

Achieving near-correct focus cues in a 3-D display using multiple image planes

Simon J. Watt^{*}^a, Kurt Akeley^b, Ahna R. Girshick^c, Martin S. Banks^c

^aSchool of Psychology, University of Wales Bangor, LL57 2AS, United Kingdom;

^bMicrosoft Research: Asia, Beijing Sigma Center, No. 49, Zhichun Road, Beijing, China 100080;

^cSchool of Optometry, 360 Minor Hall, University of California, Berkeley, CA USA 94720

ABSTRACT

Focus cues specify inappropriate 3-D scene parameters in conventional displays because the light comes from a single surface, independent of the depth relations in the portrayed scene. This can lead to distortions in perceived depth, as well as discomfort and fatigue due to the differing demands on accommodation and vergence. Here we examine the efficacy of a stereo-display prototype designed to minimize these problems by using multiple image planes to present near-correct focus cues. Each eye's view is the sum of several images presented at different focal distances. Image intensities are assigned based on the dioptric distance of each image plane from the portrayed object, determined along visual lines. The stimulus to accommodation is more consistent with the portrayed depth than with conventional displays, but it still differs from the stimulus in equivalent real scenes. Compared to a normal, fixed-distance display, observers showed improved stereoscopic performance in different psychophysical tasks including speed of fusing stereoscopic images, precision of depth discrimination, and accuracy of perceived depth estimates. The multiple image-planes approach provides a practical solution for some shortcomings of conventional displays.

Keywords: Stereoscopic displays, binocular vision, virtual reality, human-computer interaction

1. INTRODUCTION

The goal of virtual reality is to create simulated scenes that are indistinguishable from the real world. Modern computer graphics has made great progress towards this. Yet in conventional three-dimensional (3-D) displays the focal distance to points in the simulated scene is inconsistent with the portrayed scene because the light comes from a single, roughly planar display surface. This causes two distinct problems. First, differing demands are placed on the accommodation and vergence systems. Vergence and accommodation responses are normally coupled, so accommodating to one distance while converging to another (“decoupling” vergence and accommodation) can lead to discomfort, fatigue and difficulty fusing the pair of stereoscopic images^{1,2,3}. Second, focus cues to depth—from accommodation and the gradient of retinal blur—specify the geometry of the display surface and not the portrayed scene. There is now considerable evidence that the brain combines information from multiple sources (so-called “depth cues” such as binocular disparity, perspective, relative motion, etc.) to compute 3-D scene parameters⁴. Inappropriate information from focus cues can therefore result in distortions in perceived depth when viewing virtual scenes^{5,6}.

We have constructed a working prototype of a novel 3-D display. The display minimizes these problems by presenting each eye's image as a depth-weighted blend of multiple images at different focal distances. Here we report the outcome of two psychophysical experiments examining the efficacy of this approach.

1.1. Design of the display prototype

The design of the display prototype is shown schematically in Figure 1. It has been described in detail elsewhere⁷, and so only a brief outline is presented here. In each eye's view a combination of a first-surface mirror and two plate beamsplitters is used to create an image that is the sum of aligned images drawn at three discrete focal planes (see also McDowall and Bolas⁸). This creates a truly volumetric stereoscopic display, albeit with low resolution in focal depth (but see 1.2. *Depth-weighted filtering of image intensities*, below) because the light comes from sources at near-correct focal distances. The two eye's images are superimposed optically using the “periscope” arrangement shown in the lower part of Figure 1a, resulting in an overlapping, binocular visual field (Figure 1b).

*s.watt@bangor.ac.uk

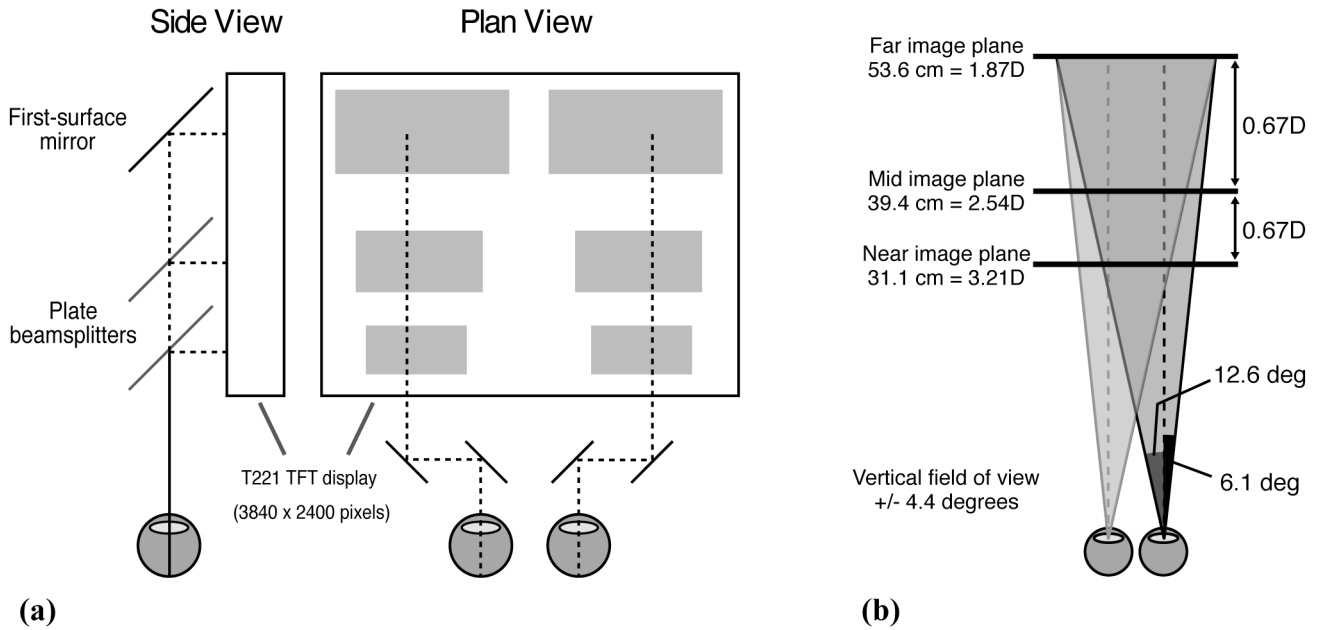


Figure 1. Design of the prototype display. Left panel: The side view has been rotated by 90 degrees clockwise and shows the arrangement of mirror and beamsplitters used to create each eye's view. These are positioned in a box and the T221 display is placed "face down" on top. The plan view shows the periscope arrangement used to superimpose the two eye's views, and shows schematically where the three viewports for each eye are located on the display. Using a single display greatly simplifies alignment of the viewports (both within and between each eye's view). Pixel locations are fixed precisely by the manufacturer, and the display is driven by a digital interface, so the position of illuminated pixels can be controlled precisely. The T221 is a high resolution Liquid Crystal display manufactured by IBM⁹. Even at the near image plane, individual pixels were not visible. Right panel: The physical distance of the three image planes from each eye. The distance between the image planes is equal in diopters (see text). This panel also shows the size of the visual field in each eye's view, as well as the extent of the overlapping (binocular) visual field. In this prototype the observer's head is positioned and restrained using a custom-fitted bite bar.

Volumetric displays are desirable because correct (or in the case of our display, near-correct) focus cues can be presented without tracking an observer's eye movements. Another critical feature of the design in Figure 1 is that each eye's viewpoint is held fixed relative to the display. The advantages of fixed-viewpoint volumetric displays over other approaches are discussed in detail by Akeley and colleagues⁷. These include: (i) presenting geometrically correct binocular disparities at the two eyes (our device takes account of individual variations in inter-ocular distance in software), (ii) preserving viewpoint-specific lighting effects of current graphics technology such as shading, specularity, reflection, and occlusion, and (iii) the capability to present multiple focal distances along a single visual line. Moreover, because the human visual system's sensitivity to focus cues is far poorer than for spatial position, a fixed-viewpoint volumetric display can be built with a comparatively low resolution in depth¹⁰. The amount of blur due to defocus is approximately proportional to the magnitude of the defocus, measured in diopters¹¹ (the reciprocal of meters). The image planes should therefore be spaced equally in dioptric distance. This means that at increasing depths fewer and fewer image planes are required, which permits considerable reductions in both hardware complexity and in the computational expense of rendering multiple images. (Note that the use of a low resolution in depth only affects focus cues, and not the precision with which other depth cues such as binocular disparity and perspective can be presented.)

1.2. Depth-weighted filtering of image intensities

For all but the extremely unlikely case in which the depth of a point in the scene coincides exactly with the depth of one of the image planes, a rule is required to assign image intensities to each depth plane. The simplest rule would be to draw each point in the scene at the depth plane to which it is nearest, measured in diopters. Continuous surfaces in depth, however, would be drawn partly on one image plane and partly on another, with a sharp boundary between regions drawn at nearer and farther planes. It is not possible to focus at two planes at once, and it can be shown that the presence

of any focus error results in visible discontinuities in the retinal image⁷. Therefore, this technique would be unsuitable for high quality images. Instead we implemented depth-based filtering of image intensities, where image intensity at each depth plane is weighted according to the distance (in diopters) of the point from that plane, determined *along lines of sight*. This is shown schematically in Figure 2a. Figure 2b shows the weights that would be given to points at each distance in the scene in a display with image planes at 1.25, 0.75, and 0.25D. As well as removing visible discontinuities due to focus error, this technique also minimizes discontinuities resulting from incorrect alignment of pixel edges at different image planes (because all three image planes are drawn on one fixed-resolution display, the angular size of the pixels is not the same at all focal distances).

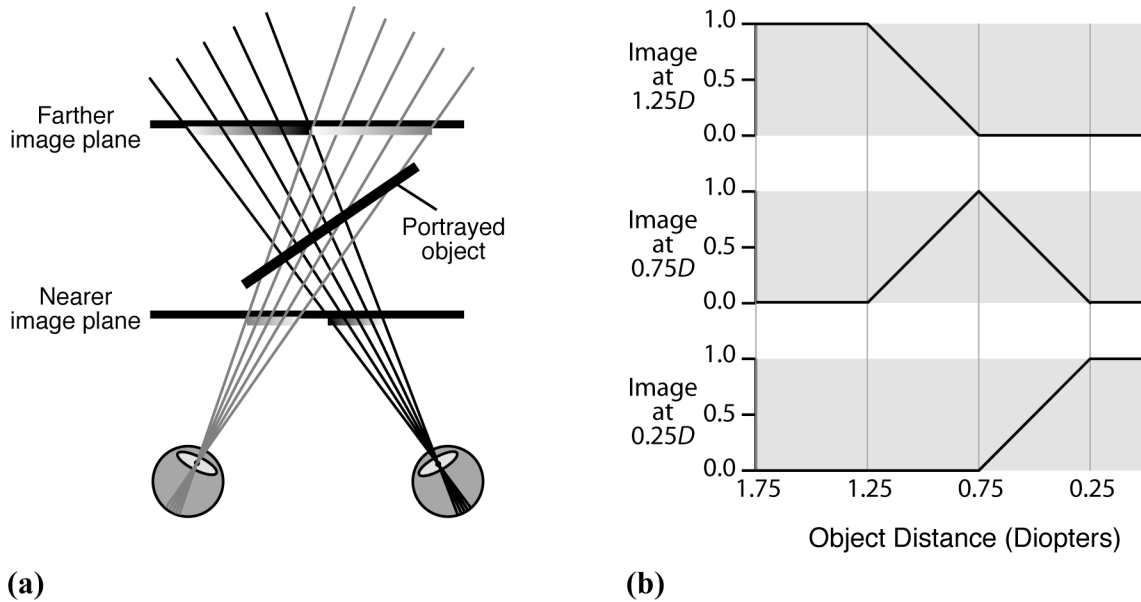


Figure 2. Intensities at each image plane are assigned by depth-weighted filtering along lines of sight. a) The principle of how this is implemented for the example of a slanted plane at a depth between two image planes. The graded bars at each image plane represent schematically the image intensity drawn at each position (the left and right eye’s images are color coded in grey and black, respectively, and do not represent different intensities). b) An example of filter weight functions for image points at different distances, drawn across three image planes at 1.25, 0.75 and 0.25D. The filter weight functions are linear with respect to distance measured in diopters.

1.3. Evaluating the prototype display

We report three experiments examining the efficacy of the prototype display device. In the first experiment we examined the effects of varying amounts of vergence-accommodation mismatch on observers’ ability to fuse close-to-threshold stereoscopic stimuli (sine-wave corrugations in depth) in order to confirm that stereoscopic fusion is affected by the need to decouple accommodation and vergence responses. In this experiment all stimuli were presented at (or very near to) each depth plane and so depth filtering was not used.

Depth filtered images differ from those in equivalent real-world scenes. To determine if the remaining differences between our display and real scenes hinders performance significantly, we conducted a second experiment. In this experiment we examined whether depth filtering led to impaired performance on the stereoscopic fusion task from Experiment 1.

Our goal is not only to produce a 3-D display with improved ergonomics, but also to present compelling 3-D percepts that result in undistorted percepts of depth. In the third experiment we examined whether depth filtering affected observers’ percept of slant from disparity and perspective.

2. EXPERIMENT 1

This experiment examined the effects of varying the mismatch between the stimulus to accommodation and to vergence on observers' ability to fuse stereoscopic images. Observers were presented with random-dot stereograms depicting sinusoidal corrugations in depth and their task was to identify the orientation of the ridges (plus or minus 15 deg from horizontal). This task cannot be completed based on the monocular images alone, but requires that the two images are fused into a percept of stereoscopic depth. Fusion performance was determined by measuring the time required to identify the stimulus orientation correctly. We assumed that observers should require more time to identify the stimulus correctly under conditions in which fusion was difficult. To ensure that the test was sensitive to changes in stereoscopic fusion performance, the dot density of the random-dot pattern and the spatial frequency of the depth modulation were chosen such that the stimulus was close to threshold for a typical observer¹². In this range, blur in the retinal image (which would arise from incorrect accommodation) has been shown to cause a reduction in stereoacuity¹².

2.1. Methods

2.1.1. Observers

Three observers participated in the experiment, aged 29 (ARG), 23 (DS) and 19 (BGS) years. All three had normal stereoscopic vision, and wore their normal correction for refractive error where appropriate.

2.1.2. Stimuli

The stimuli were random-dot stereograms depicting corrugations in depth. The dot density was ~ 45.3 dots/deg², and the spatial frequency of the depth modulation was ~ 1.35 cycles/degree. The depth varied sinusoidally in a direction either plus or minus 15 deg from vertical (therefore the corrugations in the stimulus were oriented either plus or minus 15 deg from horizontal), with a peak-to-trough disparity of ~ 4.9 arcmin at all viewing distances. The stimulus was clipped by a circular aperture with a diameter of ~ 4.2 deg. The stimulus was presented at various combinations of vergence-specified distance and focal distance (see 2.1.3. Procedure).

2.1.3. Procedure

The experiment used a forced-choice paradigm. On each trial observers were briefly presented (366msec) with a "cues-consistent" fixation cross for which the focal distance and vergence distance were both set to the mid plane. This was designed to ensure that the starting state of vergence and accommodation was the same on each trial. The stimulus was then presented in one of six vergence/focal distance configurations. The vergence-specified distance was always either at the near (31.1 cm) or far (53.6 cm) plane and the focal distance was either the near, mid (39.4 cm) or far image plane. Therefore both near and far vergence-specified stimuli were presented with varying degrees of vergence-accommodation mismatch. A mask was presented after the stimulus, consisting of dots randomly positioned in depth. On a given trial the image intensity of the mask was assigned to the different image planes using the same rule as for the experimental stimulus. Observers then indicated the orientation of the corrugation with a key press. The stimulus presentation time in each viewing condition was varied by a separate staircase procedure, using a 1-up/2-down reversal rule. Therefore each time an incorrect answer was given, the stimulus time increased for that condition on the next trial. Two correct answers in succession resulted in a shorter stimulus duration in that condition on the next trial. All viewing conditions were completed concurrently and a staircase was chosen at random on each trial. Over separate trial blocks, observers completed two repetitions of each staircase. A psychometric function (cumulative Gaussian) was fitted to all the data for each observer in each condition using a maximum likelihood criterion¹³. The 75% correct point on each function was defined as the time required to fuse the stimulus in each condition.

2.2. Results and discussion

The results for each observer in Experiment 1 are plotted in Figure 3. Results for the near and far vergence distances are shown separately. Figure 3 clearly shows that, for a given vergence distance, the time required to fuse the stimulus increased as function of the magnitude of the mismatch between vergence-specified distance and focal distance. The only exception to this pattern is the far vergence distance for observer DS. Nonetheless, this observer took significantly longer to fuse the stimulus in the condition with the largest conflict (indicated by the dark bar) than in the other two conditions. For all three observers the overall time taken to fuse the stimulus was longer when it was presented at the far vergence compared to the near vergence distance. The magnitude of vergence eye movement required in each case is

approximately equal ($\pm \sim 2.5$ deg). It seems likely therefore that this reflects the previously reported finding that convergent eye movements are faster than divergent eye movements¹⁴.

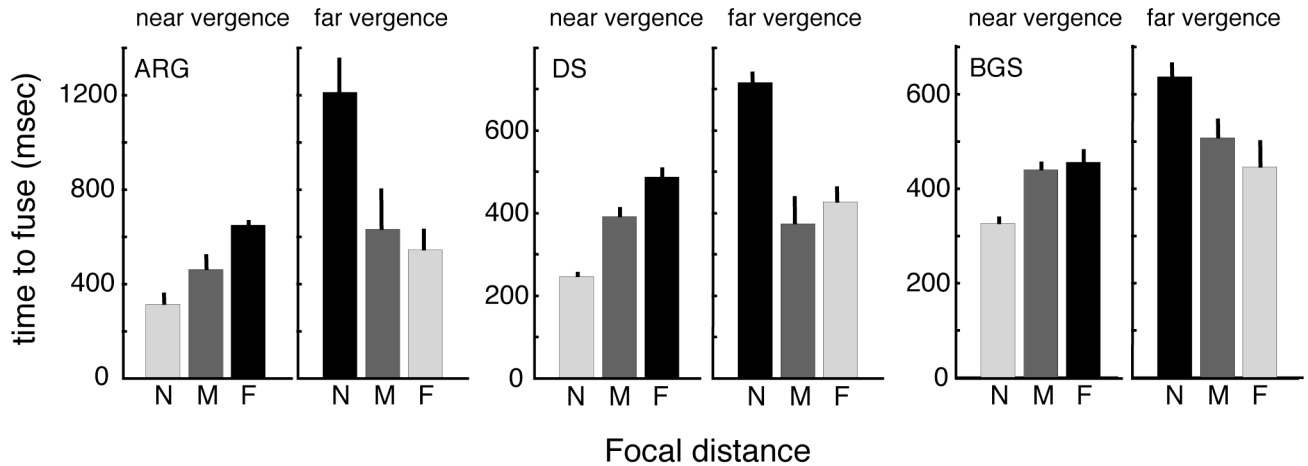


Figure 3. Time taken to fuse the stimulus (75% correct identification of grating orientation) for each observer in each viewing condition. The near and far vergence-specified distances are plotted separately for each observer, as indicated in the legends above each plot. The labels N, M and F indicate the image plane at which the stimulus was physically drawn (the focal distance), and refer to the near, mid and far planes respectively. In each case, the magnitude of the mismatch between the stimulus to vergence and to accommodation is denoted by the shading of the bars. The pale bars indicate no mismatch, the mid-grey bars indicate a moderate mismatch, and the black bars indicate the largest mismatch. For clarity, observer ARGs data are plotted on a different y-axis scale. Error bars denote 97.7% confidence interval of the 75% correct threshold estimate.

These data suggest that the presence of inconsistent stimuli to vergence and accommodation results in a significant increase in the difficulty of fusing stereoscopic images. This is consistent with previous reports that vergence-accommodation mismatch causes discomfort and fatigue^{1, 2, 3} and demonstrates that, compared to conventional displays in which focal distance is held constant, a clear advantage in stereoscopic performance can be gained by presenting correct focus cues.

3. EXPERIMENT 2

A key feature of the fixed-viewpoint, multiple depth-planes approach outlined above is that it presents near-correct focus cues using a relatively coarse resolution in focal depth. We have argued that this is only viable if depth filtering is employed to remove visible discontinuities in the retinal images⁷. Depth-filtered images differ from those resulting from geometrically equivalent real-world scenes¹⁵. To evaluate our display it is therefore necessary to determine whether depth filtering disrupts stereoscopic performance. This was the purpose of Experiment 2.

3.1. Methods

The same observers participated as in Experiment 1. The random-dot stimulus and general experimental procedure were also identical to Experiment 1. However, the stimulus to vergence and to accommodation differed from those in Experiment 1 in several ways. First, the vergence-specified distance was either (i) between the near and mid image planes, referred to here as “near-mid” at 34.8cm, or (ii) between the mid and far image planes, referred to here as “mid-far” at 39.4cm. At each vergence-specified distance, the stimulus was drawn either entirely at one of the two image planes nearest to the vergence-specified distance (at the near or mid plane in the case of the near-mid vergence distance, and at the mid or far planes for the mid-far vergence distance) or as a depth-filtered blend of image intensities at the two nearest image planes. The near-mid and mid-far distances were at the mid points between the two nearest planes, measured in diopters. This meant that for the depth-filtered stimuli, $\sim 50\%$ of the image intensity was assigned to each of the two image planes on which it was drawn.

3.2. Results and discussion

The results of Experiment 2 are plotted in Figure 4. It can be seen that for all three observers, and at both vergence-specified distances, there was no significant increase in the time taken to fuse the stimulus for depth-filtered stimuli (paler bars) compared to the stimuli drawn at one image plane (darker bars). This suggests that stereoscopic fusion was not impaired appreciably by depth filtering. Given the results of Experiment 1, which showed that stereoscopic fusion was faster when there was no mismatch between the stimulus to accommodation and to vergence, one might expect the depth-filtered stimuli to result in *improved* performance relative to the stimuli drawn at only one image plane (there was a small mismatch between vergence-specified and focal distance for the “unfiltered” stimuli). However, it is not possible to focus the eye on two different depth planes at once, so the depth-filtered stimulus must inevitably be somewhat blurred at the retina. It is possible that for our stimulus, which was close to threshold, this blurring due to defocus resulted in somewhat reduced stereoacuity which lead to slower fusion times¹².




Figure 4. Time taken to fuse the stimulus (75% correct identification of grating orientation) for each observer in each viewing condition. The near-mid and mid-far vergence-specified distances are plotted separately for each observer, as shown in the figure legends above each plot. The labels N, M and F indicate near, mid, and far planes, respectively. NM and MF refer to the near-mid and mid-far distances. In the case of focal distance, stimuli labeled NM and MF were drawn as a depth-filtered blend of image intensities at the near and mid, or mid and far planes, respectively. These depth-filtered stimuli are indicated by the paler bars; the darker bars indicate stimuli drawn entirely at one image plane. Error bars denote 97.7% confidence interval of the 75% correct threshold estimate.

4. EXPERIMENT 3

Experiments 1 and 2 suggest that our prototype display can reduce some of the problems associated with a mismatch between demands on accommodation and vergence. Our goal, however, is to achieve this while maintaining high image quality and creating compelling and accurate 3-D percepts. Experiment 3 examined whether depth-filtered stimuli, presented on a small number of discrete image planes, result in appropriate depth percepts. To do this we compared the perceived slant of simulated planar surfaces drawn entirely at one depth plane, and depth-filtered planes in which image intensity was distributed across all three image planes.

4.1. Methods

4.1.1. Observers

There were three observers, aged 18, 19, and 24 years. All had normal stereoscopic vision, and wore their normal correction for refractive error where appropriate.

4.1.2. Stimuli

The stimuli were virtual planes slanted about a vertical axis (i.e., tilt = 0 deg). Observers completed two viewing conditions: (i) a “binocular” condition in which the slanted plane was defined by a sparse random-dot stereogram, and (ii) a “monocular” condition in which the stimulus was defined by a Voronoi pattern¹⁶, and viewed monocularly. These patterns have been shown previously to be an effective cue to slant¹⁷. The stimulus dimensions were the same in both conditions. The height of the slanted plane at its centre was 6.8 deg. The width of the stimulus was not matched in terms of angular size at different slants (although a small random jitter was added to the size on each trial to prevent different slants from being discriminated by simply attending to one edge of the stimulus. The stimulus widths were 16.4, 13.2 and 9.1 deg at base slants of 30, 45, and 60 deg, respectively.

4.1.3. Procedure

The experiment employed a two-interval, forced-choice procedure. On each trial the observers saw a standard stimulus and a comparison stimulus (randomly ordered) and their task was to indicate which was more slanted. The standard stimulus was depth-filtered, with image intensity distributed across all three image planes (near, mid and far). The comparison stimulus was drawn only at the mid image plane. In each condition (disparity and texture) the standard was presented at base slants of ± 30 , 45 and 60 deg. The amount of image intensity at the near and far image planes varied as a function of the absolute magnitude of base slant because the edges of the slanted surface were closer to these image planes with increasing slant. Also, because the slanted plane always fell well within the space between the two image planes, information from focus cues actually specified parts of the surface as in front of and beyond the depth specified by geometric information from binocular disparity and texture.

The slant of the comparison stimulus was varied by both 1-down/2-up and 2-down/1-up staircases. For each base slant in each viewing condition, both reversal rules were used to sample points on the psychometric function either side of the 50% point. Staircases for each condition were randomly interleaved in a block of trials. A psychometric function (cumulative Gaussian) was fitted to the data in each condition using a maximum-likelihood criterion¹³, in order to determine the point of subjective equality (PSE) for each observer for each base slant/viewing condition.

4.2. Results and discussion

The results of Experiment 3 are plotted in Figure 5. The left panel (Figure 5a) shows the PSEs for the binocular condition, averaged across all three observers. The right panel (Figure 5b) shows the average data for the monocular condition. These PSEs correspond to the slant angle at which the non-filtered comparison, which was drawn all at one image plane, appeared equally slanted as the depth-filtered comparison. In the comparison stimulus, depth information from focus cues (accommodation and the gradient of retinal blur) was consistent with a frontoparallel plane. In the standard stimulus, however, image intensity was distributed across the near, mid and far image planes. It might be expected, therefore, that greater slant would be required in the comparison stimulus to look as slanted as the standard. Figure 5 shows, however, that this was not the case. The data all lie on the diagonal line in both viewing conditions, indicating that equal amounts of slant were seen in the depth-filtered stimuli and the stimuli drawn at only one image plane.

These results suggest that depth-filtered stimuli can create depth percepts comparable to conventional 3-D displays, in which all image intensity is at one planar surface, and that slant percepts were not adversely affected by depth filtering. It was not the case that the depth-filtered and non-filtered stimuli always appeared identical, however. Observers reported informally that in the monocular conditions the depth-filtered surface often did not appear planar. Instead the near portion of the surface appeared to curve towards the frontoparallel plane. This suggests that inappropriate information from focus cues was influencing slant percepts in this region of the image. This is consistent with previous reports that depth information from blur contributes to percepts of slant for monocularly viewed textured planes⁵, and that different weights may be assigned to depth cues in different regions of the same image¹⁸. Although we could not quantify the effect in this experiment it is consistent with the visual system giving increased weight to depth information from the gradient of retinal blur at near distances. This might be expected because the depth signal from blur should be more informative at near distances, where a given depth change results in a greater change in retinal blur.

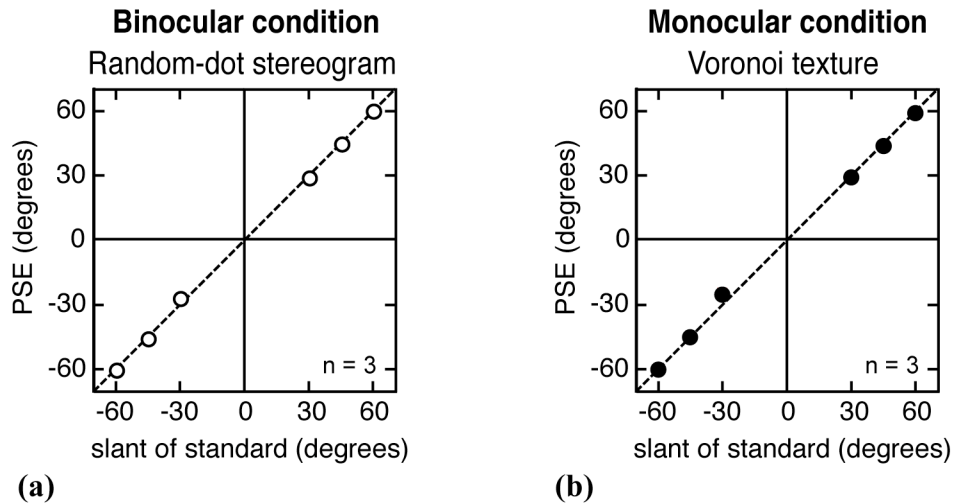


Figure 5. PSEs as a function of base slant, for the binocular (a) and monocular (b) viewing conditions. The data are averaged across all three observers. The dotted diagonal line in each plot shows where the data would lie if the depth-filtered standard and non-filtered (drawn all at one plane) stimulus appeared equally slanted. Between subject error bars (± 1 SEM) are plotted in each case but are smaller than the data points.

5. CONCLUSIONS

We examined stereoscopic fusion performance and depth perception in a novel fixed-viewpoint volumetric display that uses multiple image planes and depth-weighted filtering of image intensities to provide near-correct focus cues. Results from Experiments 1 and 2 suggest that this approach can improve stereoscopic fusion performance compared to conventional displays by minimizing the difference between demands on the vergence and accommodation systems. Significantly, no noticeable decrement in performance was observed when observers had to fuse depth-filtered stereoscopic images in which the image intensity was distributed across two image planes at different focal distances.

The results of Experiment 3 suggest that stereoscopic slant perception was not adversely affected by drawing the stimulus as a depth-weighted blend of image intensities at different focal distances. Qualitative reports suggested that for monocularly viewed texture-defined images, depth filtering resulted in distortions in perceived depth at near focal distances. However, this prototype device was designed to test the principle of multiple depth planes and depth-weighted filtering of image intensities, and as such it was built using just three widely-spaced depth planes. More practical displays could feasibly be built that use a greater number of more closely spaced depth planes, collapsed into a small enough physical size to produce a head mounted display¹⁰. Nonetheless, the results of our preliminary experiments suggest that the multiple image-planes approach provides a practical solution for some shortcomings of conventional displays.

ACKNOWLEDGEMENTS

This work was supported by NIH grant EY14194-01 (MSB).

REFERENCES

1. Judge, S. J., & Miles, F. A., Changes in the coupling between accommodation and vergence eye movements induced in human subjects by altering the effective interocular distance. *Perception*, 14, 617-629, 1985.
2. Wann, J. P., Rushton, S., & Mon-Williams, M., Natural problems for stereoscopic depth perception in virtual environments. *Vision Research*, 35, 2731-2736, 1995.
3. Wöpking, M., Viewing comfort with stereoscopic pictures: An experimental study on the subjective effects of Disparity magnitude and depth of focus. *Journal of the SID*, 101-103, 1995.
4. Landy, M. S., Maloney, L. T., Johnston, E. B., & Young, M. J., Measurement and modeling of depth cue combination: In defense of weak fusion. *Vision Research*, 35, 389-412, 1995.
5. Watt, S. J., Banks, M. S., Ernst, M. O., & Zumer, J. M., Screen cues to flatness do affect 3d percepts. *Journal of Vision*, 2(7), 297a, 2002.
6. Watt, S. J., Akeley, K. & Banks, M. S., Focus cues to display distance affect perceived depth from disparity. *Journal of Vision*, 3(9), 66a, 2003.
7. Akeley, K., Watt, S. J., Girshick, A. R., & Banks, M. S., A stereo display prototype with multiple focal distances. *ACM Transactions on Graphics*, 23, 1804-813, 2004.
8. McDowall, I. & Bolas, M., Fakespace labs accommodation display research. Unpublished report, 1994.
9. Wright, S. L., IBM 9.2-megapixel flat-panel display: Technology and infrastructure. *SPIE Proceedings 4712 (Apr.)*, 24-34, 2002.
10. Rolland, J. P., Krueger, M. W., & Goon, A. A., Dynamic focusing in head-mounted displays. *SPIE*, 3639, 463-470, 1999.
11. Smith, G., Jacobs, R. J. & Chan, C. D., Effect of defocus on visual acuity as measured by source and observer methods. *Optometry & Vision Science*, 66, 430-435, 1989.
12. Banks, M. S., Gepshtein, S. & Landy, M. S., Why is stereoresolution so low? *Journal of Neuroscience*, 24, 2077-2089, 2004.
13. Wichmann, F. A., & Hill, N. J., The psychometric function: I. Fitting, sampling and goodness-of-fit. *Perception & Psychophysics*, 63, 1293-1313, 2001.
14. Hung, G. K., Zhu, H. M. & Ciuffreda, K. J., Convergence and divergence exhibit different response characteristics to symmetric stimuli. *Vision Research*, 37, 1197-1205, 1997.
15. Akeley, K., Achieving near-correct focus cues using multiple image planes. PhD thesis, Stanford University, 2004.
16. de Berg, M., van Kreveld, M., Overmars, M., & Schwarzkopf, O., *Computational geometry: Algorithms and applications* (2nd Ed.). New York: Springer-Verlag, 2000.
17. Knill, D. C., Discrimination of planar surface slant from texture: Human and ideal observers compared. *Vision Research*, 38, 1683-1711, 1998.
18. Hillis, J. M., Watt, S. J., Landy, M. S., & Banks, M. S., Slant from texture and disparity cues: optimal cue combination. *Journal of Vision*, 4, 967-992, 2004.