Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE **2006**	2. REPORT TYPE	3. DATES COVERED **00-00-2006 to 00-00-2006**

4. TITLE AND SUBTITLE **Vertical Vergence Calibration for Augmented Reality Displays**	5a. CONTRACT NUMBER
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Virtual Research Lab,Naval Research Laboratory,Washington,DC,20375**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR’S ACRONYM(S)
	11. SPONSOR/MONITOR’S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
IEEE Virtual Reality Conference, Alexandria, VA, 12--15 March 2006.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**2**	

Vertical Vergence Calibration for Augmented Reality Displays

Mark A. Livingston*
Adam Lederer
Virtual Reality Lab
Naval Research Laboratory

Stephen R. Ellis
Human Information Processing Research Branch
NASA Ames Research Center

Sean M. White
Steven K. Feiner
Department of Computer Science
Columbia University

Abstract

Stereo and bi-ocular head-mounted displays (HMDs) require the user to fuse two images into a coherent picture of the three-dimensional world. The human visual system performs this task constantly, but when the input images contain both real and graphical depictions, the problem becomes more difficult. A vertical disparity in the graphics causes diplopia for users trying to fuse the real and virtual objects simultaneously. We implement three methods to measure and correct this disparity and assess them with a collection of a single model of optical see-through HMD.

CR Categories: H.1.2 [Models and Principles]: User/Machine Systems—Human Factors; I.3.7 [Computer Graphics]: Three-dimensional Graphics and Realism—Virtual Reality

Keywords: augmented reality, head-mounted display, vergence

1 Introduction

Many augmented reality (AR) designers consider stereo imagery necessary for users to perceive graphical elements as representations of 3D objects existing in the surrounding 3D environment. However, the human visual system can tolerate only a limited amount of vertical misalignment and still fuse stereo imagery; the compensation a user can coerce in the visual system is much smaller than the amount for horizontal disparity. If the user wants to simultaneously fuse the graphical and real environments, forcing one to fuse may cause diplopia (double vision) for the other. This will inhibit understanding the merged environment.

Proper alignment between the two eyes is a necessary, but not sufficient, condition for the user to perceive correct registration. Without correct alignment, the user may perceive a single eye to be registered, but not both. Even unregistered graphics will not fuse without sufficiently accurate vertical alignment. Furthermore, the effort required to slightly misalign the eyes in order to compensate for improper vertical alignment can lead to eye strain and headache with extended use [6].

Our motivation comes from working with optical see-through AR over long distances (Figure 1). Users had trouble fusing the graphics in our Sony Glasstron due to vertical displacement between the eyes. Thus perceiving both graphical objects and the real environment was difficult, and the problem varied between different units of the same model. Similar problems have been reported due to time delay between left and right eyes encoded in video fields [2] and difficulties in precision alignment of video-based AR eyepieces [4, 9] and our early system [10]. Tests for pilots and military applications indicated tolerances from 0.6 to 5.5 mrad (2.1 to 18.9 arcmin) [5, 8]. Oishi and Tachi [7] used a panel of LEDs and a matching procedure for thirteen points per eye in an iterative six-parameter calibration. We opt for a simpler, more direct approach.

*e-mail: mark.livingston@nrl.navy.mil

Figure 1: This stereo pair (graphics enhanced for grayscale viewing) from the NRL urban situation awareness application exhibits the problem. The left image shows excellent registration, while the right shows noticeable error. The difficulty of putting the camera in the display's exit pupil causes the displacement of the overlay fields and may add registration error; these images are representative, however.

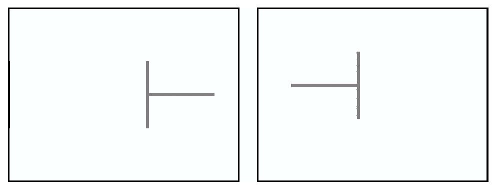

Figure 2: Simulated left and right images show a vertical misalignment; this may be perceived as properly aligned. The vertical bar and half the nonius line are in each image.

2 Measuring and Correcting Vergence Error

We measure vertical disparity with a horizontal *nonius line* (Figure 2), a line broken into two segments; each segment is visible to only one eye. When two vertical segments move laterally under fixed convergence, their relative positions indicate the depth of convergence [3]. Here, we detect a vertical vergence error, provided that the nonius lines themselves are not fused. We do not need registration with the real world, and thus ignore tracking; we only need the user to have a real object on which to focus. (We used a real crosshair that was nearly identical to our graphics, in order to minimize clutter.) If the user experiences diplopia in the graphics when focusing on the real world or vice-versa (Figure 3), we have detected a disparity. We set the IPD for rendering to eliminate horizontal disparity and then seek to correct the vertical disparity.

We may reduce apparent vertical offset between the two eyes in three ways: pitch, translate, or shift the rendered image–i.e., render the image, read it back, and draw it again with a pixel-resolution offset. The first two require a transformation between the two eyes.

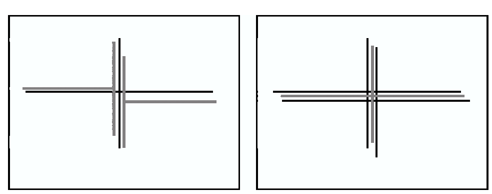

Figure 3: The user should perceive a single real and single virtual crosshair with correct vertical alignment. Otherwise, the user experiences diplopia with the virtual (left) or real (right) crosshair.

The last method is slowest, and its speed varies more with graphics hardware. All corrections are simple to apply with standard graphics library commands.

Corrections may be necessary in only one eye or in both. Since corrections are implemented in a head-referenced coordinate system, we must choose correctly. If we pitch one or both eyes erroneously and then the user tries to roll the view, the roll would induce a pitch error. Since our goal initially was to simply measure the offset, we leave this issue for future work.

For each user, we measured inter-pupillary distance (IPD) (range: 57.0–62.5mm) and screened for stereo vision. All users self-reported normal or corrected-to-normal vision and completed the stereo test without error. We then had the user look through the display and briefly described diplopia and how to recognize proper vergence for both real and virtual objects. As all users had significant experience with 3D graphics, we trusted them to recognize the situation properly. We used a chin rest to reduce head motion.

We first adjusted the IPD for the rendering so that the vertical bars converged. We then set the right image higher than the left and began adjusting the right downward. The user indicated when the right half of the nonius line first became collinear with the left half. This gave us one bound on the range in which the offset allowed for vergence of both the graphics and the real environment. This procedure was repeated for each correction method and then again with the right half of the nonius line beginning below the left half.

3 Results and Discussion

Six users tested three Sony Glasstron LDI-D100B displays. The Glasstron focuses the virtual image at an apparent distance of 1.2m; our real background was at 3.9m. The Glasstron does not enable adjusting the IPD or vergence of the displays. We manually measured the vertical field of view (FOV) of the graphics within the HMDs. Table 1 shows the average correction for all users and both starting configurations. A fourth Glasstron (LDI-D100BE, ID 234) was tested by a single user (IPD 62.5mm).

These corrections should represent a consistent visual angle. We convert translation and image shift to mrad using the Glasstron specifications of a 30-inch-wide virtual image and 600 pixels vertical resolution, modulated by the measured vertical FOV. Table 1 gives the angles for translation (α_T) and shift (α_S). To understand the lack of consistency, we measure the range for each type of correction, averaged over all users. Table 2 shows the range in which fusion occurred for each correction method, along with the value that equates to normal visual acuity of one arcminute. The range swamps visual acuity, indicating the tolerance of the visual system and its ability to make the images match a preconceived notion.

Two users perceived a roll offset for HMD 029. We apply a roll correction in one eye and with the same procedure as for the other offsets, measured four HMDs with a single user. The user was satisfied in the range of 0.9–4.7 mrad (3.1–16.1 arcmin) correction for HMD 029. Two HMDs were perceived to have no roll offset at -8.7 mrad (-30 arcmin) of roll when increasing from an extreme negative initial roll and 4.4 mrad (15.1 arcmin) when reducing from an extreme positive initial roll. This implies that no correction is absolutely required. The range for the fourth HMD was -13.9 to 0.3 mrad (-47.8 to 1.0 arcmin).

Testing a second collection of Glasstrons for vertical disparity yielded one more device that showed significant error (Table 3) and more units that appear to require a roll correction. Testing on a third collection of Glasstrons identified one (out of three) that required vertical correction; that measured -5 pixels by image shift [9].

ID	Pitch	Trans	Shift	FOV	α_T	α_S
029	6.9	27.4	19.0	19.5°	12.2	10.8
030	0.6	2.2	1.3	19.1°	1.0	0.7
060	1.9	7.1	5.1	19.2°	3.1	2.8
234	4.9	0.2	12.5	20.0°	0.1	7.3

Table 1: Summary of correction results: rotations (Pitch, α_T, α_S) are in mrad; translation is in mm; shift is in pixels.

ID	Pitch		Translation	Shift
029	0.9 mrad	(3.1′)	2.5 mm	2.7 pix
030	0.8 mrad	(2.8′)	3.0 mm	2.0 pix
060	1.0 mrad	(3.4′)	4.8 mm	4.2 pix
Acuity	0.3 mrad	(1.0′)	0.7 mm	0.5 pix

Table 2: Range of the measured correction across users. (Only one user tested HMD 234.) The range for the shift may have been enlarged by the slow update rate when this method was applied.

ID	Pitch		Translation	Shift
070	0.2 mrad	(0.6′)	1.0 mm	0 pix
231	0.3 mrad	(1.2′)	0.0 mm	-2 pix
233	-0.5 mrad	(-1.8′)	-2.0 mm	-2 pix
235	6.6 mrad	(22.8′)	23.0 mm	16 pix

Table 3: Test results for a second set of Glasstrons. The measurements for HMD 235 are not consistent.

The immediate implication for AR system designers is to test display devices for this disparity and correct it if needed. The user will find it much easier to converge the images. One explanation for the variability in our measurements is the tolerance of the human visual system. Further testing will determine whether the user's needs are met for particular applications. We must determine how the measured correction changes with the focal distance of the user. Comparing rotation and translation corrections, users observed a slight ghosting effect when applying a translation correction measured at a small focal distance to a situation in which the user maintains a large focal distance, but the effect is small and transient. No such issue was observed with the rotation correction, but there is insufficient evidence to claim it does not exist. We arrived at a simple solution for the problem originally identified; despite not testing where to apply corrections, users have an easier time fusing imagery and perceive improved registration with these corrections. More complete testing of the effects of these methods is planned. Still, our current implementation assists users in fusing stereo graphics images viewed simultaneously with the real world.

References

[1] James E. Cutting. *Reconceiving Perceptual Space*, pages 215–238. MIT Press, 2003.

[2] David Drascic and Paul Milgram. Perceptual issues in augmented reality. In Mark T. Bolas, Scott S. Fisher, and John O. Merritt, editors, *SPIE Volume 2653: Stereoscopic Displays and Virtual Reality Systems III*, pages 123–134, January/February 1996.

[3] Stephen R. Ellis and Brian M. Menges. Localization of virtual objects in the near visual field. *Human Factors*, 40(3):415–431, 1998.

[4] Henry Fuchs, Mark A. Livingston, Ramesh Raskar, D'nardo Colucci, Kurtis Keller, Andrei State, Jessica R. Crawford, Paul Rademacher, Samuel H. Drake, and Anthony A. Meyer MD. Augmented reality visualization for laparoscopic surgery. In *Proceedings of First International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, October 1998.

[5] James E. Melzer and Kirk Moffitt, editors. *Head Mounted Displays*. McGraw-Hill, 1997.

[6] Kirk Moffitt. *Designing HMDs for Viewing Comfort*. McGraw-Hill, 1997.

[7] Takashi Oishi and Susumu Tachi. Methods to calibrate projection transformation parameters for see-through head-mounted displays. *Presence: Teleoperators and Virtual Environments*, 5(1):122–135, 1996.

[8] Clarence E. Rash, editor. *Helmet-Mounted Displays: Design Issues for Rotary-Wing Aircraft*. 2001. SPIE Press Monograph Vol. PM93.

[9] Andrei State. Personal communication, August 2005.

[10] Blaine Bell. Personal communication, August 2001.