

# HMD calibration and its effects on distance judgments

Scott A. Kuhl\*  
University of Utah

William B. Thompson  
University of Utah

Sarah H. Creem-Regehr  
University of Utah

## Abstract

Most head-mounted displays (HMDs) suffer from substantial optical distortion, and vendor-supplied specifications for field-of-view often are at variance with reality. Such displays do not present perspective-related visual cues in a geometrically correct manner, which has the potential to affect applications of HMDs which depend on precise spatial perception. This paper provides empirical evidence for the degree to which these resulting distortions affect one type of spatial judgment in virtual environments. We show that minification in the HMD that would occur from an overstated HMD field of view results in a significant change in distance judgments. Incorrectly calibrated pitch and pincushion distortion, however, did not cause statistically significant changes in distance judgments for the degree of distortions examined. While the means for determining the optical distortion of display systems are well known, they are often not used in non-see-through HMDs due to problems in measuring and correcting for distortion. As a result, we also provide practical guidelines for creating geometrically calibrated systems.


**CR Categories:** H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems—Artificial, augmented, and virtual realities; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Virtual reality

**Keywords:** immersive virtual environment, perception, minification, field of view, pincushion distortion, pitch

## 1 Introduction

Virtual environment (VE) systems are computer interfaces that provide users with the sensory experience of being in a simulated space. These systems frequently consist of a head-mounted display (HMD) that allows people to view and move within a computer generated environment. HMD-based VEs can be used for a variety of applications including training, prototyping, education, entertainment, and research. If a VE is used for a particular application, the overall utility of the system often depends on the ability of the VE to correctly convey the virtual world to the user. For instance, if an architect wants to view a virtual version of a building in a VE, it is important that the architect can perceive the same building as the computer is displaying. Uncalibrated HMD systems can distort perspective-related visual cues that may prevent people from

\*e-mail: skuhl@cs.utah.edu



**Figure 1:** The angle of declination  $\theta$  and eye-height  $h$  can be used to calculate absolute distance to objects on the floor.

properly perceiving the virtual world.

We are particularly interested in learning how geometric distortions from uncalibrated HMDs might influence people's ability to make distance judgments in virtual environments. It is well documented that people frequently underestimate absolute distances to targets on the floor in HMD-based VEs at distances of 3 to 15 meters [Witmer and Kline 1998; Knapp 1999; Loomis and Knapp 2003; Thompson et al. 2004; Sahm et al. 2005; Willemsen et al. 2008; Willemsen et al. in press]. Although different studies have found varying amounts of compression, people typically judge these targets to be 70 to 90% of their intended distance in HMDs. None of these papers provide detailed information about if and how their HMDs were calibrated before being used for distance judgments. Since calibration procedures are not reported, it is difficult to determine if differences in calibration might cause some differences between laboratories. Therefore, understanding if and how HMD calibration might affect distance judgments can help us compare previous studies and can indicate to researchers which type of calibrations can result in significant changes in distance judgments.

The *angle of declination* can provide sufficient information for people to judge the absolute distance to targets on the floor if eye height is known (Figure 1). There is evidence that manipulating this cue can change the perceived distance to targets in both real [Stoper 1999; Gardner and Mon-Williams 2001; Ooi et al. 2001] and virtual [Messing and Durgin 2005] environments. We will focus, therefore, on measuring and calibrating three aspects of the HMD which can cause changes to the angle of declination. First, virtual worlds can easily become incorrectly pitched up or down if the orientation sensor for the HMD is not accurately aligned with the optical axis of display. Second, the optics in HMDs often introduce distortions which cause straight lines to appear curved. Third, the graphics displayed to the user can become minified or magnified if the field of view of the display is not accurately known. We will discuss each of these distortions in detail and present methods to measure and correct them. We will also compare distance judgments in a calibrated HMD and compare them to judgments in an incorrectly calibrated HMD.

See-through HMDs are usually carefully calibrated to allow items in the virtual world to properly align with objects seen in the real world [Azuma and Bishop 1994; McGarrity and Tuceryan 1999; Gilson et al. in press]. Non-see-through displays are more difficult to calibrate because one cannot easily use objects in the real world and directly compare them to the virtual world. Unlike see-through displays where poor calibration could be distracting, it can be diffi-

cult to recognize incorrect calibration in non-see-through displays without careful measurement.

## 2 Distance judgments in distorted environments

We measured how several geometric distortions in an HMD might influence distance judgments with a series of experiments. All experiments used an NVIS nVisor SX HMD and an IS-900 tracking system. Unless otherwise stated, we used our HMD with the pitch, pincushion distortion, and field of view calibrated to produce geometrically correct displays, using the methods described in Section 3.

### 2.1 Distance judgment procedure

All of the distance judgments in this paper used a direct blind walking task to measure the perceived distance to targets on the floor. During the entire experiment, subjects wore noise-canceling headphones that allowed them to hear the experimenter who was speaking into a wireless microphone. The subjects could also hear static in the headphones which was designed to block out any of the ambient sounds in the lab. Prior to the experiment, subjects were assisted in walking around blindfolded for several minutes outside the laboratory to familiarize them with walking without vision. Next, the experimenters brought the subjects into the lab blindfolded so that they would not see the lab prior to the experiment. Once in the lab, the subjects removed the blindfold and kept their eyes closed until the experimenter placed the HMD on their head. Then the subjects could open their eyes to see a blank screen in the HMD. The experimenter then assisted the subject in adjusting the HMD to comfortably fit on their head.

Next, we unblanked the screen and allowed subjects to view a virtual hallway (shown in Figure 2). Subjects viewed the target on the floor in front of them for as long as necessary to create a mental image of the space. They were allowed to rotate their heads to view the space but were not allowed to translate their heads or move their heads by bending at their waist. When the subjects felt confident in their mental image of the space, they closed their eyes, verbally stated that they were ready to walk, and walked to the target with their eyes closed. The HMD screens were blanked by the experimenter immediately after the subject stated that they were going to walk toward the target. The subjects stopped when they thought that they were standing directly on top of the target and the experimenter recorded their location with the tracking system. The subjects were not given feedback on their performance. Next, subjects kept their eyes closed and they were assisted in walking back to a starting location. We used three different virtual world starting locations which were each separated by one meter to make it more difficult for people to determine the target distance simply by looking at the target's location relative to other objects in the hallway. Subjects walked to four different target distances (3, 4, 5, and 6m) once from each of the three starting locations. In addition to these 12 trials, we included two practice trials at distances of 3.5 and 5.5 meters at the beginning of the experiment for subjects to familiarize themselves with the blind walking task. No feedback was provided for these practice trials. We also randomly inserted three other trials within the 12 trials at distances of 3.5, 4.5, and 5.5m to make it more difficult for subjects to memorize the distances we were showing them. As a result, only the data from 12 of the 17 trials the subjects performed was included in the analysis. The subjects were

not told that some of the trials would not be used. We also randomized the target shape, size, and color throughout the experiment.


### 2.2 Pitched environment

Tracked head-mounted displays typically use some type of sensor to measure the orientation of the HMD. If the sensor is perfectly aligned with the optical axis of the display, one can simply read the pitch, yaw, and roll values directly from the sensor to determine the HMD orientation. It is often unknown, however, if the sensor is accurately aligned with the optical axis of the display. We will focus on the calibration of pitch (up and down rotations) because there is evidence [Gardner and Mon-Williams 2001; Ooi et al. 2001] that pitching the world with prism glasses in the real world can change distance judgments. Franz et al. [2004] also showed that a similar manipulation that both pitches the world and preserves vertical lines can change perceived distance in a desktop display. Pitching the world does not change the angle between the targets on the floor and the visual horizon. The angle is changed, however, between the horizontal with respect to gravity (i.e., true horizontal in the real world) and the targets. Therefore, pitching the world up or down would change the angle of declination (Figure 1) with respect to gravity and could subsequently make objects on the floor appear farther or closer than they actually are.

In this experiment, we used a between-subject experimental design with three different visual conditions that all used the same virtual hallway environment. In the first visual condition, we displayed the virtual world using all of the calibrations discussed in this paper (Figure 2(a)). We also had two pitch conditions which pitched the world either up or down 5.7 degrees around the eye point (Figure 2(b)). This amount of pitch corresponds to the Ooi et al. [2001] experiment that used prism glasses to pitch the real world down 5.7 degrees. Thirteen subjects participated in the baseline (no pitch) condition and twelve participated in each of the pitched conditions.

#### 2.2.1 Results and discussion


The results of the baseline and pitched conditions are shown in Figure 3. Subjects walked 80, 84, and 80 percent of the actual target distance in the baseline, pitch up, and pitch down conditions respectively. A 3(condition) x 4(distance) ANOVA confirmed that there was no significant difference in judged distance due to the pitched environment ( $F(1, 23) = 0.45, p = 0.51$ ). Although Ooi et al. [2001] found that pitch affected distance judgments in a real environment by using prism glasses, these results indicate that this effect does not occur in HMD-based VEs for the range of distances and pitch values used in this experiment. There are a few possible explanations for this difference. First, people rapidly adapt to a pitched environment and this adaptation may occur more quickly in HMD-based virtual environments. Second, the weight of the HMD may add additional uncertainty to the ability of people to use proprioceptive and vestibular information to measure pitch. This uncertainty may result in an increased reliance on visual information to determine the location of the horizon. Third, Ooi et al. found that prism glasses had the greatest effect on distance judgments at distances of 6m and larger and produced relatively smaller effects at shorter target distances. Although we found no significant change in distance judgments at 6m, it may be possible that the minimum distance at which pitch affects distance judgments is greater for virtual environments than for the real world.



(a) Baseline

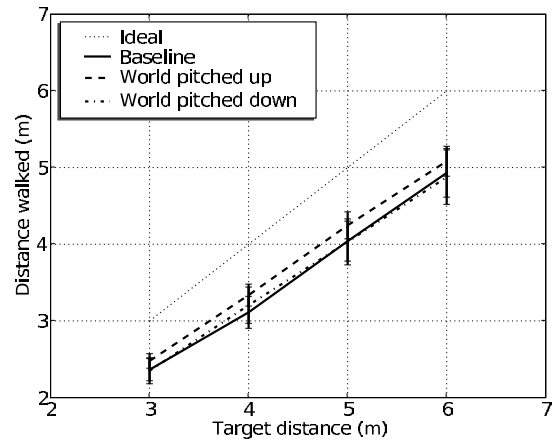


(b) World pitched up



(c) Minified

**Figure 2:** Screenshots of the virtual environment used for distance judgments



**Figure 3:** Pitching a virtual world up or down does not influence distance judgments measured with direct blind walking.

### 2.3 HMD with pincushion distortion

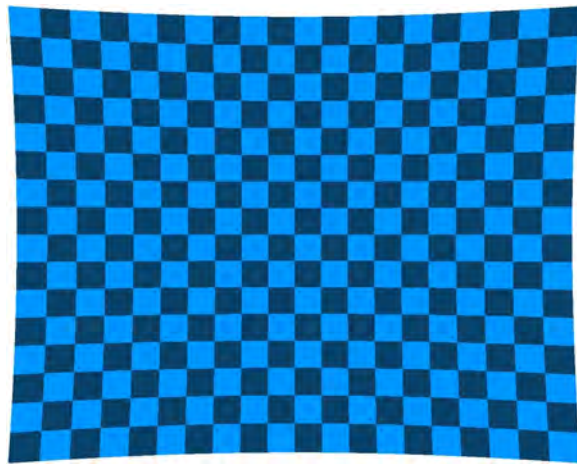
HMDs frequently have collimated optics between the display and the viewer to allow the user's eyes to focus as if the screens were farther away than they actually are. These optics can cause pincushion distortion—a radial distortion that occurs when points farther from the center of the image are magnified more than the points near the center. This distortion causes every line that does not cross the center of the image to be curved inwards toward the center of the image. Figure 4 illustrates how a checkerboard image appears in our HMD before and after pincushion distortion correction.

When there is not a visible horizon in an environment, linear perspective cues can provide information that can be used to determine the location of the effective horizon. For example, the hallway environment used for the present work has converging lines on edges of the floor and ceiling that can be used to determine where the effective horizon is. When pincushion distortion is applied to such a scene, these lines will become curved and, depending on where people look, can make the effective horizon appear higher or lower in the scene.


We used the same experiment procedure described in Section 2.1 to measure the effect of pincushion distortion on distance judgments. In this experiment, we performed a new visual condition that did not correct for the pincushion distortion in our HMD. Twelve subjects participated in the uncorrected condition. See Section 3.2 for a detailed discussion of our HMD's pincushion distortion. We compared the results of this experiment to the same baseline data used for the pitch experiment. Both the baseline and uncorrected pincushion conditions used the same black frame at the corners of the image (shown in Figure 4(b)) to ensure that the image displayed to the subject was rectangular in both conditions.

#### 2.3.1 Results and discussion

Figure 5 shows the results of the baseline and pincushion distorted conditions. Subjects walked 80 and 86 percent of the actual target distance in the baseline and uncorrected pincushion condition respectively. Although the data suggests that there may be a difference in these two conditions, a 2(condition) x 4(distance) ANOVA found that there was no significant difference ( $F(1, 23) =$



(a) Uncorrected



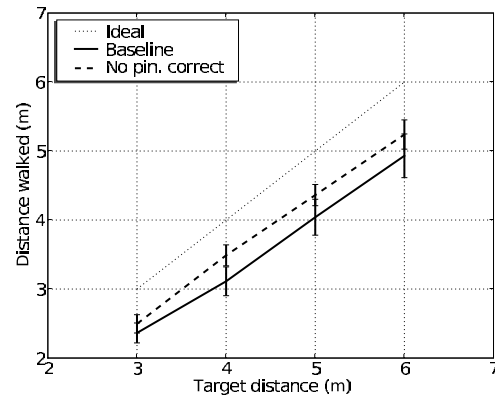
(b) Corrected

**Figure 4:** Simulation of the pincushion distortion seen in our HMD before and after correction.

1.65,  $p = 0.31$ ). The small and statistically non-significant increase in perceived distance in the distorted condition provides weak evidence that pincushion distortion might have the potential to change perceived distance in some circumstances. Future work could test this with an additional condition that introduces more distortion than the display normally has.

## 2.4 Environments with minification

The term *field of view (FOV)* refers to the horizontal and vertical angles subtended by a display or imaging device. A related term, *geometric field of view (gFOV)*, is the field of view used by the computer to render 3D graphics. In a virtual environment system, gFOV must match the FOV of the display to prevent unnecessary minification or magnification of the graphics. For example, if the gFOV is larger than the FOV, the final image will be minified. In a real environment, minifying or magnifying lenses can be used to perform the same manipulations. Previous work shows that minification and magnification can change perceived distance in real environments [Campos et al. 2007], and minification can change perceived distances in HMD-based VEs [Kuhl et al. 2006]. Since the HMD calibrations described in the present work were not performed in the



**Figure 5:** Pincushion distortion correction did not significantly change distance judgments.

previous study [Kuhl et al. 2006], we wanted to replicate the result with a more accurately calibrated baseline condition.

Minification and magnification change several visual cues that contain distance information. Minification changes three specific cues in a way that can potentially increase perceived distance to objects. First, minification reduces the visual angle between objects in a scene and can potentially decrease the angle of declination and increase perceived distance. Minification does not, however, change the angle of declination if it is measured proprioceptively with head rotations. Second, since minification will make objects appear smaller, familiar size cues may also make objects appear more distant. Finally, minification will cause binocular convergence to indicate that objects are also more distant. This cue, however, can only provide accurate distance information for objects within a few meters. Magnification changes these cues in an opposite way and can potentially decrease perceived distance.

We used the same experiment design described in Section 2.1 to compare minification to the baseline condition used in the previous two experiments. Twelve subjects participated in a minification condition which scaled the graphics by a factor of 0.70 relative to the baseline. Figures 2(a) and 2(c) illustrate the difference between the baseline and minified conditions.

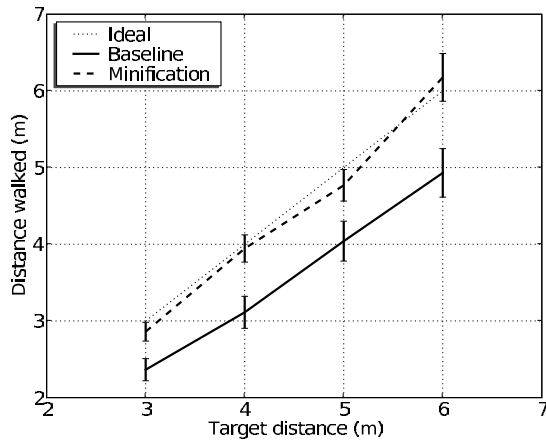
### 2.4.1 Results and discussion

In the minification condition, we found that subjects blind walked significantly further than the subjects in the baseline condition—a result similar to that of Kuhl et al. [2006]. A 2(condition) x 4(distance) ANOVA confirmed that these conditions were significantly different from each other ( $F(1, 23) = 8.06, p < 0.01$ ). This result provides additional evidence that the minification of graphics can cause large changes in perceived distance.

## 3 HMD calibration

### 3.1 Pitch calibration

There has been little information published about how people might calibrate pitch in their non-see-through HMD displays. Measuring



**Figure 6:** *Minification significantly changed distance judgments measured by direct blind walking.*

pitch in an HMD can be difficult since the tracker orientation sensor is typically at some unknown orientation relative to the optical axis of the display. An informal survey we sent to several laboratories indicated that many people attempt to calibrate their displays by having individual users look straight ahead, measure the pitch indicated by the sensor, and then compensate for that amount of pitch for that particular user. There are a couple problems with this approach.

First, little is known about how accurate people are at looking straight ahead in a HMD when asked. We had subjects who participated in our distance judgment experiments (Section 2) also perform a quick task to determine how accurate they were at looking straight ahead. The task was performed after they put on the HMD and before the virtual hallway was displayed in the HMD. Subjects were told to close their eyes, stand up straight, imagine that they were standing in front of a mirror, and point their head straight ahead as if they were looking at their own eyes in the imaginary mirror in front of them. Then, we measured the pitch of the HMD with our tracking system. For the twelve subjects that participated in one of our visual conditions, we found that people pitched their head up an average of 1 degree (relative to level as determined by our pitch calibration procedure discussed below). While average positioning was near accurate, there was a 22 degree range in pitch across subjects. Although we found that a 5.7 degree pitched world does not affect distance judgments (Section 2.2), it is possible that greater amounts of pitch may produce an effect. This experiment demonstrates that relying on people’s ability to look straight ahead for pitch calibration may introduce a significant and unnecessary between-subject variable.

A second problem with calibrating HMD pitch based on the pitch people look when asked to look straight ahead is related to the way the HMD is sitting on a person’s head. If HMD is not fitted level on the subject’s head the virtual world will be pitched by the same amount. For example, consider the situation where the HMD is pitched down by five degrees on a subject’s head but the subject is pointing their head straight ahead during the calibration. In this situation, the virtual camera used to render the graphics should be pitched down since the HMD is pitched down on the subject’s head. If we set level to the angle the HMD sensor is at (pitched five degrees down) the horizon will also be displayed as being pitched down five degrees relative to the real horizon.

We used the following procedure to measure the pitch of our orientation sensor relative to the optical axis of the HMD. First, we measured subject’s eyeheight and placed a horizontal strip of tape on a wall in the lab such that the top of the tape matches the measured eyeheight. In the HMD, we displayed a colored rectangle that corresponded to the area of the real wall below the top of the tape. I.e., the rectangle was the same width as the tape and was at a position in the virtual world that matched the real world. To verify that the placement of the rectangle was correct, subjects compared the virtual rectangle with the tape while standing as close as possible to the tape. Subjects performed this comparison by repeatedly raising and lowering the HMD onto their head. At short distances, the top of the tape and the top of the virtual rectangle should appear to line up almost exactly even if the virtual world is incorrectly pitched by a small amount.

Next, subjects stood approximately 9 meters directly in front of the tape and the virtual rectangle. By standing at a distance from the tape, the effect of any small errors in the placement of the tape or in the tracking system’s height measurements are reduced. Subjects then viewed the virtual rectangle and tape by repeatedly lifting and lowering the HMD onto their head. Subjects were told to keep the HMD as level as possible as they raised and lowered the HMD. It is not critical, however, that the HMD be perfectly level during this procedure as long as there is no minification, magnification, or pincushion distortion in the display. For example, consider the situation where the subject is raising and lowering the HMD pitched up ten degrees. Since the HMD is pitched up according to the tracking system, the virtual camera used to render the scene will also be pitched up and the colored rectangle will appear ten degrees lower in the HMD. If the subject matched the virtual rectangle with the real-world tape with the HMD tilted in this way, it would still provide an accurate measure of pitch. To account for this potential problem, we performed our pitch calibration after we measured the field of view of our HMD (Section 3.3) and corrected for pincushion distortion (Section 3.2).

Each subject performed four matching trials for each eye. For half of the trials, the virtual world was initially pitched up by more than 5 degrees. For the other trials, it was pitched down. Subjects then indicated to the experimenter if the virtual environment had to be pitched up or down. The experimenter followed the instructions of the subject and pressed buttons on the computer to pitch the world up or down in quarter degree increments. Once the subject was satisfied that the top of the tape matched the top of the virtual rectangle, the amount of pitch adjustment was recorded.

We had six lab members perform this calibration procedure. Only lab members who had normal vision without eyeglasses performed the procedure since eyeglasses can introduce a small amount of pitch. We found that the virtual world needed to be pitched up in both eyes for the world to be level. The median value of the averages for each subject were 1.5 and 2.2 degrees for the left and right eyes respectively. A paired t-test showed a statistically significant difference between the left and right eyes ( $t(5) = -3.57, p < 0.02$ ). One-sample t-tests shows that both the left ( $t(5) = 4.60, p < 0.01$ ) and right ( $t(5) = 6.66, p < 0.01$ ) eye results differed significantly from zero. The difference between the left and right eye indicate that the optical axes are not perfectly aligned in our display. By independently pitching the virtual world in the left and right eyes by the amounts we measured, we can compensate for this misalignment.

## 3.2 Pincushion distortion correction

Several methods can be used to correct for pincushion distortion [Tsai 1987; Weng et al. 1992; Bax 2004]. For example, one real-time method uses a textured mesh that is warped to compensate for the distortion [Bax 2004]. We used a fragment program, instead of a textured mesh, to correct for the distortion in real-time. The first step in our method is to render a normal image of the scene to a texture. In our implementation, we render to a 1280x1024 rectangular texture that is the same size as one of the screens in our HMD. We apply this texture with a fragment program to a quad which covers the entire screen. The fragment program is effectively run for every pixel on the screen and corrects for the pincushion distortion using the following steps:

1. The program looks up the coordinate on the screen the fragment program is being run on. Next, we calculate the distance between the current screen coordinate and the center of the screen. We normalize this distance such that the distance from the center of the screen to the corner is 1. We also record which direction the point is from the center of the screen.
2. Using the radial distance from the center of the screen ( $r_s$ ), we calculate the radial distance we should use to lookup into the texture ( $r_t$ ):

$$r_t = r_s + kr_s^3$$

This equation is an approximation of pincushion distortion where  $k$  is a number that represents the amount of pincushion distortion. Next, we reverse the normalization procedure on  $r_t$  and use it, along with the direction recorded in step 1, to calculate the texture coordinate of the pixel we want to display.

3. Next, we look up the color of the pixel in the texture at the texture coordinate and use that color at the current position on the screen. If we look up a pixel that extends past the edge of the texture, we simply display black. This results in a black border on the screen which will frame the image in a way that will look rectangular when the image is displayed through the HMD optics (see Figure 4).

Although the procedure described above does correct for pincushion distortion, it results in part of the screen being unused. Specifically, the whole image will be shrunk such that there are unused black pixels around the perimeter of the screen. Ideally, all of the pixels in the center vertical column and in the center horizontal row should be used. Since many HMDs already have limited fields of view, we want to ensure that our pincushion distortion correction allows us to use as much of the screen as possible. We can solve this problem by calculating a horizontal and vertical scaling needed to make the image fit to the screen. To calculate the horizontal scaling, we follow the steps above for the pixel in the middle of the right edge of the screen to find the corresponding texture coordinate to be used. In this case, the texture coordinate will fall outside the texture. Next, we will determine the scaling factor necessary to make the texture coordinate correspond to the middle of the right edge of the texture. Then, we use that horizontal scaling factor immediately before all of our texture lookups. This procedure is repeated in a similar way to determine the vertical scaling factor for texture lookups. After this scaling, the resulting image will fit the screen with a minimal black border in the corners of the screen as shown in Figure 4.

The extra rendering pass and the fragment program required for pincushion distortion correction introduces extra work that can reduce the rendering performance of a system. The hallway environment we used for the experiments in Section 2 was rendered on a pair of computers (one computer for each eye) with nVidia GeForce


7900 GTX video cards. Without pincushion distortion correction the hallway was rendered at approximately 750 frames per second. With pincushion distortion correction turned on, the hallway was rendered at approximately 350 frames per second. In all of the experiments discussed in this paper, however, the frame rate was synchronized to the 60Hz refresh rate of our HMD. Our pincushion correction implementation rendered both passes of the distortion correction algorithm to the frame buffer with multisampling and anisotropic texture filtering. Other implementations that render directly to textures in the first distortion correction rendering pass would likely improve performance.

To help us measure the pincushion distortion in our HMD, we created a program which allowed us to dynamically correct for varying amounts of pincushion distortion. The amount of pincushion distortion we are correcting for in the program is represented by a variable  $k$ . The program displayed three vertical lines and three horizontal lines. The lines were evenly spaced on the screen. In other words, one of the vertical lines was at the center of the screen and the other two were midway between the center and left and right edges of the screen. Subjects were told to press the buttons on the keyboard until all of the lines appeared as straight as possible. The buttons allowed subjects to increase or decrease  $k$  in 0.01 increments. They were also told that the center horizontal and center vertical lines should always appear to be straight and should be used for reference. When the subjects were satisfied that all of the lines looked straight, we recorded  $k$ . The software initially started at  $k = 0$  or  $k = 0.2$  which caused the image to have significant pincushion or barrel distortion respectively. After performing this procedure with five lab members, we found that the average  $k$  value for all people was 0.098 and 0.102 for the left and right eyes respectively. There was no significant difference in the results for the left and right eyes (paired t-test:  $t(4) = -0.86, p = 0.43$ ). One sample t-tests found that results differed significantly from  $k = 0$  (no pincushion correction) for both the left ( $t(4) = 10.21, p < 0.01$ ) and right ( $t(4) = 8.67, p < 0.01$ ) eyes. Since the optics for the left and right eyes should be the same and there was no significant difference between the left and right eye results, we averaged these two numbers together and used  $k = 0.1$  for both eyes. Figure 4 shows a simulated image of what this amount of distortion looks like in our HMD.

## 3.3 Field of view calibration

Nominal field of view for most HMDs is specified by the manufacturers in visual angle across the diagonal of the screen. Most graphical programs instead require the horizontal and vertical FOVs of the display. Almost all HMD users determine geometric horizontal and vertical fields of view by converting from the nominal diagonal field of view assuming a square-pixel aspect ratio and using calculations equivalent to Equations 1 and 2, where  $d$  is the diagonal FOV in degrees and  $a$  is the aspect ratio. The horizontal and vertical geometric fields of view so determined can fail to match the actual display FOV for three reasons: (1) The nominal diagonal field of view can differ from the actual display diagonal FOV. (2) The actual aspect ratio of the display may not match the ratio of horizontal to vertical pixel counts. (3) Over most of the display area, pincushion distortion leads to field of view values larger than actually occur (See Figure 7). Effective HMD calibration methods must be able to account for all three sources of error.

$$\text{hFOV} = \frac{180}{\pi} \cdot 2 \cdot \arctan \left( a \cdot \frac{\tan(d \cdot \pi/360)}{\sqrt{a^2 + 1}} \right) \quad (1)$$



**Figure 7:** Calculating the horizontal and vertical FOVs from the diagonal FOV in a display that has pincushion distortion produces artificially large values.

$$\text{vFOV} = \frac{180}{\pi} \cdot 2 \cdot \arctan\left(\frac{\tan(d \cdot \pi/360)}{\sqrt{a^2 + 1}}\right) \quad (2)$$

There are several different methods published by others which discuss methods of measuring field of view. Ellis and Nemire [1993] displayed vertical poles in the HMD and had subjects point at the perceived locations of the poles. The actual angle of the virtual pole off of center should have a match the angle the subject pointed at. If not, one can calculate how the gFOV should be changed so that the virtual angles and pointed angles do match. Another method [Rinalducci et al. 1996] uses a camera flash to give people an after-image of a known visual angle. This afterimage can then be used while wearing the HMD to measure the FOV of the display. Finally, it may be possible to calibrate non-see-through HMDs by adapting methods used for see-through display calibration [Gilson et al. in press]. Although none of these methods are particularly complex, we used a simpler method to measure the FOV of our display.

To calibrate the horizontal field of view, we placed two vertical white strips onto a wall in the lab 1m apart from each other. Next, we displayed two vertical strips in the HMD whose centers were separated by two thirds of the display (i.e. 853 pixels between the centers of the strips during the first rendering pass of the pincushion distortion correction process). Subjects stood between and directly in front of the two strips on the wall. Subjects compared the real world strips with the strips displayed in the HMD by repeatedly raising and lowering the HMD on their head. By walking forward and backward, subjects could change the visual angle between the real world strips in order to match the virtual world strips. Because the vertical strips were 2.5cm wide in the real world and roughly 16 pixels wide in the display, subjects were told to line up the centers of the strips. When the subjects were satisfied that the real and virtual strips line up, we measured the distance from their eyes to the wall.

We also performed the same measurement for the vertical field of view. For this measurement, we placed two horizontal strips on the wall 50 cm apart (half the distance used to measure the horizontal field of view) and centered about the our eye height. In the HMD, we also halved the distance between the horizontal bars (i.e. 427 pixels between the centers of the strips). As a result, the distances subjects stand from the wall should be the same for both the horizontal and vertical fields of view.

$$\frac{180}{\pi} \cdot 2 \cdot \arctan\left(\frac{w}{2dp}\right) \quad (3)$$

Several members of our research group performed this task with both eyes open and stood an average of 196.6 cm and 197.0 cm from the wall to match the real and virtual lines for the horizontal and vertical field of view measurements respectively. We can use Equation 3 to calculate the entire field of view of the display in degrees (let  $w$  be the distance between the real strips on the wall,  $d$  be the distance the subjects stood from the strips, and  $p$  be the number of pixels between the centers of the bars divided by the number of pixels across the entire screen prior to pincushion distortion compensation). With this equation, we calculated that the actual horizontal and vertical FOVs of our display were 41.7 and 33.9 degrees respectively. Equation 4 shows how to compute the scaling factor applied to the image when a gFOV of  $g$  degrees (either horizontal or vertical) is used on a device with a FOV of  $f$  degrees.

$$\text{ImageScale} = \frac{\tan(f \cdot \pi/360)}{\tan(g \cdot \pi/360)} \quad (4)$$

Our FOV calibration procedure must occur after pincushion distortion is compensated for. If FOV is measured without pincushion distortion correction, the percent of the screen between the bars as seen by the subject ( $p$ ) cannot be easily determined due to the non-linear pincushion distortion.

## 4 Conclusions

HMD calibration is often overlooked by researchers studying distance perception in virtual environments, and poorly calibrated displays have the potential to change distance perception. Although the present work found no statistically significant effects from pitch and pincushion distortion on distance judgments, different amounts of distortion or different environments could produce different results. The lack of strong effects for the pitch and pincushion distortion manipulations introduces further perceptual questions about the information used to determine absolute distance and the differences between perception in virtual environments versus the real world. The present work also provides additional evidence that researchers using HMDs must be particularly careful when measuring the FOV of their displays. Since trusting published FOV values for displays can be problematic, the calibration procedure presented in the present work is one way researchers can independently measure an HMD's FOV.

## Acknowledgements

This material is based upon work supported by the National Science Foundation under grant numbers IIS-0121084 and IIS-0745131.

## References

- AZUMA, R., AND BISHOP, G. 1994. Improving static and dynamic registration in an optical see-through HMD. In *Proc. ACM SIGGRAPH*, 197–204.
- BAX, M. R. 2004. Real-time lens distortion correction: 3D video graphics cards are good for more than games. *Stanford Electrical Engineering and Computer Science Research Journal* (spring).
- CAMPOS, J., FREITAS, P., TURNER, E., WONG, M., AND SUN, H.-J. 2007. The effect of optical magnification/minimization on

- distance estimation by stationary and walking observers. *Journal of Vision* 7, 9, 1028a. Abstract.
- ELLIS, S. R., AND NEMIRE, K. 1993. A subjective technique for calibration of lines of sight in closed virtual environment viewing systems. In *Proc. Society for Information Display*.
- FRANZ, G. M., VON DER HEYDE, M., AND BÜLTHOFF, H. H. 2004. The influence of the horizon height on spatial perception and experience in VR. In *Proc. 7th Tübingen Perception Conference*, H. H. Bühlhoff, K. R. Gegenfurtner, H. A. Mallot, R. Ulrich, and F. A. Wichmann, Eds., 149.
- GARDNER, P. L., AND MON-WILLIAMS, M. 2001. Vertical gaze angle: absolute height-in-scene information for the programming of prehension. *Experimental Brain Research* 136, 3, 379–385.
- GILSON, S. J., FITZGIBBON, A. W., AND GLENNERSTER, A. in press. Spatial calibration of an optical see-through head mounted display. *Journal of Neuroscience Methods*.
- KNAPP, J. 1999. *The Visual Perception of Egocentric Distance in Virtual Environments*. PhD thesis, University of California at Santa Barbara.
- KUHL, S. A., THOMPSON, W. B., AND CREEM-REGEHR, S. H. 2006. Minification influences spatial judgments in virtual environments. In *Proc. Symposium on Applied Perception in Graphics and Visualization*, 15–19.
- LOOMIS, J. M., AND KNAPP, J. 2003. Visual perception of egocentric distance in real and virtual environments. In *Virtual and Adaptive Environments*, L. J. Hettinger and M. W. Haas, Eds. Erlbaum, Mahwah, NJ, ch. 2, 21–46.
- MCGARRITY, E., AND TUCERYAN, M. 1999. A method for calibrating see-through head-mounted displays for AR. In *Proc. IEEE and ACM International Workshop on Augmented Reality*, 75–84.
- MESSING, R., AND DURGIN, F. 2005. Distance perception and the visual horizon in head-mounted displays. *ACM Transactions on Applied Perception* 2, 3, 234–250.
- OOI, T. L., WU, B., AND HE, Z. J. 2001. Distance determination by the angular declination below the horizon. *Nature* 414 (Nov.), 197–200.
- RINALDUCCI, E. J., MAPES, D., CINQ-MARS, S. G., AND HIGGINS, K. E. 1996. Determining the field of view in HMDs: A psychophysical method. *Presence: Teleoperators and Virtual Environments* 5, 3, 353–356.
- SAHM, C. S., CREEM-REGEHR, S. H., THOMPSON, W. B., AND WILLEMSSEN, P. 2005. Throwing versus walking as indicators of distance perception in real and virtual environments. *ACM Transactions on Applied Perception* 1, 3, 35–45.
- STOPER, A. E. 1999. Height and extent: Two kinds of perception. In *Ecological approaches to cognition: Essays in honor of Ulric Neisser*, E. Winograd, R. Fivush, and W. Hirst, Eds. Erlbaum, Hillsdale, NJ.
- THOMPSON, W. B., WILLEMSSEN, P., GOOCH, A. A., CREEM-REGEHR, S. H., LOOMIS, J. M., AND BEALL, A. C. 2004. Does the quality of the computer graphics matter when judging distances in visually immersive environments? *Presence: Teleoperators and Virtual Environments* 13, 5, 560–571.
- TSAI, R. Y. 1987. A versatile camera calibration technique for high-accuracy 3D machine vision metrology using off-the-shelf TV cameras and lenses. *IEEE Journal of Robots and Automation* 3, 4 (Aug.), 323–344.
- WENG, J., COHEN, P., AND HERNIOU, M. 1992. Camera calibration with distortion models and accuracy evaluation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14, 10 (Oct.), 965–980.
- WILLEMSSEN, P., GOOCH, A. A., THOMPSON, W. B., AND CREEM-REGEHR, S. H. 2008. Effects of stereo viewing conditions on distance perception in virtual environments. *Presence: Teleoperators and Virtual Environments* 17, 1, 91–101.
- WILLEMSSEN, P., COLTON, M. B., CREEM-REGEHR, S. H., AND THOMPSON, W. B. in press. The effects of head-mounted display mechanical properties and field-of-view on distance judgments in virtual environments. *ACM Transactions on Applied Perception*.
- WITMER, B. G., AND KLINE, P. B. 1998. Judging perceived and traversed distance in virtual environments. *Presence: Teleoperators and Virtual Environments* 7, 2 (Apr.), 144–167.