Evaluating Depth Perception of Volumetric Data in Semi-Immersive VR

Isaac Cho* Wenwen Dou† Zachary Wartell‡ William Ribarsky§ Xiaoyu Wang¶

University of North Carolina at Charlotte

ABSTRACT

Displays supporting stereopsis and head location based motion parallax can enhance human perception of complex three dimensional datasets. This has been demonstrated for datasets containing 3D surfaces and 3D networks. Yet many domains, such as medical imaging, weather and environment simulations and fluid flow, generate complex volumetric data. This poster present results of an initial formal experiment that examines the effectiveness of various display conditions on depth perception of volumetric data. There is an overall benefit for stereoscopy with head-tracking in enhancing depth perception. Further, familiarity with 3D games and VR-like hardware improves the users'ability to perceive such data.

Index Terms: H.5.1 [Multimedia Information Systems]: Artificial, augmented, and virtual realities—;

1 INTRODUCTION

Previous research has demonstrated the utility of computer displays that provide stereopsis and structure-from-motion for enhancing human perception of complex three-dimensional datasets. This includes fully immersive displays such as CAVE's and HMD's and semi-immersive displays such as desktop VR and the virtual workbench. For example, studies by Ware et al. examine the effect of the stereoscopic and kinetic depth for understanding 3D networks which are represented by tubes or lines [4]. Their results demonstrate improved user performance at finding paths in a complex 3D networks when using stereopsis and structure-from-motion.

A significant number of visual analytic domains, however, also heavily use 3D volumetric data. Volumetric data is characterized by large amounts of transparency, occlusion and ambiguous spatial structure. There has been a fair amount of evaluation of perception of volumetric data under different rendering conditions and parameterizations such as different transfer functions (for example see [3]) but somewhat less on evaluation of perception of volume under sterescopic display (for example see [2]). Prior study of surface and 3D networks shows that as 3D geometry grows more complex, VR display capabilities can further improve shape and depth perception. One would expect similar results for volumetric data. Further, the addition of VR display technology could be especially important with time-varying volumetric datasets that are viewed in real-time where extensive preprocessing for optimizing transfer functions and volume rendering parameters is not possible. An example would be real-time, streaming doppler weather radar data. With the increasing affordablility of semi-immersive VR displays and GPUs capable of advanced volume rendering, there is a pressing need to quantify the effectiveness of stereoscopy and structure-from-motion on volumetric data and also to quantify how these display parameters interact with other volumetric rendering conditions. The large number of potential display hardware and rendering variables make such evaluations particularly challenging. In this paper, we take a first step by using a fixed set of volumetric rendering parameters–chosen through pilot studies–and then varying the VR display hardware employed.

*e-mail: icho1@uncc.edu
†e-mail:wdou1@uncc.edu
‡e-mail:zwartell@uncc.edu
§e-mail:ribarsky@uncc.edu
¶e-mail:xwang25@uncc.edu

Figure 1: Similarity comparison between our artificial dataset and actual MRI blood vessel scan. Top : Maximum Intensity Projection rendering of blood vessels [1]. Bottom : Our artificial dataset

This poster presents Part I of a two part experiment on the benefits of stereoscopy and head-tracking for a person's correct perception of depth ordering of volumetric objects. The experimental design is motivated by datasets such as the MRI scan of blood vessels shown in Fig 1a reproduced from [1]. As is typical of volumetric data, this dataset is characterized by a heavy presence of transparency, occlusion and highly ambiguous spatial structure. In Fig 1a, it is particularly challenging to determine the depth order of the blood vessels designated by the red square. As discussed in [1], the volume rendering technique used here makes it appear that the square-shaped vessel is in front of the diagonal one. However, in fact the diagonal one is in front of the square-shaped one. We mimic this type of ambiguity by generating controlled experimental datasets such as Fig 1b where the user's task is to determine the depth ordering of various occluding, transparent cylinders. The subjects view the datasets under a variety of display conditions including combinations of stereoscopic display, head-tracking, and small object rotations.

2 EXPERIMENTAL DESIGN

Our study tests the effectiveness of a semi-immersive VR display for facilitating depth perception of volumetric data. We examine the effects of display environments on two tasks: a depth ordering task, in which participants describe the general depth ordering between 6 volumetric cylinders with no time limit; and a depth discrimination task, in which participants must distinguish the depths of a pair of cylinders with a short exposure time (2 sec). Both tasks use a desktop VR system which consists of a stereo display and a tracked pair of stereo glasses. The synthetic volumetric dataset contains six overlapping cylinders of varying diameters and transparency (Figure 1 bottom). We recruited 16 participants total.

3 EXPERIMENT 1: DEPTH DISCRIMINATION

Experiment 1 evaluates how stereoscopy and structure-from-motion affect users when performing a depth discrimination task. The participant must determine which of two cylinders, one horizontal and the other vertical, is in front of the other. The volumetric dataset actually contains 6 cylinders, but in each trial a pair of cylinders is designated as the target pair for the depth discrimination task. Participants are exposed to the volumetric dataset for a short amount of time (2s) so that they do not have time to cognitively reason about the depth order based on factors such as transparency, window size, etc. On each trial, the first screen displays a 2D picture which designates which of the 9 intersections of the 6 cylinders is the target pair. Next, the screen displays the volumetric dataset for 2 seconds. Finally the screen displays a menu with three choices: "the horizontal cylinder is in front", "the vertical cylinder is in front", or "I don't know". Figure 2 shows the 3 screens display during the task. The 2D picture screen is shown on the left, the volumetric data screen in the middle and the question screen on the right. Note, we deliberately do not used a forced-choice protocol in this experiment because we want to gather data on how often a user reports they can not determine the depth ordering. A force-choice protocol would have conflated results for trials where participants were guessing at the depth order with those trials in which they felt they could determine a specific ordering.

Figure 2: Illustration of the depth discrimination task procedure.

The experiment has six conditions: non-stereo without head-tracking, stereo with no head-tracking, no stereo with head-tracking, stereo with head-tracking, no stereo with head-tracking simulation and stereo with head-tracking simulation. The last two conditions were added because in pilot tests not all users utilized the head-tracking when limited to the 2 second exposure time. The head-tracking simulation condition automatically rotates the cylinders left and right by 10 degrees. Technically this is a kinetic depth manoeuvre, but for a small range of motion the visual effect is similar to the participant quickly moving her head side to side. When not using head-tracking, a participant uses a chin rest. In this condition, the view frustums are calibrated for this fixed head position.

3.1 Experiment 1 Results

The display condition is the independent variable and the dependent variable is accuracy, measured by the percentage of correct depth judgements. A one-way ANOVA does not show significant effect of the display condition on accuracy. However, post-hoc LSD test comparisons show that the mean accuracy for condition SH (stereoscopy with head tracking) (M = 0.75, SD = 0.16) is significantly better than the condition of no stereoscopy no head- tracking (M = 0.67, SD = 0.2) with $p = 0.01$. However, the magnitude of the effect is fairly small.

3.1.1 Experienced vs. less experienced observers

During the experiment, we noticed that computer science major participants reported more confidence with their performance, while participants of other majors reported less confidence and seemed less comfortable using the semi-immersive VR environment. Further, a one-way ANOVA showed that CS majors (M = 5.6, SD = 1.5) scored significantly higher than other majors (including history, nursing, psychology) (M = 2.9, SD = 1.9) on questions indicating experience with computer games $p = 0.006$. This lead us to categorize the participants in two groups–those with more gaming experience and those with less–with 8 participants in each group. We then evaluated if depth discrimination performance differed significantly between these groups.

First, we compare the overall performance of the more and less experienced group over all display conditions. We use a 6X2 (condition X experience) factorial analysis of variance to evaluate the effects of condition and experience on accuracy for the depth-discrimination task. Results indicate a significant main effect for the experience factor, F(1, 132) = 10.57, p = 0.001. As hypothesized, the accuracy of more experienced users is higher. The interaction, experience X condition, is not significant.

One-way ANOVA was used to test for accuracy of both groups among the six display conditions. The effect of condition on accuracy for the less experienced user group yielded no significant results. The effect of condition on accuracy for the more experienced group also had no significant main effect (F(5, 66) = 2.047, $p = 0.08$). However, post hoc LSD test comparisons indicate that the mean accuracy for condition stereoscopy with head-tracking) (M = 0.83, SD = 0.12) is significantly better than the condition of non-stereoscopy with head-tracking simulation (M = 0.67, SD = 0.2) $p = 0.025$, and non-stereoscopy, non-head-tracking (NSNH) (M = 0.64, SD = 0.2) $p = 0.009$. Different from the findings from overall population, the mean accuracy for the condition, non-stereoscopy with head tracking, (M = 0.79, SD = 0.1) $p = 0.042$ is significantly better than the NSNH condition.

In summary, for all participants stereoscopy with head-tracking led to higher accuracy than non-stereo, non-head-tracking in the depth-discrimination task on volumetric datasets, but showed no advantage over other conditions. For experienced participants it is also the case that non-stereoscopy with head-tracking led to better accuracy than non-stereo, non-head-tracking.

4 Future Work

Part II of our study is an depth ordering task where the participants must state which of six cylinders is the farther one, the middle one or the closest one. Preliminary results indicate a more statistically significant effect of display condition on accuracy in this finer grained task.

References

[1] R. Maciejewski, S. Choi, D. S. Ebert, and H. Z. Tan. Multi-modal perceptualization of volumetric data and its application to molecular docking. In *WHC '05: Proceedings of the First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 511–514, Washington, DC, USA, 2005. IEEE Computer Society.

[2] B. Mora and D. S. Ebert. Instant volumetric understanding with order-independent volume rendering. *Computer Graphics Forum*, 23(3):489–497, September 2004.

[3] A. J. Stewart. Vicinity shading for enhanced perception of volumetric data. In *Visualization, 2003. VIS 2003. IEEE*, pages 355–362, 2003.

[4] C. Ware and P. Mitchell. Visualizing graphs in three dimensions. *ACM Trans. Appl. Percept.*, 5(1):1–15, 2008.