We had seven users complete the experiment, and they each ran ninety-six trials [1]. The users were not paid and were volunteers. Between trials, the user was shown the blank background with the goal for a few seconds before proceeding with the next trial. We implemented this using the Blender Game Engine^{TM}, which is supported by Brown's CAVE^{TM}.

We had sixteen possible locations for the obstacles in the scene – eight unique sets of locations backwards and forwards. We duplicated these six times to give us our ninety-six trials. After gathering data from these users, we were able to plot their navigational trajectories for each unique set of obstacle locations. We also asked for gender, handedness, previous video game experience, level of qualitative 'sense of naturalness', and degree to which the user felt sick. Using this data, we were able to come up with some qualitative and quantitative results.

3 DISCUSSION

Our first few users we ran had to be thrown out because of an intense feeling of nausea. We re-evaluated our software and realized that we had made some mistakes with scale and positioning of the camera, which in turn made users feel like the obstacles were very large and that they were flying many feet above the ground. An interesting result of moving from the Desktop (where we developed in Blender) to Immersive VR is that the user then used their bodies as a baseline for scale, which does not exist on the Desktop.

We fixed these scale issues and had a user re-try the software (but did not include the user in the final data). The user noted that he felt much less sick, which leads us to conclude that appearing to fly *over* the environment can lead to intense motion sickness. We will now discuss results we found after running the other seven users.

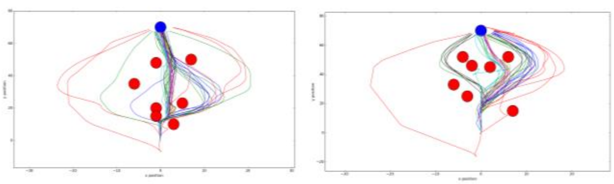

Figure 2: a) An example of a bored user, b) an example of users taking consistent routes to the goal

3.1 Consistency with Warren

Professor Warren's model takes into account distance from the obstacle, velocity of the user, angle of approach to the goal, and angle of approach to the obstacle to generate a set of possible predicted paths that humans might navigate. One interesting result he finds is that within a user, generally the same path is taken for similar obstacle positions [1]. Looking at Figure 2 a) and b), one can see that users (separated by color) within trials seem to take similar trajectories, but different users do not necessarily take similar paths, which agrees with Warren et. al's findings.

3.2 Sickness

After we fixed the aforementioned problems with our system, we still had some sickness. In fact, three out of seven users felt some level of discomfort running our system. We have a few ideas as to why this is happening. One is that the user is flying through the environment. Clare Regan notes in a 1995 paper that the conflict between the user's visual system (motion) and vestibular system (static) is a major contributing factor in motion sickness [4]. In our experiment, the environment translates and rotates without any vestibular input for the user. This is a direct result of the mode of navigation (flying) chosen. A few users specifically noted this as a reason they thought they felt sick.

One interesting potential reason for feeling sick that a user brought up was that he could see obstacles out of his peripheries outside the scope of the glasses, which resulted in the obstacles not being in stereo. This suggests possibly adding blinders to the sides of the glasses for use in a CAVE™ to avoid motion sickness. Another interesting result we found is that there seemed to be little to no correlation between video game use and motion sickness – some users felt sick that had video game experience, and some without experience felt no motion sickness at all.

3.3 Boredom

The reader may notice the seemingly erratic user in red in both of the data plots. This user spoke openly about how bored he was and began to take more 'interesting' routes to the goal despite being prompted to take the shortest path to the goal. This resulted in erratic and unrealistic paths. Our users were not paid, which gave them little incentive to complete the trials. However, even an incentive might not be enough to cure boredom in these monotonous navigational experiments. An idea we had to limit the boredom would be to show a timer during these tasks to give the user some incentive within trials to finish more quickly. Perhaps making it more game-like – having a high-score time that the user is trying to beat with penalties for colliding with obstacles – could help as well. A natural extension to our work would be to introduce more levels of realism into the scenery, something Professor Youssef Barhomi of Brown's CLPS department is interested in. In this extension, the scenery could be changed between trials to make the trials more interesting and engaging for the users.

3.4 Sense of Naturalness

One aspect we were interested in was the users 'sense of naturalness' in completing these tasks. This can be thought of as a qualitative level of immersion. We hypothesized that this 'sense of naturalness' would be directly correlated with the level of naturalness in the user's navigational response. Users generally felt little 'sense of naturalness' in the environment, most citing the method of navigation (flying) and feeling sick as reasons for this. This is a problem intrinsic to the CAVE™ – mimicking natural navigation through a large environment is impossible due to space constraints.

A potential extension to this experiment is to allow the user to 'walk-in-place' – in this case, the user's movement would be directly correlated with his or her motion of the legs. We could even use an omni-directional treadmill to more closely simulate human locomotion. However, because the user is still staying in one place in space, this still does not give a direct equivalence between vestibular and visual stimuli, and therefore we believe this would again lead to motion sickness and a lower level of qualitative immersion.

4 Conclusion

The CAVE™ does not offer full proprioceptive immersion like a HMD does. When using a HMD, users, can walk freely, and their vestibular system is given stimuli that agree with the visual stimuli the user is experiencing. In the CAVE™, a user is limited to navigating while standing in place. There is a correlation between this and sickness as well as a lower qualitative sense of immersion. In order to have natural navigational responses, quantitative and qualitative measures of immersion should be high. Because of this, we feel that a CAVE™ is not the best tool for conducting navigational experiments, and that the ability to walk around and have full proprioceptive immersion is important in this area of experimentation.

References

[1] Brett R. Fajn, William H. Warren. Behavioral dynamics of steering, obstacle avoidance, and route selection. Journal of Experimental Psychology: Human Perception and Performance, 29(2): 343-362, April 2003

[2] Heinrich H. Bulthoff and Hendrik A. H. C. van Veen. Vision and Action in Virtual Environments: Modern Psycophysics in Spatial Cognition Research. In Vision and Attention, 233-252, 2001.

[3] A. Ekstrom. Navigation in Virtual Space: Psychological and Neural Aspects. Encyclopedia of Behavioral Neuroscience. 286-293, 2010.

[4] Clare Regan. An Investigation into Nausea and other Side-Effects of Head-Coupled Immersive Virtual Reality

Evaluating Attributes of Virtual Reality Displays for Interacting with Isosurface Renderings

Wesley R Miller*　　Johannes Novotny†

Brown University

ABSTRACT

We present a user study to evaluate different components of visual fidelity in a virtual reality environment and how they affect a user's ability to interact with scientific isosurface renderings of volume datasets. In the user study, we evaluate field of regard and stereoscopy by varying these components and measuring different aspects of users' completion of tasks requiring interaction with isosurface renderings of micro-CT scans of beetle tracheal systems. We carried out these tasks in the Brown University CAVE (CAVE Automatic Virtual Environment), and we varied the visual fidelity components between users. Using our results, domain scientists can choose which virtual reality environments and configurations are best suited for their research, and visualization scientists can be better informed of the importance of different features during the construction of new virtual reality displays.

Keywords: Virtual reality, isosurface rendering, CAVE.

1 INTRODUCTION

Considering the amount and complexity of the data the can be collected in many different scientific domains [7], immersive virtual reality for scientific visualization is becoming increasingly important. With the growing importance of virtual reality in scientific visualization, much work has been done to evaluate different components of virtual reality displays and their effectiveness for different visualizations [1, 2, 4, 5, 6]. One important kind of scientific visualization that can be done in virtual reality is biological isosurface data visualization. In our user study, we evaluated the effect of varying different components of fidelity in a virtual reality environment on the users' interactions with the isosurface data; specifically, we evaluated field of regard and stereoscopy with respect to the Brown University CAVE. We borrowed the experimental design from an earlier study completed at Virginia Tech to be able to compare results across the different displays used for each study [4].

Figure 1: Isosurface rendering of the micro-CT scan of a beetle of the Platynus genus used for the main tasks in the study.

*e-mail: wmiller@cs.brown.edu
†e-mail: novotny@cs.brown.edu

2 USER STUDY

2.1 Goals

The goals of this study were to replicate the previous study [4] in the Brown University CAVE. We hoped to either replicate the results, providing evidence for their generalization across virtual reality displays, or to significantly differ from them, showing that there may be unconsidered variables between the two displays that cause the differing results.

2.2 Datasets

The isosurfaces were obtained from John Socha of the Virginia Tech department of biomedical engineering and mechanics. Socha frequently uses isosurface visualizations in his own research, as do many others in his field. These specific isusurfaces, which are datasets he uses in his research, were generated from micro-CT scans of beetle tracheal systems. Two different scans were used, one from a beetle of the Pterostichus genus and one from a beetle of the Platynus genus.

2.3 Environment

The evaluations took place in the Brown University CAVE, a CAVE system with three rear-projected walls and a front-projected floor. It is capable of displaying stereo images as well as tracking a user's head and a wand input device. To test the different components of visual fidelity, the CAVE was switched between various configurations. Field of regard was varied between 90° and 270°, and stereo was turned on or off. Combining these variables created four distinct configurations, and each user was tested in exactly one configuration for the duration of the experiment. Unlike Laha [4], we did not test head-tracking and left that on for all user groups.

Users were given head tracking stereo glasses and a tracked wand with buttons. They were able to walk around the environment and, using the wand, grab the virtual space at any point and orientation and manipulate it to place it at any other point and orientation. They were also able to reset the model to its original location. The software used in testing made use of the VRUI toolkit [3]. This allowed for duplication of the interaction methods and visuals from the original experiment.

2.4 Tasks

There were twenty tasks overall, 5 training tasks using the Pterostichus dataset and 15 main tasks using the Platynus dataset. The tasks were organized into five separate categories: search, pattern recognition, spatial judgement, quantitative estimation, and shape description. Figure 3 shows which of the main tasks were in each category; it should be noted that counting tasks are a specific form of search task, so some tasks are marked as search and counting. A more detailed description of each individual task can be found in the previous study's paper [4].

2.5 Participants

We recruited graduate students at Brown University for this study. To conform with the practices in the previous study [4], our participants were unpaid volunteers. We recruited eight users in total, and each user was put into one of the four groups mentioned earlier, so there were two participants used to evaluate each fidelity configuration. Participants were asked about prior experience with isosurface visualization and biomechanics, and they were not permitted to be a part of the study if they had prior experience in these areas. Users would also be removed if they did not pass the spatial reasoning test given to them so that all users could be known to have the same baseline understanding of 3D space.

2.6 Procedure

Upon arrival, each participant was given and asked to read and sign our IRB approved informed consent form; the participant was then given a background questionnaire about demographic information to fill out. The spatial ability test was then administered, and the participant was guided to begin the training tasks, followed by the main tasks.

The user was given each task while facing away from the data, which was in a default position. The user was then permitted to turn around to begin the task as a timer was started, and the timer stopped when the user said, "done." The participant then gave the answer, which we recorded; each answer was scored on a continuous scale from 0 to 1 based on the same rubric used in the previous study, where higher grades corresponded to better answers. We then asked the user to give a rating of difficulty for the task on an integer scale from 1 to 7 and to give a rating of confidence in his/her answer on the same scale. These ratings and the time were recorded alongside the answer(s) given.

After the tasks were complete, the user was given a post-experiment questionnaire asking general questions about how he/she felt during the experiment and how difficult it was to perform the tasks.

3 Results

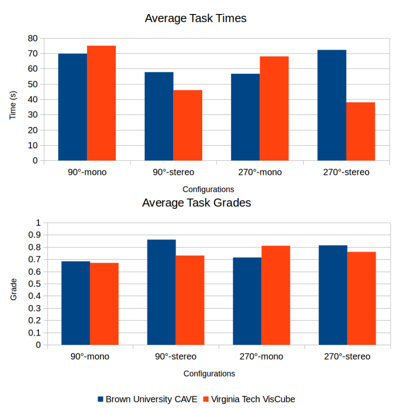

Figure 2: Charts showing the average task times and grades, respectively, for all members of each configuration at the Brown University CAVE for our study and at the Virginia Tech VisCube for Laha's study [4].

Figure 2 shows the average of the recorded task times and the average of the calculated task grades for each display configuration in our study compared to those values for Laha's study. We see a decreasing trend in average time replicated between 90° mono and 90° stereo, but we see contradicting trends between the two studies when looking at the average task times for the 270° field of regard cases. Looking at average task grades, there is an upward trend between grades and presence of stereoscopy replicated across the two 90° setups; however, there again appear to be conflicting trends between our data and that of Laha for the 270° configurations.

	Grade		Time		Difficulty		Confidence	
	FOR	ST	FOR	ST	FOR	ST	FOR	ST
T01 (s, c)								
T04 (s, c)		X						
T05 (s)			O	O				
T08 (s)	O	O		X				
T09 (s)		X		X	O	O		
T02 (sj)			X					
T06 (sj)		X						
T10 (sj)				X				
T11 (sj)	X				O	O		
T13 (sj)								
T14 (sj)		X				X		
T03 (qe)								
T15 (qe)								
T07 (sd)				X			O	O
T12 (pr)								

Figure 3: A table showing significant effects and interactions found by Laha and how our data compares. Each cross (X) represents a significant effect of the given component of fidelity on the given metric for the given task. Connected circles (O) represent a significant interaction between the components for the metric and task of concern. Green, yellow, and red indicate a strong agreement, a weak agreement, and no agreement with the finding from our data, respectively.

In Figure 3, we see a recreation of a figure in Laha's paper[4]. Fields for head-tracking have been removed, as is was a condition tested only in Laha's study, and color was added to indicate the level of agreement found in our data to the given finding in Laha's data. With our data, we confirmed 6 of Laha's findings. We had expected to confirm many more; however, there are some likely reasons as to why our results could have differed like this. Most notably, our differing results could come from our very low number of users; with so few users, each user has much more of an impact on the data, so the presence of any outliers can be very disruptive. Furthermore, there may be uncontrolled differences between the two displays that led to the results discrepancy; resolution is another component of visual fidelity, and it was not controlled across the two displays for these experiments.

4 Conclusion

In our study, we did not confirm a majority of the findings of the previous study, but we did confirm some. Because of the sensitivity to outliers with our current number of users, to get significant results confirming or contradicting those of Laha, we must include more participants in our study to increase the sample size and the significance of our results. Until then, we cannot make claims about generalizing trends across multiple displays or otherwise.

References

[1] K. Arthur. *Effects of field of view on performance with head-mounted displays*. PhD thesis, University of North Carolina at Chapel Hill, 2000.

[2] c. Demiralp, C. D. Jackson, D. B. Karelitz, S. Zhang, and D. H. Laidlaw. CAVE and fishtank virtual-reality displays: A qualitative and quantitative comparison. *IEEE Transactions on Visualization and Computer Graphics*, 12(3):323–330, 2006.

[3] O. Kreylos. Environment-independent vr development. In *Advances in Visual Computing*, pages 901–912. Springer, 2008.

[4] B. Laha, D. A. Bowman, and J. J. Socha. Effects of VR system fidelity on analyzing isosurface visualization of volume datasets. *IEEE Transactions on Visualization and Computer Graphics*, 20(4):513–522, 2014.

[5] B. Laha, K. Sensharma, J. D. Schiffbauer, and D. A. Bowman. Effects of immersion on visual analysis of volume data. *IEEE Transactions on Visualization and Computer Graphics*, 18(4):597–606, 2012.

[6] E. D. Ragan, R. Kopper, P. Schuchardt, and D. A. Bowman. Studying the effects of stereo, head tracking, and field of regard on a small-scale spatial judgment task. *IEEE Transactions on Visualization and Computer Graphics*, 19(5):886–896, 2013.

[7] A. van Dam, D. H. Laidlaw, and R. M. Simpson. Experiments in Immersive Virtual Reality for Scientific Visualization, 2002.