Sensitivity to Scene Motion for Phases of Head Yaws

Jason Jerald[†] Tabitha Peck[†] Frank Steinicke[‡] Mary Whitton[†]

[†] Department of Computer Science, University of North Carolina at Chapel Hill, USA
email: (jjerald|tpeck|whitton)@cs.unc.edu

[‡] Visualization and Computer Graphics Research Group, University of Münster, Germany
email: fsteini@math.uni-muenster.de

Abstract

In order to better understand how scene motion is perceived in immersive virtual environments and to provide guidelines for designing more useable systems, we measured sensitivity to scene motion for different phases of quasi-sinusoidal head yaw motions. We measured and compared scene-velocity thresholds for nine subjects across three conditions: *visible With head rotation* (*W*) where the scene is presented during the center part of sinusoidal head yaws and the scene moves in the same direction the head is rotating, *visible Against head rotation* (*A*) where the scene is presented during the center part of sinusoidal head yaws and the scene moves in the opposite direction the head is rotating, and *visible at the Edge of head rotation* (*E*) where the scene is presented at the extreme of sinusoidal head yaws and the scene moves during the time that head direction changes.

The *W* condition had a significantly higher threshold (decreased sensitivity) than both the *E* and *A* conditions. The median threshold for the *W* condition was 2.1 times the *A* condition and 1.5 times the *E* condition. We did not find a significant difference between the *E* and *A* conditions, although there was a trend for the *A* thresholds to be less than the *E* thresholds. An Equivalence Test showed the *A* and *E* thresholds to be statistically equivalent.

Our results suggest the phase of user's head yaw should be taken into account when inserting additional scene motion into immersive virtual environments if one does not want users to perceive that motion. In particular, there is much more latitude for artificially and imperceptibly rotating a scene, as in Razzaque's redirecting walking technique, in the same direction of head yaw than against the direction of yaw.

The implications for maximum end-to-end latency in a head-mounted display is that users are less likely to notice latency when beginning a head yaw (when the scene moves with the head) than when slowing down a head yaw (when the scene moves against the head) or when changing head direction (when the head is near still and scene motion due to latency is maximized).

CR Categories: H.1.2 User/Machine Systems—Human Factors; I.3.6 Methodology and Techniques—Ergonomics and Interaction Techniques; I.3.7 Three-Dimensional Graphics and Realism—Virtual Reality; I.3.m [Computer Graphics]: Miscellaneous—Perception; J.4 [Computer Application]: Social and Behavioural Sciences—Psychology;

Keywords: human factors, psychophysics, motion perception, immersive virtual environments, latency, reorientation techniques

1 Introduction

A fundamental task of an immersive virtual environment (IVE) system is to present images that move appropriately on the display surface as the user's head moves, so that the elements of the virtual scene appear stable in space. Visual instability can be caused by latency, incorrect field of view, or by intentionally injecting scene motion into the IVE as is done in redirected walking [Razzaque *et al.* 2001], a technique that allows users to walk in IVEs larger than the tracked lab space. Noticeable visual instability can lessen an IVE experience by causing simulator sickness [Draper 1998], degraded task performance [So and Griffin 1995], degraded visual acuity [Allison *et al.* 2001], and a decreased sense of presence [Meehan *et al.* 2003].

Whereas error in IVEs are well defined mathematically [Adelstein *et al.* 2005; Holloway 1997], perception of these errors is not well understood. Motion perception when the head is held still is fairly well understood by vision researchers, but less is known about motion perception when the head is moving.

We use psychophysics methods to better understand perception of scene motion when turning the head for the purposes of improving the usability of IVEs and determining design requirements for future IVE systems. Specifically, we are interested in finding out how much scene motion can be present for different parts of head turns in IVEs without users noticing that scene motion. I.e., we measure sensitivity to motion for different phases of quasi-sinusoidal head turns.

The results of this work provide

- an improved understanding of of how we perceive scene motion with head movement
- a step toward better understanding of latency perception for users of head-mounted displays (HMDs), and
- guidelines on how much scene motion can be intentionally injected into IVEs (HMDs or projected systems) for reorientation purposes.

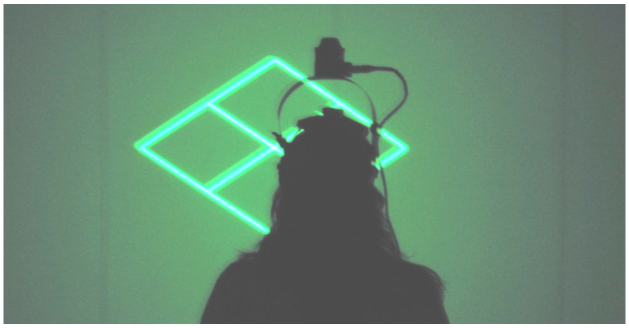

Figure 1. A subject and scene as the subject yaws her head. Note the motion blur occurs from the exposure time of the camera and the image has been brightened substantially.

APGV 2008, Los Angeles, California, August 9–10, 2008.
© 2008 ACM 978-1-59593-981-4/08/0008 $5.00

2 Background

From an egocentric perspective, the real world appears stationary even as we rotate our head and eyes. Extra-retinal cues come from other parts of the mind/body that help us perceive a stable world. These cues come from the vestibular system, proprioception, our cognitive model of the world, intent, etc [Razzaque 2005]. When one or more of these cues conflicts with other cues, as is often the case for IVEs, the virtual world may appear to be spatially unstable.

The primate brain has visual receptor systems dedicated to motion detection [Lisberger and Movshon 1999], which humans use to perceive optical motion [Nakayama and Tyler 1981]. While visual velocity (i.e., speed and direction) is sensed directly in primates and humans, visual acceleration is not, but is instead inferred through processing of velocity signals [Lisberger and Movshon 1999]. Most theorists assume that acceleration is not as important as velocity for motion perception [Regan *et al.* 1986]. Thus, in this experiment we set scene velocity to be constant (zero scene acceleration) on the display for each individual presentation. I.e., scene motion is independent of head motion and latency.

For this study we define the *scene-velocity threshold* to be the scene velocity (degrees/second) at which a subject is able to detect the presence of scene motion 75% of the time. We define *sensitivity* to be the inverse of threshold.

Sensitivity to scene motion is well studied when the head is stationary. It is generally understood that head motion diminishes sensitivity to scene motion, but scene-motion thresholds while the head moves have not been adequately explored.

The most relevant work is Wallach's summary of his own perceptual stability research [1987], which found that healthy subjects report visual environments appear unstable when the environments move more than ±2-3% of head motion. He reported that these values are the same whether the scene moved against or with head rotations. He did not report the types of head motion in his summary or in earlier work [Wallach and Kravitz 1965]. Steinicke *et al.* [2008] propose thresholds for human's sensitivity to scene-velocity when walking in general without focusing on head rotations. Their subjects did not notice up to 30% compression of head rotations, where the scene rotates with head rotation.

Probst *et al.* [1986] found reaction times to detection of a moving light spot increased as head motion increased. Loose and Probst [2001] found visual-motion thresholds to be significantly higher when the visual motion was with the direction of the head turn. They used random-dot kinematograms where visual-motion thresholds were measured in percentage of coherently moving pixels. Loose and Probst presented the moving visual stimuli in head-centric coordinates (stimuli moved relative to the head), and were judged object-relative to a head-stabilized target. Their conditions contrast with conditions in an HMD, where visual stimuli are judged to be moving in world coordinates and judgments are subject-relative.

Adelstein *et al.* [2006] and Li Li *et al.* [2006] also showed head motion suppresses perception of visual motion. Whereas an HMD was used for these experiments, no head tracking was used so that the image moved relative to head-centric coordinates.

One way users perceive end-to-end IVE latency is by scene motion that results from latency [Adelstein *et al.* 2003]; when users turn their heads, latency causes the visual scene to appear to "swim". Several NASA experiments [Adelstein *et al.* 2003; Adelstein *et al.* 2005; Ellis *et al.* 1999; Mania *et al.* 2004] measured latency thresholds for quasi-sinusoidal head yaws.

The NASA measurements showed no differences in latency thresholds for different scene complexities, ranging from single, simple virtual objects [Ellis *et al.* 2004] to detailed photorealistic rendered environments [Mania *et al.* 2004]. Steinicke [2008] also found no significant effect of visual appearances on motion thresholds. Thus, for this study we used a very simple geometric pattern for our scene and expect the results to generalize to more complex scenes.

Adelstein *et al.* [2005] found subjects to be more sensitive to latency at the edge of head yaws, i.e., when head speed is lowest and head direction changes, than the center of head yaws, when head velocity peaks. Latency perception for the different phases of sinusoidal head rotations consists of two components:

- Scene velocity due to latency peaks at the edge of quasi-sinusoidal head rotations and is near zero at the center of quasi-sinusoidal head rotations.
- Subjects may have different sensitivities to scene velocity for different phases of head rotations.

In this work, we specifically study subject sensitivity to scene velocity by having the scene velocity controlled by the experiment, rather than by head motion and latency.

3 Methods

We designed an experiment to emulate a zero-latency HMD. A BARCO CRT projector displayed a simple 2D visual scene (a rotated green monochrome square with diagonals and a 20º horizontal span from the subjects point of view) to subjects seated four meters from the screen (see Figure 1) as they yawed their heads in a sinusoidal pattern for four full head cycles (see Figure 2—middle portion). The simple test scene was chosen for the following reasons:

- Minimal drawing time is important for future latency experiments and we want the stimuli to be consistent across experiments.
- Prior work has not been able to find significant differences in scene-motion thresholds across scene complexities.
- A 2D scene was chosen to eliminate the possibility of depth issues (stereo cues, motion parallax, etc) that could confound our results.
- Vertical lines are more subject to tearing effects than diagonal lines. We eliminated vertical lines in the stimulus figures to maximize the probability that participants respond to real scene motion and not other visual artifacts caused by image tearing when the system does not wait on vertical sync.

The CRT projector was chosen for its fast phosphor response time so that no ghosting was present and no light was projected for black pixels, which is not the case for LCD and DLP projectors. The green monochrome illuminance was approximately one lux.

A Virtual Research V8 HMD was modified by removing the display elements so that subjects could see through the casing to the projector screen. This was done to limit the field of view to 48º horizontal by 36º vertical. All light sources in the lab were turned off in order to remove all object-relative cues; only the computer-generated scene was visible. A bright green screen (approximately ten lux) was shown between trials to prevent dark adaptation, so that brightness sensitivity would be consistent across trials. Subjects sat four meters away from the planar screen. The left and right edges of the screen were curved inward so that the screen surrounded the subjects by nearly 180 degrees. The floor was covered with the same material. Subjects confirmed they could not see any visual elements other than the scene presented by the projector.

The horizontal starting position of the stimulus for each presentation was randomly set within ±3.2° of the screen center. Scene motion was controlled by the experiment and was independent of head movement; head position/orientation did not affect the position or motion of the visual scene. Head orientation was determined by a 3rdTech HiBall 3000 tracking system. The tracking data were used to control auditory cues, to check for acceptable head rotations, and to record motion for potential post analysis. Total weight of the modified HMD and tracker was 0.6 kg.

We asked subjects to turn their heads from side-to-side for four full head cycles, with a head amplitude of 22° off center, to the pacing of metronome beeps. We trained subjects using both visual and auditory cues; the experiment room was otherwise kept quiet. During training and data collection, the system played a clicking sound when the subject changed head direction. We asked the subjects to sync the clicking sound with the experiment-controlled metronome beeps. If the subject-controlled clicking got out of sync from the metronome beeps by more than 0.25 second, the trial was cancelled and a new trial was selected. A second tone sounded when their head yaw exceeded the 6° minimum head-yaw amplitude. Subjects were trained to move their head far enough to hear this tone at the edge of head yaws. If this minimum head-yaw amplitude was not reached when head rotation direction changed, then the trial was canceled. If the subject yawed their head beyond 14° at any time, a buzzing sound occurred and the trial was cancelled. No communication between the experimenter and the subject occurred during trials.

To increase our ability to generalize our results across head frequencies, we measured scene-velocity thresholds for three head frequencies: 0.35 Hz, 0.5 Hz, and 0.65 Hz corresponding to one side-to-side head swing in 2.9 seconds, 2.0 seconds, and 1.5 seconds. Head frequency variation was a controlled between-subjects variable, with three subjects per head frequency. The top element of Figure 2 shows specified and actual yaw-head orientation over time for several head rotations for a single subject.

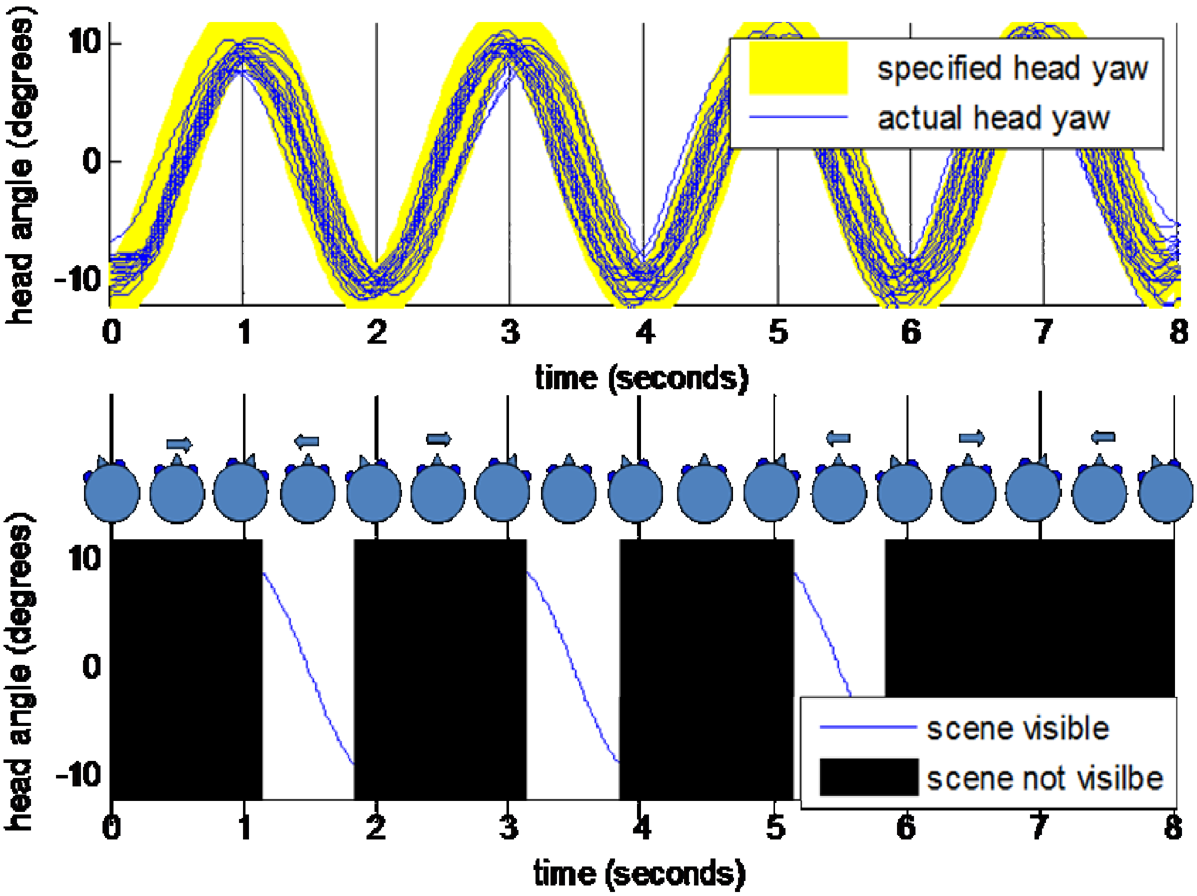

Figure 2. The top part of the diagram shows specified and actual yaw-head rotations for 20 trials with a head frequency of 0.5 Hz. The subject starts the trial with the head turned to the left (-11° head yaw angle) and turns the head to the right then back to the left with a sinusoidal head yaw. The icons in the middle of the diagram represent a top view of head yaw and the direction the head is turning. The bottom part of the diagram shows a *visible With head rotation* (*W*) example trial. The scene is presented three times for 0.5 seconds each as the head yaws from the right (+11° head yaw angle) to the left (+11° head yaw angle). The scene is not visible for the black area of the diagram. The scene moves from the right to the left for the second or third presentation, but not for both presentations. At the end of the trial, the subject selects the second or third presentation that she believes contains scene motion.

We measured scene-velocity thresholds for different phases of sinusoidal yaw-head rotation within subjects. We did this by making the constant-velocity scene visible for only parts of the head rotation. These visibility conditions were

- ***Visible With Head Rotation* (*W*),** where the scene is presented during the center part of quasi-sinusoidal head yaws and the scene moves in the same direction the head is rotating,
- ***Visible Against Head Rotation* (*A*),** where the scene is presented during the center part of quasi-sinusoidal head yaws and the scene moves in the opposite direction the head is rotating, and
- ***Visible at the Edge of Head Rotation* (*E*),** where the scene is presented at the extent of quasi-sinusoidal head yaws and the scene moves during the time that head direction changes.

The bottom part of Figure 2 shows an example trial of when the scene was visible for the *W* condition. We asked subjects to keep their eyes centered at the center of the target so that the scene was seen mostly in the foveal region of the eye. However, since the scene appeared randomly within ±3.2° of center and we were not able to measure eye movements, we do not know how well subjects kept their eyes centered.

Trial procedures were similar to previous NASA latency experiments [Adelstein *et al.* 2003; Adelstein *et al.* 2005; Ellis *et al.* 1999; Mania *et al.* 2004] that used a two-down, one-up adaptive descending staircases algorithm [Levitt 1970] (so that judgments theoretically converged to 70.7% correct responses) to determine perceptual thresholds. However, we used a selection task instead of the same-different task used in the NASA experiments. A selection task asks the question "Which of multiple presentations contains the stimulus?", whereas a same-different task asks the question "Are the presentations the same or different?" A same-different task contains subject bias for absolute thresholds, since subjects' definitions of *different* are subjective. A selection task, as used in this study, has little bias for absolute thresholds when presentation order is randomized[1].

For each trial, the two-alternative forced-choice selection task consisted of three stimulus presentations. The first presentation was a reference stimulus containing no scene velocity so that subjects would know what a stable scene looks like. Some scene velocity, determined by the adaptive staircase algorithm, was randomly assigned to the second or third presentation, with the remaining presentation containing no scene velocity. The subject then selected, via a mouse, which of the latter two presentations she believed contained scene velocity (i.e., which of the latter two presentations was different from the first presentation?). To encourage good performance, we rewarded subjects $0.05 for every correct response. After each response the system informed the subjects whether they were correct.

If subjects could not determine which presentation contained scene motion, then they had equal probability of choosing either of the two latter presentations. This 50% guessing rate is called the *point of subjective equality* (PSE). We define the detection threshold to be the halfway point between the PSE and 100% detection. I.e., a subject's detection threshold is the scene velocity at which she correctly chooses the stimulus presentation that contains scene velocity 75% of the time.

For each subject, six sessions were conducted over one or more days. We randomly interleaved three staircases, one staircase for each visibility condition, during each session. This interleaving of conditions minimized order effects, and made it difficult to differentiate among the conditions. Each staircase terminated after eight staircase reversals, resulting in each subject judging a total of 148 to 219 trials for each of the three visibility conditions.

Nine subjects (age 18-44, 7 male and 2 female, average age = 30) participated. Subjects came from backgrounds ranging from students to professionals. One of the authors served as a subject[2]; all other subjects were naïve to the experimental conditions. Subjects were allowed to take breaks at any time. Total time per subject, including consent form, instructions, training, experiment sessions, breaks, and debriefing, took three to six hours hours with an average time of approximately five hours.

4 Results

For each subject and condition, proportions of correct responses were computed for the different scene-velocities presented. A cumulative Gaussian psychometric function was fit to these proportions. A minimum of two judgments at a scene velocity were required for the proportion to contribute to the fit. In addition, a theoretical proportion of 0.5 for zero scene velocity contributed to the fit. A finger error (the probability of a subject pushing the wrong button by mistake) of 2% was used for the fit. The Gaussian distribution's mean yields the scene-velocity threshold value of 75%. The Gaussian variation is related to the steepness of the psychometric function.

Scene-velocity thresholds were computed for each of the three visibility conditions for each of nine subjects. Figure 3 shows a single subject's proportion of correct responses for scene-velocities in the *W* condition along with the psychometric function fit to these proportions. This subject had a *W* scene-velocity threshold of 1.87 degrees/second.

Figure 4 shows percentage of correct responses and psychometric functions for all three visibility conditions for a single subject at a head frequency of 0.35 Hz. The 27 psychometric function fits to a cumulative Gaussian function (9 subjects times 3 visibility conditions) were all statistically significant (all $p < 0.05$) with the single lowest goodness of fit Pearson correlation of r = 0.56, the next eight with $0.71 < r < 0.87$, and the remainder at $r > 0.90$.

Scene-velocity thresholds for the three visibility conditions from all nine subjects are reported in Figure 5. Although the scenes were presented on a planar surface with scene velocity in units of meters per second, thresholds have been converted to degrees per second for a more intuitive understanding. It is visually evident that thresholds are higher (i.e., lower sensitivity) for the *W* condition. The slopes of the scene velocity threshold curves increase (not shown) as thresholds increase, implying that the equal variance requirements for parametric tests do not hold. Thus non-parametric tests are reported throughout.

Friedman analyses of variance (ANOVA) shows that scene-velocity thresholds were significantly affected by the visibility conditions (Q = 16.22, $p < 0.001$). Scene-velocity thresholds were greater for the *W* condition than for both the *A* and *E* conditions (both Wilcoxon Matched-Pairs Signed-Ranks Tests: S = 0, $p < 0.01$ two-tail, effect size $Z/\sqrt{N} = 0.89$). The scene-velocity thresholds median was 2.1 times larger under the *W* condition than the *A* condition and 1.5 times larger than the *E*

---

[1] A same-different task was used in the NASA experiments because difference thresholds were the primary interest, whereas we were interested in absolute thresholds.

[2] Experimenters often serve as their own subjects in psychophysics studies.

condition. There was a trend for the *A* condition to be smaller than the *E* condition but this difference was not significant at the $\alpha = 0.05$ level (Wilcoxon Matched-Pairs Signed-Ranks Test: $S = 6$, $p < 0.055$ two-tail).

Since there was no statistical difference between the *A* and *E* conditions, we conducted an equivalence test across these two conditions. An equivalence test does not check if two conditions are the same, but instead checks that the results are "close enough" where "close enough" is chosen by the researchers based on a subjective judgment of importance. We defined "close enough" to be the difference of the peak theoretical scene velocities for the *A* and *E* conditions that would occur in an HMD with 50 ms of latency (typical latency for such a system) at 0.5 Hz of quasi-sinusoidal head motion. This value is 1.9°/sec. Thus thresholds are considered to be equivalent if the threshold differences are within a ±1.9°/sec range. The *A* and *E* conditions were statistically equivalent (Wilcoxon Matched-Pairs Signed-Ranks Equivalence Test: $S = 0$, $p < 0.01$).

We computed the ratio of scene-velocity thresholds to theoretical peak head velocities for the *A* and *W* conditions. Figure 6 shows these ratios. These ratios ranged from 2.2% to 7.7% (median of 5.2%) for the *A* condition and 7.7% to 23.5% (median of 11.2%) for the *W* condition. The ratio threshold median for the *W* condition was 2.1 times larger than the ratio threshold median for the *A* condition.

5 Discussion

5.1 External Validity Required

We believe the specific threshold values found in this experiment are not the significance of this research, instead the relative difference between conditions is what is important. Different conditions such as training effects, type of head turns, contrast, scene size, field of view, etc. might result in different thresholds, but we suspect that the ordering of thresholds for the visibility conditions would remain consistent across conditions other than visibility.

A previous study did not find a significant effect of visual appearances on the evaluation of redirected walking, but there was a significant effect of optical flow (subjects realized redirected walking more when more detail was apparent in the scene) [Steinicke *et al.* 2008]. Further experimentation would be required to know if our scene-velocity thresholds are externally valid (i.e., do the results hold across other settings?).

Although we instructed subjects to keep their eyes centered at the approximate center of the screen, we did not record eye movements. It is possible subjects may have used more of their periphery vision to make judgments of the *E* condition.

Subjects took three to six hours to complete the experiment. Although we encouraged subjects to take breaks at any time, we suspect fatigue may have inflated threshold values. We noticed that performance did not seem to consistently relate to the session number across subjects: some subjects performed better during the early sessions, some performed better during the middle sessions, and some performed better during later sessions.

We measured scene-velocity thresholds for sinusoidal head yaws, so that the scene-velocity thresholds for the *E* condition was for a change of head direction. Thresholds, especially for the *E* condition, might be quite different for half-cycle head turns where the head starts accelerating from a stopped angle and then decelerates to a resting angle.

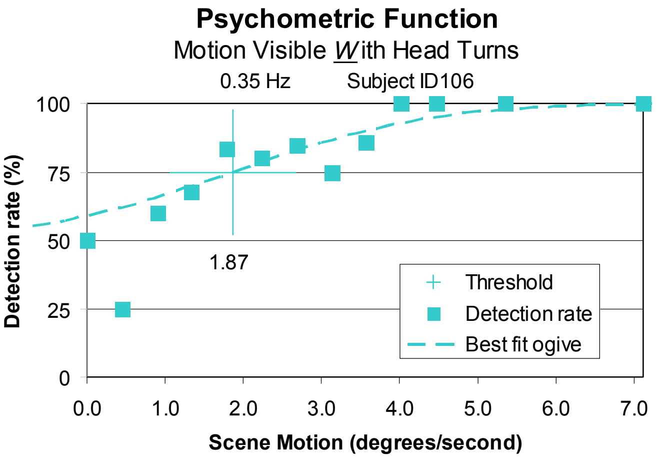

Figure 3. Data collected from a single subject for the *W* condition. The curve is a psychometric function determined by the best fit cumulative Gaussian function. The curve fits the data with a correlation of $r = 0.95$.

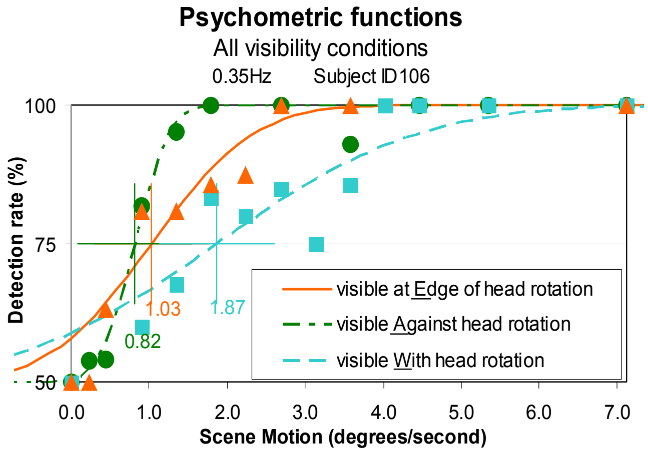

Figure 4. Psychometric functions for all visibility conditions for a single subject Note that the Y range of the graph is 50%-100%.

5.2 Psychometric Fits

Varying the finger error and/or the minimum number of judgments per scene velocity required to contribute to the psychometric fit would result in slightly different thresholds. Other psychometric functions could have been used that would have also resulted in different thresholds. We could have also constrained the psychometric function to cross the y-axis within some value near 50% since the psychometric function has a theoretical value of 50% for zero scene velocity. However, such a constraint would come at the cost of a lower correlation of the psychometric function to the proportions.

5.3 Comparison to Previous Results

The results of our study differ from Wallach's finding that sensitivity to scene motion is the same whether scene motion is with head rotation or against head rotation (a range of 2% to 3% of head motion) [Wallach 1987]. Wallach provided little detail on the specifics of the head motion or other aspects of his experiment, so we cannot speculate why our results differ.

Steinicke *et al*. [2008] found subjects did not notice up to 30% compression of head rotations, which was similar to our *W* condition. Several factors could contribute to the large 30% threshold compared to our *W* median threshold of 11.2% and Wallach's 2-3% threshold including the following:

- The wording of the question may have biased results. Subjects were asked if scene motion was not perceivable, slightly perceivable, noticeable, or strongly perceivable.
- Scene motion due to latency in the HMD system could have confounded scene motion thresholds.

5.4 How Subjects Made Their Judgments

We expected subjects to judge scene motion directly from the visual stimuli. However, approximately half the subjects reported making judgments about the smaller motions based on a vection-like sensation (i.e., self motion). Some subjects commented on this orally and some reported this in written form for the exit questionnaire. Two subjects wrote in response to "Explain in your own words how you went about making the scene-motion judgments.":

> *"Most of the times I detected the scene with motion by noticing a slight dizziness sensation. On the very slight motion scenes there was an eerie dwell or "suspension" of the scene; it would be still but had a floating quality." (Subject ID103)*
>
> *"At the beginning of each trial [sic:session—when scene velocities were greatest], I used visual judgments of motion; further along [later in sessions] I relied on the feeling of motion in the pit of my stomach when it became difficult to discern the motion visually." (Subject ID109)*

5.5 Application to Reorientation Techniques

Reorientation techniques enable users to walk in larger-than-lab sized IVEs by imperceptibly rotating the IVE around the user. Razzaque's maximum rotation values [Razzaque et al. 2001; Razzaque 2005] were determined by informal evaluation of imperceptibility for a small group of users. The amount of added rotation could be determined by an algorithm using maximum rotation parameters. The thresholds reported in this paper, or thresholds obtained through similar psychophysical measures, could be used for these maximum rotation parameters.

The ratios of scene-velocity thresholds to head velocities for the different conditions reported at the end of Section 4 provide guidelines for a maximum imperceptible rotation parameter. The amount of injected scene rotation would be a function of head rotation velocity (speed and direction). We could choose the minimum thresholds obtained for each condition to be the maximum amount of rotation that would be allowed to be used in a reorientation technique. Based on our results, the scene could be rotated against head rotation at a speed up to 2.2% of head rotation speed (the minimum threshold value reported for condition *A*), or with head rotation at a speed up to 7.7% of head rotation speed (the minimum threshold value for condition *W*).

Alternatively the median values could be used such that the scene would be rotated against head rotation at a speed up to 5.2% of head rotation speed and with head rotation at a speed up to 11.2% of head rotation speed.

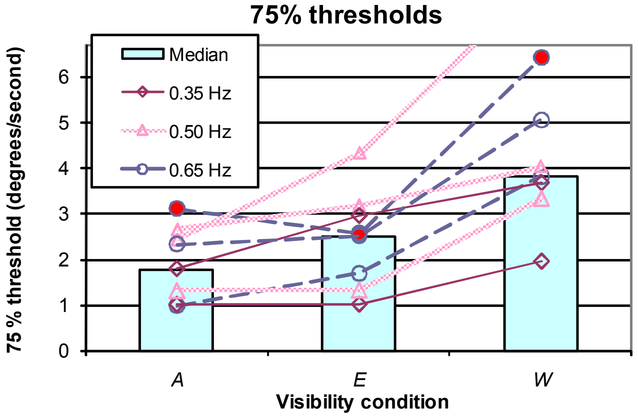

Figure 5. Thresholds for all nine subjects for the three visibility conditions. Note the filled Os indicate the sole subject whose thresholds decreased from the *A* condition to the *E* condition. Bars indicate medians for each condition.

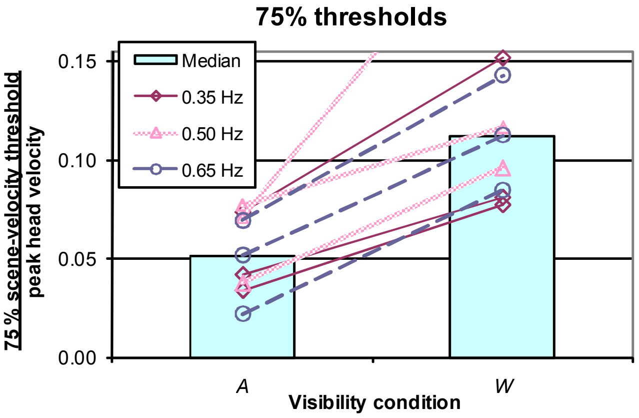

Figure 6. Ratios of scene-velocity thresholds to peak head velocities.

The findings that subjects are least sensitive to a scene moving with the head suggest distractors used for reorientation purposes [Peck *et al.* 2008] could be used most effectively by moving the distractor in the same direction we wish to rotate the world. We also suspect the thresholds would be larger if subjects pay attention to a distractor or some distracting task—further experimentation would be required to confirm this.

5.6 Latency Perception

Latency in an HMD causes the visual scene to move with the user's head until the system "catches up". The scene moves with the head until shortly after head acceleration goes to zero (constant head velocity) or the head starts to decelerate. With some constant head velocity (zero head acceleration), the scene appears to be stable in space (i.e., zero scene-velocity) with a constant offset from where it would correctly appear with no head motion. As the head decelerates, the scene starts to move back to where it should be in space, moving against the head turn. Maximum scene velocity occurs near the edge of head turns, when head acceleration peaks. These facts combined with our results suggest users are less likely to notice latency in a head-mounted display when beginning a head turn than when slowing down a head turn or changing head direction.

Our results show subjects to be equally sensitive (or more sensitive as the trend suggests) to the *A* condition as the *E* condition. This suggests that the reason subjects are most sensitive to latency at the edge of head yaws [Adelstein et al. 2005] is because scene velocity due to latency is maximized at the edge of head yaws, not because sensitivity to scene motion is maximized at the edge of head yaws. Although Adelstein et al. compared a condition similar to our *E* condition to a condition similar to our *W* and *A* conditions combined, we suspect their combination was more equivalent to our *A* condition, because a person is going to perceive, and thus judge, the most apparent portion of the scene motion (the *A* portion).

5.7 Future Work

We plan to continue studying motion thresholds as a function of various parameters through experimentation. In particular we are interested in further relating scene-motion thresholds to latency thresholds. We also plan to integrate our results into reorienting systems.

6 Acknowledgements

Jason Jerald has been supported by a LINK Foundation Fellowship and an NC Space Grant Fellowship. Professor Whitton is supported by grants from the Office of Naval Research and the NIH National Institute for Biomedical Imaging and Bioengineering. We thank Fred Brooks, Neeta Nahta, and the anonymous reviewers for their reviews of this paper. We thank David Harrison and John Thomas for their technical and facilities support. We thank Steve Ellis and Bernard Adelstein of NASA Ames Research Center for their numerous discussions and suggestions. We also thank the members of the UNC Effective Virtual Environments (EVE) research group for their support.

7 References

ADELSTEIN, B. D., LEE, T. G., AND ELLIS, S. R. 2003. Head Tracking Latency in Virtual Environments: Pyschophysics and a Model. Proceedings of *the 47th Annual Meeting of the Human Factors and Ergonomics Society*, pp. 2083-2987.

ADELSTEIN, B. D., BURNS, E. M., ELLIS, S. R., AND HILL, M. I. 2005. Latency Discrimination Mechanisms in Virtual Environments: Velocity and Displacement Error Factors. Proceedings of *the 49th Annual Meeting of the Human Factors and Ergonomics Society,* pp. 2221-2225.

ADELSTEIN B.D., LI L., JERALD, J, AND ELLIS, S.R. 2006. Suppression of Head-Referenced Image Motion During Head Movement. Proceedings of *the 50th Annual Meeting of Human Factors and Ergonomics Society*, pp. 2678-2682.

ALLISON, R. S., HARRIS, L. R., JENKIN, M., JASIOBEDZKA, U., AND ZACHER, J. E. 2001. Tolerance of Temporal Delay in Virtual Environments. Proceedings of *IEEE Virtual Reality*, pp. 247-253.

DRAPER, M. H. 1998 *The Adaptive Effects of Virtual Interfaces: Vestibulo-Ocular Reflex and Simulator Sickness*, PhD Thesis, University of Washington.

ELLIS, S. R., YOUNG, M. J., ADELSTEIN, B. D., AND EHRLICH, S. M. 1999. Discrimination of Changes in Latency During Head Movement. Proceedings of *Computer Human Interaction*, pp. 1129-1133.

ELLIS, S. R., MANIA, K., ADELSTEIN, B. D., AND HILL, M. I. 2004. Generalizeability of Latency Detection in a Variety of Virtual Environments. Proceedings of *the 48th Annual Meeting of the Human Factors and Ergonomics Society*, pp. 2083-2087.

HOLLOWAY, R. 1997. Registration Error Analysis for Augmented Reality. *Presence: Teleoperators and Virtual Environments,* 6(4), pp. 413-432.

LEVITT, H. 1970. Transformed up-Down Methods in Psycho-Acoustics. *Acoustical Society of America,* 49(2), pp. 467-477.

LI, L., ADELSTEIN, B.D., AND ELLIS, S.R. 2006. Perception of Image Motion During Head Movement. Proceedings of *ACM Applied Perception in Graphics and Visualization*, pp. 45-50.

LISBERGER, S. G. AND MOVSHON, J. A. A. 1999. Visual Motion Analysis for Pursuit Eye Movements in Area MT of Macaque Monkeys. *The Journal of Neuroscience,* 19(6), pp. 2224-2246.

LOOSE, R. AND PROBST, T. 2001. Velocity Not Acceleration of Self-Motion Mediates Vestibular-Visual Interaction. *Perception,* 30(4), pp. 511-518.

MANIA, K., ADELSTEIN, B. D., ELLIS, S. R., AND HILL, M. I. 2004. Perceptual Sensitivity to Head Tracking in Virtual Environments with Varying Degrees of Scene Complexity. Proceedings of *ACM Applied Perception in Graphics and Visualization*, pp. 39-47.

MEEHAN, M., BROOKS, F., RAZZAQUE, S., AND WHITTON, M. 2003. Effects of Latency on Presence in Stressful Virtual Environments. Proceedings of *IEEE Virtual Reality*, pp. 141-148.

NAKAYAMA, K. AND TYLER, C. W. 1981. Psychological Isolation of Movement Sensitivity by Removal of Familiar Position Cues. *Vision Research,* 21, pp. 427-433.

PECK, T., WHITTON, M., AND FUCHS, H. 2008. Evaluation of Reorientation Techniques for Walking in Large Virtual Environments. Proceedings of *IEEE Virtual Reality,* pp. 121-172.

PROBST, T., BRANDT, T., AND DEGNER, D. 1986. Object-Motion Detection Affected by Concurrent Self-Motion Perception: Psychophysics of a New Phenomenon. Behavioral Brain Research, 22, pp. 1-11.

RAZZAQUE, S., KOHN, Z., AND WHITTON, M. 2001. Redirected Walking. Proceedings of *Eurographics*, pp. 289-294.

RAZZAQUE, S. 2005 *Redirected Walking*, PhD Thesis, University of North Carolina at Chapel Hill.

REGAN, D. M., KAUFMAN, L., AND LINCOLN, J. 1986 "Ch. 19, Motion in Depth and Visual Acceleration." In K. R. Boff, L. Kaufman, & J. P. Thomas (Eds.), *Handbook of Perception and Human Performance,* Wiley-Interscience, New York.

SO, R. H. Y. AND GRIFFIN, M. J. 1995. Effects of Lags on Human Operator Transfer Functions with Head-Coupled Systems. *Aviation, Space and Environmental Medicine,* 66(6), pp. 550-556.

STEINICKE, F., BRUDER, G., ROPINSKI, T., HINRICHS, K. 2008. Moving Towards Generally Applicable Redirected Walking. Proceedings of *Virtual Reality International Conference,* pp. 15-24.

WALLACH, H. AND KRAVITZ, J. H. 1965. The Measurement of the Constancy of Visual Direction and of Its Adaptation. *Psychonomic Science,* 2, pp. 217-218.

WALLACH, H. 1987. Perceiving a Stable Environment When One Moves. *Annual Review of Psychology,* 38, pp. 1-27.