Assessing Multiple Participant View Positioning in Virtual Reality-Based Training

Jonathan W. Kelly, Eliot Winer, Stephen Gilbert, Michael Curtis, Eduardo Rubio, Kenneth Kopecky, Joseph Holub

Iowa State University

Ames, IA

jonkelly; ewiner; gilbert; mkc; erubio10; kennyk; jholub@iastate.edu

Julio de la Cruz

ARL HRED STTC

Orlando, FL

julio.delacruz@us.army.mil

ABSTRACT

As cost, time, and other challenging resource requirements are placed on U.S. Joint forces training, the role of simulations will play an even greater role than it does today. To effectively aid a Warfighter in gaining critical skills and to assess the proficiency of those skills, computer-based training must advance beyond traditional desktop simulations and monoscopic projection technology. Virtual Reality (VR) based training has been proven in fields such as medical and engineering to increase a trainee's level of immersion, and increase training performance in several metrics including accuracy and efficiency, while simultaneously decreasing cost.

Warfighter training offers a unique set of challenges that demand additional studies before they can be correctly addressed in a VR environment. Primary among them is the ability to have multiple Warfighters train together. While VR systems typically include monocular and binocular depth cues, the actual imagery is only drawn correctly for a single viewer. Imprecision in Warfighter training can result in incorrect acquisition of an enemy avatar's position and/or target location. These errors can carry over into future training as well as actual missions.

In this paper, a formal method to produce a combined viewpoint, suitable for multiple participants, in a VR simulation-based training is presented. The concepts of monocular and stereoscopic depth cues and their effect on Warfighter training will be discussed. A comprehensive review of current research into simulation-based training environments will also be presented. Lastly, we will present new results from a formal user study comparing the proposed combined viewpoint with that of a typical VR system in a Warfighter training task involving shooting of virtual targets. Initial results of this study show significant advantages to using the combined viewpoint. Analyses of the results show the maximum shooting error committed by an individual participant was reduced by up to 47%.

ABOUT THE AUTHORS

Jonathan Kelly, Ph.D., is an assistant professor in the department of Psychology and a faculty affiliate of the Virtual Reality Applications Center (VRAC) at Iowa State University. He is an expert in spatial tasks in virtual environments, working on understanding how visual cues such as the presence of landmarks, object shapes, and view distortions affect Warfighters in simulation-based training.

Eliot Winer, Ph.D., is the associate director of the Virtual Reality Applications Center and associate professor of mechanical engineering at Iowa State University. He is currently co-leading an effort to develop a next-generation mixed-reality virtual and constructive training environment for the U.S. Army. Dr. Winer has over 15 years of experience working in virtual reality and 3D computer graphics technologies.

Stephen Gilbert, Ph.D., is the associate director of the Virtual Reality Applications Center at Iowa State University and assistant professor of industrial and manufacturing systems engineering in the human factors division. His research focuses on intelligent tutoring systems and emerging technologies for training. He is currently co-leading an effort to develop a next-generation mixed-reality virtual and constructive training environment for the U.S. Army.

Michael Curtis, Eduardo Rubio, Ken Kopecky, and Joseph Holub are pursuing graduate degrees in human computer interaction and mechanical engineering at Iowa State University.

Julio de la Cruz, M.S. is the Chief Engineer for Synthetic Environments Research at the U.S. Army Research Laboratory, Human Research and Engineering Directorate, Simulation and Training Technology Center in Orlando, Florida and is responsible for the research and development of applied simulation technologies for learning, training, testing and mission rehearsal. Mr. de la Cruz has a Bachelor of Science degree in Electrical Engineering from the City College of New York and a Master of Science degree in Industrial Engineering from Texas A&M University. He is a graduate of the U.S. Army School of Engineering & Logistics and alumnus of the Advanced Program Managers Course at the Defense Systems Management College.

Assessing Multiple Participant View Positioning in Virtual Reality-Based Training

Jonathan W. Kelly, Eliot Winer, Stephen Gilbert, Michael Curtis, Eduardo Rubio, Kenneth Kopecky, Joseph Holub

Iowa State University

Ames, IA

jonkelly; ewiner; gilbert; mkc; erubio10; kennyk; jholub@iastate.edu

Julio de la Cruz

ARL HRED STTC

Orlando, FL

julio.delacruz@us.army.mil

INTRODUCTION

Virtual Reality (VR) has been proven useful for applications such as medical training simulations (Grantcharov et al., 2004), physical rehabilitation (Jack et al., 2001), and psychological rehabilitation (Glantz, Rizzo & Graap, 2003). VR-based training has been proven to increase training performance in several metrics including accuracy and efficiency, while simultaneously decreasing cost.

With the increasing role of simulations in US joint forces training comes an increasing need for research on the effectiveness of those simulations for conveying the intended information. Displaying a virtual environment on a projection screen visible to multiple individuals is often a training requirement, as it allows group training of Warfighter skills and evaluation of training effectiveness by observers. However, displaying the virtual environment to multiple individuals on the same screen demands an understanding of the effects of the rendering technique on the perceptual experiences of the viewers. Past work indicates that the traditional rendering technique can result in distortion of perceived spatial properties of the virtual environment for most viewers. Imprecision in Warfighter training caused by misperception of the environment can result in incorrect acquisition of an enemy avatar's position and/or target location. These errors can carry over into future training as well as actual missions. The primary goals of the current project were to 1) quantify the perceptual experiences of viewers using the traditional rendering technique, and 2) test the feasibility of a novel rendering technique intended to reduce distortions of perceived spatial properties of the virtual environment. The novel technique involves rendering the virtual environment from a combined viewpoint, suitable for multiple participants. This technique is implemented and evaluated in a simulation-based training scenario.

Rendering Virtual Environments and the Center of Projection (CoP)

When presenting the same scene on a projection screen to multiple viewers, the scene is typically rendered from one location, referred to as the CoP. In the traditional leader-follower approach to scene rendering, the CoP is linked to the position of one person, referred to as the leader. The leader's head position and orientation is continuously tracked, allowing for dynamic updating of the scene from his or her perspective. All other viewers are referred to as followers, and they view the scene from a position displaced from the CoP. This arrangement can lead to perceptual problems for the followers. For example, when followers change their position or orientation, the scene does not change to reflect their movements. Conversely, when the leader changes position the scene can change dramatically, even if the followers remain completely still. Moreover, even when the leader and followers are static, the followers' perceptual experiences of the virtual environment are spatially distorted relative to the intended scene. This paper focuses on a new technique for mitigating the perceptual distortion experienced by followers in the traditional leader-follower approach.

Existing alternatives to the leader-follower approach include frame-rate splitting, image blending and view clustering. Frame-rate splitting involves presenting each user with images rendered from his or her actual location by dividing the frame rate among multiple users. However, this approach is only feasible with a high frame rate and few users; lower frame rates or larger groups of users result in unacceptably low frame rates and dim images. Image blending (Marbach, 2009) renders the environment from each viewer's position and orientation and creates a composite image with a blend zone where the images overlap on the screen. Image blending is primarily useful when multiple viewers are looking at different portions of the screen. View clustering (Marbach, 2009), herein referred to as view averaging, is based on calculation of gaze-intersection points when multiple viewers look at the same portion of the screen. The gaze-interaction points are grouped together and the scene is drawn using the average position. The result of view averaging is that the view is rendered from a location at which nobody is standing, meaning that all viewers are displaced from the CoP to some extent. View averaging has been proposed and implemented elsewhere, but its effects on users' perceptual experiences have never been formally tested, and the current project was designed to fill this research gap.

Perceptual Distortion after Displacement from the CoP

Viewing the virtual environment from a position displaced from the CoP causes measurable distortions of perceived space. Forward and backward displacement causes perceived depth extents to contract and expand, respectively. Leftward and rightward displacement causes the perceived vanishing point at the horizon to shift in the direction opposite the displacement, which causes directions and angles in the environment to appear distorted. These effects of displacement are large, and are grounded in the viewing geometry: displacement from the CoP changes the monocular and binocular characteristics of the images that are incident on the viewer's eyes. Based on transformations of the binocular and monocular cues in the perspective images, models of the virtual environment can be used to generate predictions about the perceived distortions of depth and direction (Banks, Held & Girshick, 2009; Held and Banks 2008; Sedgwick 1991; Woods, Docherty & Koch, 1993).

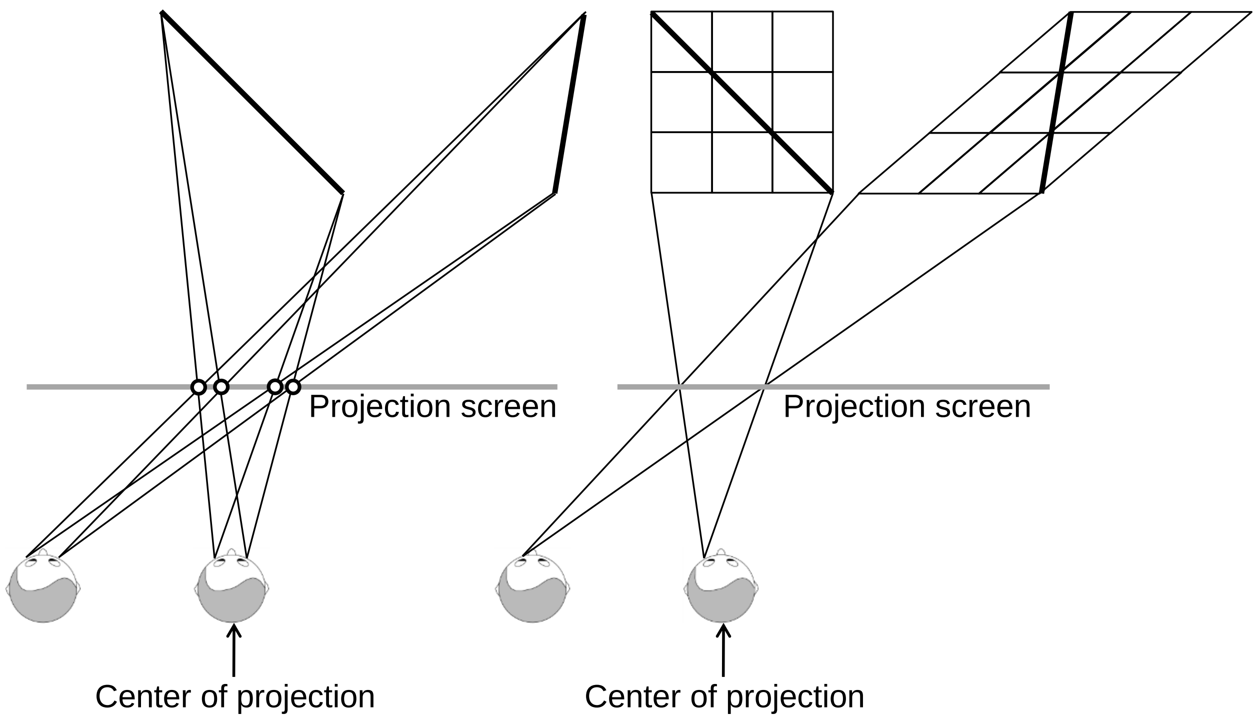

Figure 1. Predictions of perceived angle after leftward displacement from the CoP. Model predictions based on binocular (left) or monocular (right) cues are identical.

Figures 1 and 2 show the predicted perceptual distortion, based on models of the viewing geometry, when viewing an oblique line in the virtual environment after leftward displacement (Figure 1) and forward displacement (Figure 2) from the CoP. Whereas the viewer at the CoP should perceive the line to be directly in front of him/her and oriented at a 45° angle, the leftward displaced viewer should perceive the line to be shifted rightward and rotated in depth (Figure 1), and the forward displaced viewer should perceive the line to be shifted closer and rotated in depth (Figure 2). The left side of each figure shows the predicted distortion based on binocular cues, whereby the perceived locations of points on the virtual object are calculated by projecting rays from the two eyes through the image points on the screen and out into space. The right side of each figure shows the predicted distortion based on monocular cues, whereby displacement from the CoP causes changes to the projective geometry. The two methods (using binocular or monocular information) of predicting perceived direction and angle produce exactly the same predictions. Therefore, these predictions are herein referred to as cue-based model predictions.

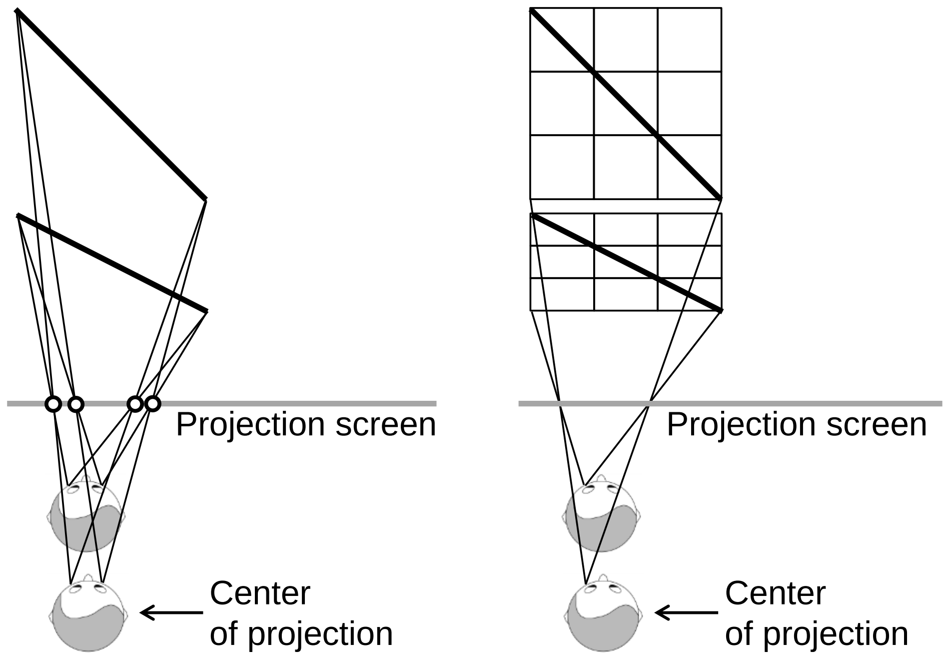

Figure 2. Predictions of perceived angle after forward displacement from the CoP. Model predictions based on binocular (left) or monocular (right) cues are identical.

Perceptual Compensation

The human perceptual system can compensate for some of the distortion caused by displacement from the CoP. For example, viewing a photograph from an angle doesn't cause the depicted scene to look completely distorted, despite the displacement of the viewing position from the CoP. Kelly et al. (in press; see also Pollock, Burton, Kelly, Gilbert & Winer, 2012 and Burton et al., 2012) conducted a study to quantify the perceptual distortion experienced by viewers of a virtual environment after displacement from the CoP. Participants made judgments of object angle and object depth, viewed in stereo on a six-walled virtual reality (VR) system. For angle judgments, participants viewed a line on the ground plane and judged its angle (similar to the stimulus shown in Figure 1) while standing at the CoP or after leftward or rightward displacement from the CoP. For depth judgments, participants viewed a rectangle on the ground plane and judged its depth (i.e., the distance from the front to back edge of the object) while standing at the CoP or after forward or backward displacement from the CoP. Both types of judgments showed perceptual distortions in the same direction as the model predictions but of a significantly smaller magnitude. In some cases, the perceptual distortion reported by viewers was only 50% of the predicted distortion based on model predictions like those shown in Figures 1 and 2. The authors concluded that viewers can compensate, to some extent, for the perceptual distortions induced by displacement from the CoP. This perceptual compensation is similar to the common experience that photographs look spatially undistorted when viewed from an angle, which is essentially a displacement from the CoP.

Effects of Perceptual Distortion on Training in Virtual Environments

Perceptual distortion of any type poses potential problems for VR as a training tool. For example, if the virtual world is perceived to be only half of its intended size, then training of spatial skills in VR will produce large errors when attempted on the battlefield. Furthermore, different viewers typically have differently distorted perceptual experiences, leading to difficulties in communication about the virtual environment. For example, a virtual object that looks like a Humvee to one person will appear to be the size of a Mini Cooper to someone standing much closer to the screen (see Figure 2). Such discrepancies can lead to breakdowns in communication between multiple viewers of the same environment (Pollock et al., 2012). Therefore, methods for mitigating perceptual distortions are critical. But in addition to mitigating perceptual distortions, it is also important that multiple viewers' experiences of the virtual environment be as similar to one another as possible.

Mitigation of Perceptual Distortion through View Averaging

Perceptual errors in judging spatial properties of the virtual environment will impact transfer of training from virtual to real environments, and will interfere with communication about the virtual environment. One approach to mitigating the perceptual distortions experienced by viewers displaced from the CoP is to render the virtual environment from a location that is the average of all viewers' head positions. For example, if two viewers stand at positions with (X,Y) coordinates of (-5,0) and (+5,0), then the view averaging technique would be to render the environment from (0,0), which is the average of the two viewers' positions. The predicted effect of view averaging is that perceptual distortions caused by displacement from the CoP will be distributed across all viewers. In contrast, the leader-follower approach places all of the perceptual distortion on the followers and none on the leaders. When using view averaging, reducing the maximum perceptual distortion experienced by any individual reduces the amount of perceptual compensation that must occur in order to perceive the scene accurately.

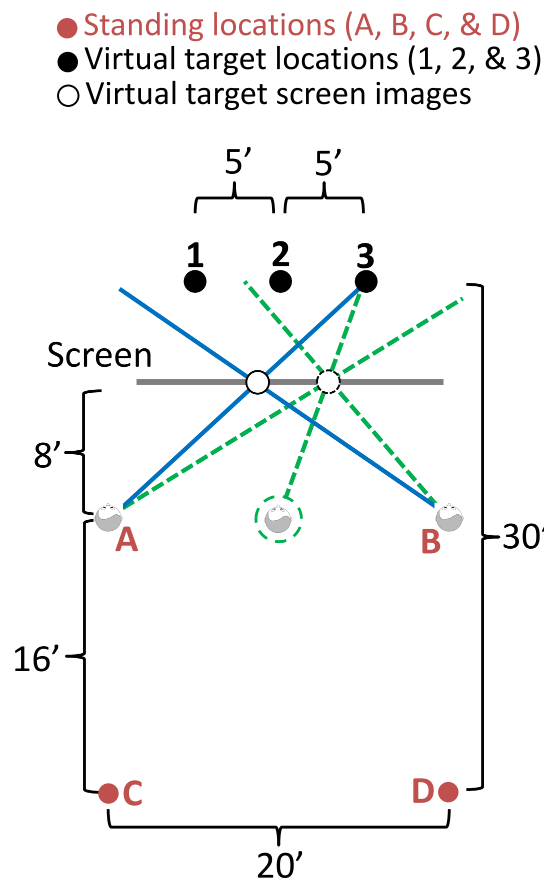

Figure 3. Predictions of perceived target direction in an aiming task for two viewers standing at positions A and B. Solid blue lines indicate predicted aiming when the scene is rendered from position A. Dashed green lines indicate predicted aiming when the scene is rendered from the average of position A and B, indicated by the dashed green circle.

Figure 3 shows an example of the predicted effects of view averaging on the aiming task used in the current study. In the example, two viewers stand at positions A and B attempt to shoot at virtual target 3. When the view is rendered from position A (i.e., A is the leader and B is the follower), the viewer standing at position A aims accurately (solid blue line extending from position A, intersecting target 3). However, the viewer at position B aims at the image of target 3 on the screen, which is shown as a circle in Figure 3. This results in an error of 36° to the left of the actual target direction (solid blue line extending from position B). Dashed green lines in Figure 3 illustrate the predictions when the view is rendered from the average of positions A and B (indicated by the head surrounded by the dashed green circle in Figure 3). The image of target 3 is drawn on the screen from the average head position. The viewer at position A aims at the image of target 3 on the screen, shown as a dashed circle in Figure 3, producing an error of 10° to the right of the target. The viewer at position B aims at the same image of target 3 on the screen, producing an error of 23° to the left of the target. Although view averaging should cause aiming errors for all viewers, the maximum error experienced by any one individual is reduced.

QUANTITATIVE EVALUATION OF VIEW AVERAGING IN A GROUP AIMING TASK

The goal of the current research study was to evaluate the view averaging technique and to measure its effect on an aiming task. In order for view averaging to be a viable technique for scene rendering, it must produce measurable benefits over the traditional technique of rendering the scene from the location of one viewer. Pairs of participants aimed and shot Airsoft rifles at targets in a virtual environment displayed on a large projection screen. The shooting task was performed when the scene was rendered from one participant's position (leader-follower) and when the scene was rendered from the average of the two participants' head positions (view averaging).

Participants

Twenty students affiliated with Iowa State University participated in exchange for monetary compensation. Four Soldiers from the Iowa National Guard volunteered to participate.

Stimuli and Procedure

The experiment took place in the MIRAGE (Mixed Reality Adaptive Generalizable Environment) facility at Iowa State University. The MIRAGE supports integration of physical and virtual objects through high precision motion tracking within a 40' × 40' area. Within this space, there is a 33' × 11' fixed screen display with six active-stereo WUXGA projectors. Tracking is accomplished using a MotionAnalysis optical tracking system which provides sub-millimeter accuracy.

The virtual environment consisted of a ground plane covered with a grass texture along with a bull's eye mounted on a vertical post (see Figure 4). The bull's eye appeared at one of 18 target locations, which varied in their X, Y, and Z positions, and participants' task was to shoot the center of the bull's eye. Shooting was always done from a standing position. Participants stood in each of six paired positions (each position is marked in Figure 3): A-B, B-A, C-D, D-C, B-D, and D-B. For each position pair, participants shot at all 18 target locations, presented in a random sequence. The order of position pairs was counterbalanced to prevent order effects. The environment was rendered from one participant's location (leader-follower condition) or from the average of both participants' positions (averaging condition). Rendering condition was blocked, such that participants completed all position pairs under one rendering condition before beginning the next, and order was counterbalanced.

Figure 4. Perspective view of the virtual environment.

After providing informed consent, both participants were outfitted with a tracked helmet and Airsoft rifle. Prior to beginning the experiment, each participant individually stood in the center of the room and was given an unlimited amount of practice shots at one target. During practice, the scene was rendered from the participant's position and visual feedback (showing the location that was shot on a given practice trial) was provided to help the participant learn the dynamics of the rifle. Feedback was restricted to practice, and no feedback was provided once practice concluded. The rifle did not fire actual projectiles, but shot location was determined by the tracked location and orientation of the rifle.

After practice, participants were led to the first position pair. In the leader-follower rendering condition, one participant was randomly assigned to be the leader, but this distinction was not relevant in the view averaging condition. For each rendering condition, participants stood at each position pair and shot at all targets from that position, resulting in 216 shots fired by each participant. No feedback about accuracy was provided during testing. Since the research objective was to understand how rendering technique influences accuracy, participants were encouraged to take their time and to be as accurate as possible. The screen went blank when participants were walking between the different positions, so as to prevent them from learning about the rendering technique being tested on a given trial. The experiment took approximately 30 minutes to complete.

Results

The primary dependent measure was the aiming error in the yaw dimension. Errors committed by civilians and Soldiers were highly similar across almost all conditions, although statistical comparisons using the Soldiers' data were not conducted due to the small sample size. Errors for position pairs in which the participants were separated in the left-right dimension (i.e., pairs A-B, B-A, C-D, and D-C) are shown in Figure 5. In the leader-follower rendering condition (left half of Figure 5), the leader was generally quite accurate. This is not surprising, because

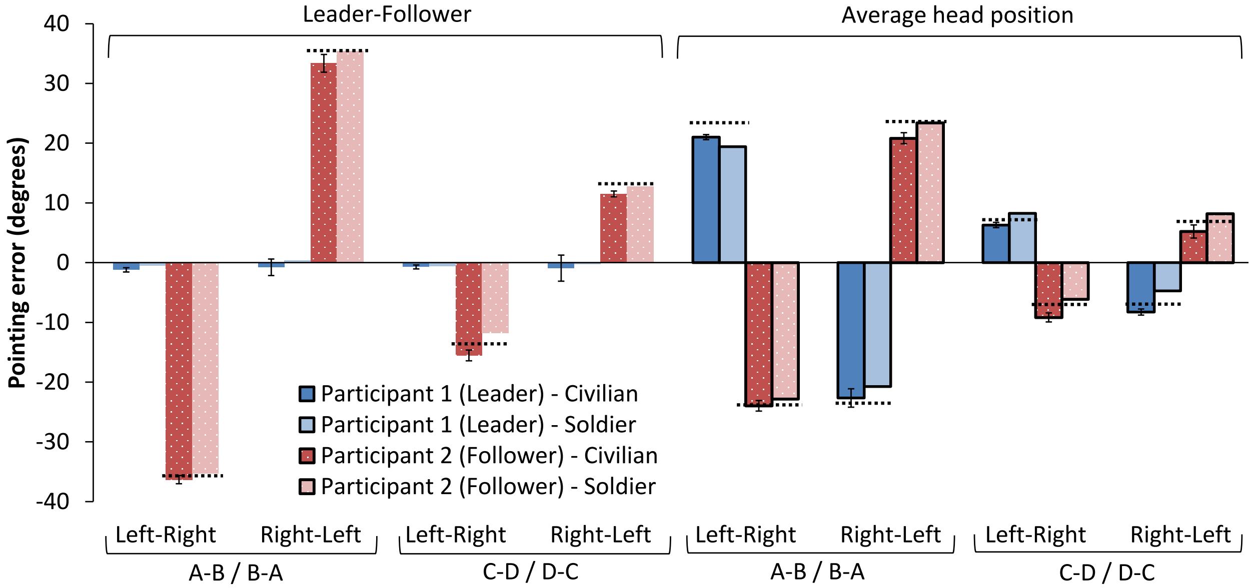

Figure 5. Pointing errors (i.e., aiming errors) when participant pairs were displaced left-right. Results from the leader-follower rendering condition are shown on the left, and results from the view averaging condition are shown on the right. Dashed lines indicate predicted errors based on the viewing geometry.

the view was rendered from the leader's perspective. However, the follower produced large errors, and these errors were consistent with predictions based on the viewing geometry (dashed lines in the data figures show model predictions). Actual errors differed from predicted errors in some conditions, but these differences were inconsistent and could be explained by a general tendency toward negative pointing errors (which could be due to minor errors in tracking calibration).

In the averaging condition, in which the environment was rendered from the average of the participants' head positions, both participants made relatively large errors and these errors were distributed relatively evenly across the two participants. Furthermore, errors did not differ reliably from predictions based on the viewing geometry (shown as dashed lines in the data figures).

Errors for position pairs in which the participants were separated in the front-back dimension (i.e., pairs B-D and D-B) are shown in Figure 6. In the leader-follower rendering condition (left half of Figure 6), the leader was accurate but the follower produced large errors, consistent with predictions based on the viewing geometry (see horizontal dashed lines).

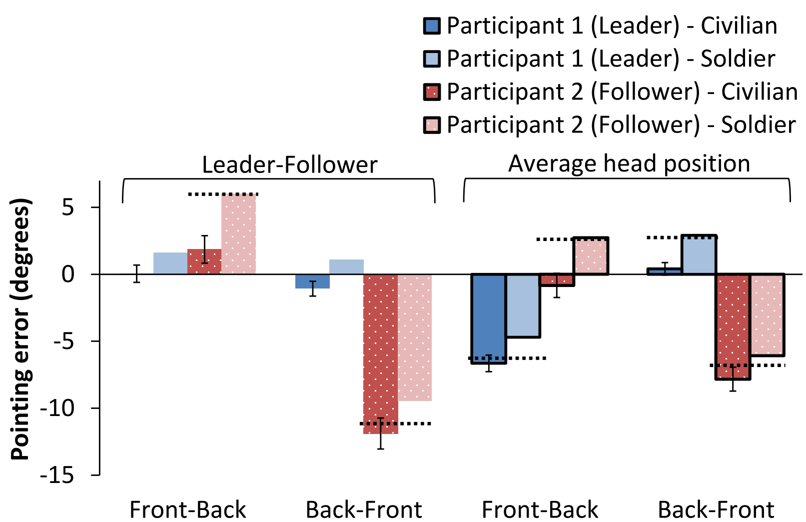

Figure 6. Pointing errors (i.e., aiming errors) when participant pairs were displaced front-back. Results from the leader-follower rendering condition are shown on the left, and results from the view averaging condition are shown on the right. Dashed lines indicate predicted errors based on the viewing geometry.

Actual errors differed from predicted errors in some conditions, but these differences were inconsistent and could be explained by a general tendency toward negative pointing errors (possibly due to minor calibration errors).

In the averaging condition, (right half of Figure 6), the participant who stood in front made larger errors than the participant in back, which was predicted by the viewing geometry.

Discussion

The current study was designed to evaluate a new method of rendering the virtual environment from the average position based on multiple viewers' head positions. This view averaging technique was compared to the traditional leader-follower technique, in which the scene is rendered from the location of one viewer (the leader).

When shooting at virtual targets, the leader-follower technique produced accurate performance for leaders but large errors for followers. For example, in the A-B and B-A position pairs in which the follower was laterally displaced from the leader, followers committed an average error of 36.3°. In contrast, the view averaging technique resulted in errors for both participants, but the errors were of smaller magnitude than in the leader-follower conditions. For example, in the A-B and B-A position pairs, both participants committed an average error of 22.1°. This represents a large increase in error relative to the leader's performance in leader-follower, but a large decrease in error relative to the follower's performance.

Errors committed by civilians and Soldiers were equally affected by manipulation of viewing position and rendering technique. This indicates that all groups of individuals are similarly influenced by the perceptual errors associated with displacement from the CoP. However, it is possible that performance could be influenced by familiarity with the standing positions tested. To that end, participants in the current study were always separated laterally (e.g., A-B in Figure 3) or front-back (e.g., B-D) but never in diagonal formations (e.g., B-C) that are commonly used in the militaries. The current study was designed to manipulate just one dimension (X or Y) at a time, but using the more familiar diagonal configurations in future studies could produce different results.

The current project did not measure or control for the viewer's height (or any other physical attributes), which could be another source of rendering errors in both the leader-follower and view averaging techniques. For example, rendering the scene from the perspective of a very tall viewer will make the environment appear distorted when seen by a shorter viewer. Based on preliminary analyses, rendering the environment from the wrong height will lead to errors in pitch when shooting at targets.

CONCLUSIONS

The traditional leader-follower technique to rendering the virtual environment led to very large aiming errors by participants in a virtual target-shooting task. Depending on the positions of the leader and follower, follower errors were as large as 35°, which seems unacceptably high for the purpose of training a skill like aiming. The view averaging technique, whereby the environment was rendered from the average position of the two viewers, effectively distributed the follower's errors across both individuals. In this way, the maximum aiming error by either individual was around 25°. However, the resulting aiming errors were still quite large, and errors equally affected both viewers (as opposed to the leader-follower technique, in which the leader was typically very accurate). Therefore, the current study indicates the feasibility of the view averaging technique, and identifies some of the trade-offs associated with these two rendering techniques.

One clear implication of these results is that positioning viewers farther from the screen reduces perceptual distortion of the virtual environment. Viewers in the A-B / B-A configurations (which were 8' from the screen) produced shooting errors that were three times larger than those produced by viewers in the C-D / D-C configurations (which were 24' from the screen). Placing viewers farther from the screen is a relatively simple but effective way to reduce errors by taking advantage of the projective viewing geometry. Keeping viewers close together will also reduce perceptual distortion.

Past research has shown that displacement from the CoP causes distortion of properties of perceived shape, such as object depth or the angle formed by two intersecting lines. However, those studies also indicate that viewers compensate, to some degree, for the effects of displacement. In other words, perceptual distortion of perceived shape is not as large as predicted by geometric models. The results of the aiming task, however, show no evidence for such correction. It is possible that correction for the effect of displacement only occurs for shape properties but not location properties, similar to other research showing dissociation between perceived shape and location

(Loomis, Philbeck, & Zahorik, 2002). However, a more controlled study which tests both perceived shape and location after displacement from the CoP is needed in order to verify this hypothesis.

ACKNOWLEDGEMENTS

This work was supported by a grant from the United States Army RDECOM – STTC, and is based in part upon work supported by the National Science Foundation under CNS-1156841. Special thanks to Timothy Honeywell, Andrew Mendez, and Amanda Skinner for help with pilot research leading to the current project.

REFERENCES

Banks, M., Held, R., & Girshick, A. (2009). Perception of 3-D layout in stereo displays. *Information Display, 25*(1), 12-16.

Burton, M.B., Pollock, B.B., Kelly, J.W., Gilbert, S.B., Winer, E. & de la Cruz, J. (2012). Diagnosing perceptual distortion present in group stereoscopic viewing. *Proceedings of the SPIE 8291*, 82910P.

Glantz, K., Rizzo, A.A. & Graap, K. (2003). Virtual reality for psychotherapy: Current reality and future possibilities. *Psychotherapy: Theory, Research, Practice, Training, 40*, 55–67.

Grantcharov, T.P., Kristiansen, V.B., Bendix, J., Bardram, L., Rosenberg, J., & Funch-Jensen, Pl (2004). Randomized clinical trial of virtual reality simulation for laparoscopic skills training. British Journal of Surgery, 91, 146-150.

Held., R. T. & Banks, M. S. (2008). Misperceptions in stereoscopic displays: A vision science perspective. *Proceedings of the 5th symposium on applied perception in graphics and visualization* (pp. 23-31). New York: ACM.

Jack, D., Boian, R., Merians, A., Tremaine, M., Burdea, G. & Adamovich, S. (2001). Virtual reality-enhanced stroke rehabilitation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering, 9*, 308–318.

Kelly, J.W., Burton, M., Pollock, B., Rubio, E., Curtis, M., de la Cruz, J., Gilbert, S., & Winer, E. (in press) Space perception in virtual environments: Displacement from the center of projection causes less distortion than predicted by cue-based models. *ACM Transactions on Applied Perception*.

Loomis, J. M., Philbeck, J. W., & Zahorik, P. (2002). Dissociation of location and shape in visual space. *Journal of Experimental Psychology: Human Perception and Performance, 28*, 1202-1212.

Marbach, J. (2009). Image blending and view clustering for multi-viewer immersive projection environments. *Proceedings of the IEEE Virtual Reality Conference*.

Pollock, B., Burton, M., Kelly, J.W., Gilbert, S. & Winer, E. (2012). The right view from the wrong location: Depth perception in stereoscopic multi-user virtual environments. *IEEE Transactions on Visualization & Computer Graphics, 18*(4), 581-588.

Sedgwick, H. A. (1991). The effects of viewpoint on the virtual space of pictures. In S. R. Ellis, M. K. Kaiser, & A. C. Grunwald (Eds.), *Pictorial Communication in Virtual and Real Environments* (pp. 460–479). London: Taylor & Francis.

Woods, A., Docherty, T., & Koch, R. (1993). Image distortions in stereoscopic video systems. *Proceedings of SPIE: Stereoscopic Displays and Applications IV* (pp. 36-48).