Psychophysical Exploration of Stereoscopic Pseudo-Transparency

Mai Otsuki
Ritsumeikan University, Japan.
otsuki@rm.is.ritsumei.ac.jp

Paul Milgram
Dept. Mechanical & Industrial Engineering,
University of Toronto, Canada.
milgram@mie.utoronto.ca

IEEE International Symposium on Mixed and Augmented Reality 2013
Science and Technology Proceedings
1 - 4 October 2013, Adelaide, SA, Australia
978-1-4799-2868-2/13/$31.00 ©2013 IEEE

Abstract
We report a two part experiment related to perceiving (virtual) objects in the vicinity of (real) surfaces when using stereoscopic augmented reality displays. In particular, our goal was to explore the effect of various visual surface features on both perception of object location and perception of surface transparency. Surface features were manipulated using random dot patterns on a simulated real object surface, by manipulating dot size, dot density, and whether or not objects placed behind the surface were partially occluded by the surface.

Author Keywords
Human factors, stereoscopic augmented reality, pseudo-transparency, transparency perception

ACM Classification Keywords
H.5.1. Information interfaces and presentation (e.g., HCI): Artificial, augmented, and virtual realities.

General Terms
Human Factors, Design, Experimentation

Intro: Stereoscopic Pseudo-Transparency
One of the challenges facing the practical application of augmented reality (AR) in domains such as endoscopic surgery [1-3, 7, 8] is how to cause a virtual (computer generated) image to appear behind a real object

Figure 1. Experimental setup

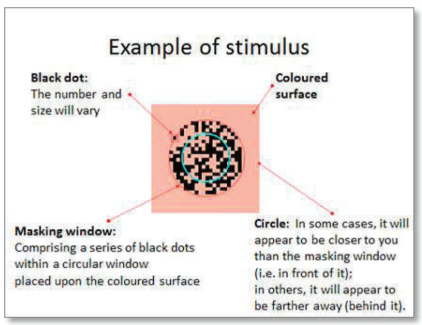

Figure 2. Typical stimulus

surface. When using video based stereoscopic displays (that is, either see-through or monitor based video), a conflict occurs between the binocular disparity depth cue, which tells the observer (correctly) that the virtual object is farther away than the real object surface, and the occlusion depth cue, which tells the observer that the virtual object surface must be closer than the intervening real surface. The net effect of overcoming this conflict and successfully creating the desired impression occurs when the intervening (real object) surface appears to be semi-transparent, with the effect often referred to as dynamic transparency [5], or pseudo-transparency[1][10].

Our first goal, carried out in Part 1 of the experiment, was to explore whether there is a perceptual bias in either direction when placing a virtual cursor near a surface, as well as to estimate what degree of sensitivity is to be expected. Our second goal, carried out in Part 2, was both to explore the extent to which the various simulated surface textures were amenable to being perceived as transparent and to create a quantitative scale for assigning perceived transparency values to various surface features.

Experimental Setup

Image Generation

All stimuli were generated on a desktop computer (Windows 7 Professional OS, Intel Core i5 2310 2.8 GHz CPU, 8G RAM, with NVIDIA Quadro 600), coded using Visual C++ 2010 and were presented on a 23-inch LCD screen (ASUS VG236HE, 1920 x 1080 resolution, 120-Hz refresh rate) with a black background. Stereo images were observed using a nVidia 3D vision system, with 3D Vision 2 glasses. The lighting in the room was kept at a level of between 350 and 500 lux. The experimental setup is shown in Fig. 1.

Fig. 2 shows the image shown to participants as part of their training for the experiment. The size of the entire stimulus area shown was 200 x 490 pixels. Although the paragraphs above refer to real surfaces and virtual objects, as if we are dealing with an augmented reality display, it is important to point out that the "real surface" that we used in our experiment was in fact a *simulated real surface*. We did this for the sake of expediency, by using a monochrome plane – the 'coloured surface' in the figure. For our virtual object, we used a blue circle, whose position in the depth direction could be produced at various distances in front of (closer to the participant) or behind the coloured surface. For the textured surface, we used a pattern of random dots, for which we were able to vary the size of the dots and the density of dots. We surmised that participants would be comfortable with the term 'masking window' to describe the entire circular pattern of dots and surface elements. The size of the masking window was 140 x 140 pixels.

To investigate a range of surface textures (the second goal of our research) we varied the random dot patterns in terms of the size and density of the dots. Fig. 3 shows all stimuli used in both parts of the experiment. *Dot size* refers to the fraction into which each dimension is divided. For example, 1/20 means that a 20x20 grid was used to generate the random dot pattern. *Density* refers the percentage of the entire masking window that consisted of dots. Note that the dot size and the density are independent of each other.

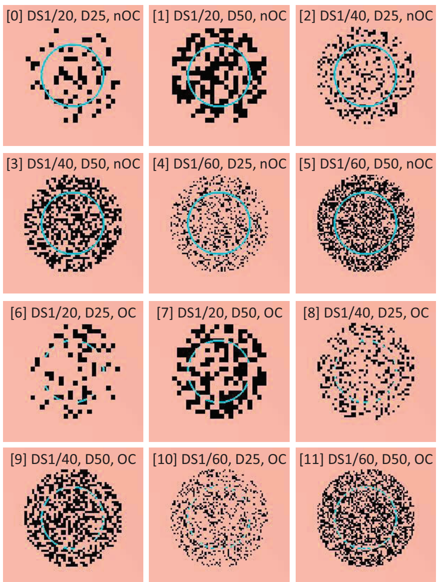

Figure 3. Stimuli for experiment. Only 8 stimuli: #2-5 and 8-11, were used in Part 1. All Stimuli shown here were used in Part 2.
DS=Dot size, D=Density, OC=With Occlusion; nOC = No Occlusion

All images were rendered stereoscopically, with the coloured surface held at a constant distance corresponding to zero disparity. The virtual circle was produced to be either in front of the surface, or behind the surface. For the sake of generality, all of the longitudinal distances used in this report are in program units, where one unit = 116.5 mm. In other words, a displacement of 0.01 units corresponds to 1.165 mm. (Note that we did not constrain participants to maintain a constant viewing distance from the screen.)

Two different methods were used to render the interactions between the circle and the coloured surface. In one condition, illustrated in the top two rows of Fig. 3, *No Occlusion (nOC)*, we essentially ignored any depth relationship between the circle and the surface, such that the circle pixels would occlude all elements of the coloured surface, regardless of whether it was drawn in front of or behind the surface. This corresponds to the condition we hypothesised should be conducive to perception of *stereo-transparency* [1, 10, 11]. (It also corresponds to the case in which one does not have a model of the real surface, which would be necessary in order to implement occlusion in practice.)

For the opposing condition, shown in the bottom two rows of Fig. 3, the *Occlusion (OC) condition*, we treated the black dots and the remaining portions of the coloured surface differently. Whenever the circle was supposed to appear behind the coloured surface, the circle pixels continued to occlude the black dots, as in the nOC case. The coloured surface pixels, on the other hand, *occluded* the circle pixels. This was hypothesised to correspond to the case of *pseudo-transparency*, in the sense of simulating the case of light passing through gaps in non-transparent ('lacy') objects [10] [11]. Whenever the (virtual) circle appeared in front of the surface, there was no difference in how it was rendered for both conditions.

Participants
We recruited 15 University of Toronto students over the age of 21 for the study (18 male, 3 female). All claimed to have normal or corrected to normal visual acuity, as well as no problems with stereoscopic vision. To confirm the latter, the nVidia 3D stereo vision test was administered. All participants were compensated $15 for the privilege of participating in the experiment (both parts), which lasted approximately one hour in total.

EXPERIMENT: PART 1

Purpose and Procedure
Part 1 examined the relationship between the psychophysical threshold of perception of depth sensitivity, using the *method of constant stimuli* [6]. Our aim was to investigate both accuracy, in terms of determining whether there was a perceptual bias in longitudinally localising the (virtual) circle within the vicinity of the coloured surface, and precision, in terms of the sensitivity of perceived location of the circle.

Following a brief training session to familiarise themselves with the interface and procedure, participants were shown a series of stimuli, to each of which they responded whether, according to their judgement, the circle was in front of or behind the masking window. The virtual circle was presented at four distances relative to the masking window: two in front {+0.01, +0.02} and two behind {-0.01, -0.02}. The windows comprised two levels of dot size {1/40, 1/60} and two levels of dot density {25%, 50%}.

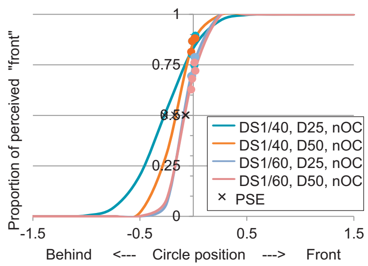

(a) No Occlusion

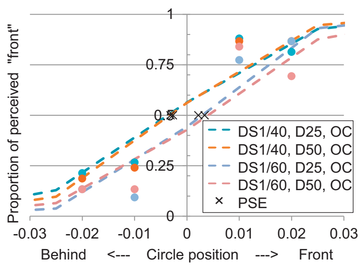

(b) With Occlusion

Figure 4. Psychometric functions for all stimuli. Vertical axis: proportion of times that circle was judged to be closer to front. Horizontal axis: actual position of circle relative to surface.

(NB: *Note different scales: top graph has 50 times width of bottom graph.)*

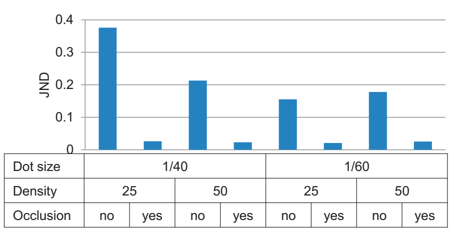

Figure 5. JND Values for each condition

These particular values were selected following extensive pilot testing. Together with the two occlusion conditions {OC, nOC}, and with 5 trials for each combination of conditions, this led to 160 trials (4x2x2x2x5) for each participant. The order of stimuli was randomised Unlimited time was given for responding to each stimulus, with 500 ms between stimuli. In Part 1 we used stimuli 2-5 and 8-11 in Fig.3.

Results and discussion

Fig. 4 shows results from Part 1, averaged over participants, in terms of proportion of times that the circle was perceived as being in front of the coloured surface, as a function of where the circle really was (relative to the participant). Two important patterns can immediately be observed from the figure.

(1) All graphs conform well to the cumulative normal probability function models fit to the data. Note, however, that the horizontal axis in the top graph is *50 times as wide* as that of the bottom graph. In other words, the *sensitivity* for perceiving the circle relative to the coloured surface is *markedly less for the No Occlusion case* (top graph) than it is for the Occlusion case (bottom graph). The just noticeable difference (JND) values (estimated as half of the difference between the 25% and 75% points of the psychometric function) for all of the No Occlusion cases are much larger than those for the With Occlusion cases (Fig. 5).

(2) There is clearly a large *perceptual bias* for the top No Occlusion results, in contrast to the relatively unbiased results at the bottom of Fig. 4. This is illustrated by the X symbols, (labelled PSE, for Point of Subjective Equality) that show where each of the fitted psychophysical functions cross the 0.5 level. In Fig. 4(b), for the With Occlusion conditions, the PSE values indicate that there is very little, if any, bias (corresponding to the centre of the horizontal axis) in participants' ability to perceive the circle in the vicinity of the coloured surface. For the nOC conditions of Fig. 4(a), on the other hand, the PSEs are *much farther behind* (recalling the different scales) the centre of the graph, meaning that for surface textures comparable to those tested here, one should expect to encounter a bias towards perceiving the cursor as being *behind* the surface, under the belief that it is *on* the surface.

The results for the four surface texture conditions tested are less dramatic in comparison with the occlusion / no occlusion conditions. Other than a slight indication that the smallest dot size (1/40) resulted in a somewhat larger JND, Fig. 5 suggests that any potential influence of either dot size or dot density was much less than whether or not occlusion of the circle by the coloured surface was present.

EXPERIMENT: PART 2

Purpose and Procedure

Part 2 examined the relative effectiveness of the different kinds of masking windows for creating the illusion of pseudo-transparency. In this part, the virtual blue circle was always at the same distance *behind* the coloured surface + masking window, at a constant distance of -0.1 units for all conditions. This distance was chosen on the basis of pilot tests.

Thurstone's paired comparison scaling method [9] was used, deemed to be the most efficient way of gauging participant's relative experience of a percept that was not otherwise (easily) objectively quantifiable. The approach was to present participants with pairs of

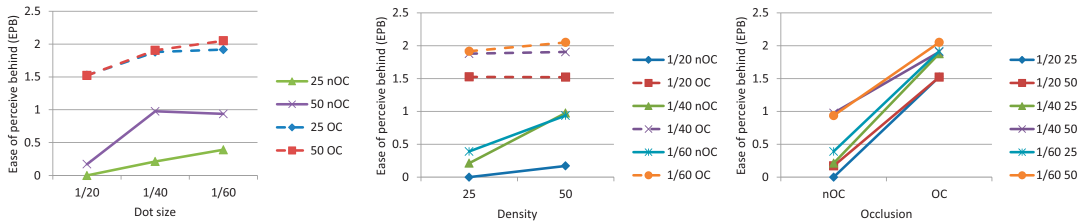

(a) Ease of perceiving behind (EPB) scale

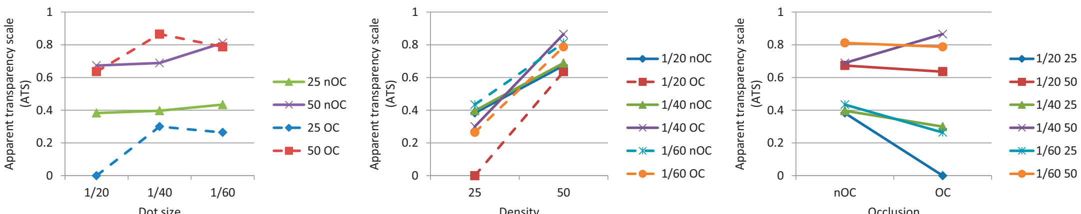

(b) Apparent transparency (ATS) scale

Figure 6. Result of Part 2.

stimuli, for each of which they were required to answer the following two questions. (1) "In which image is it easier to perceive that the circle is behind the masking window?" and (2) "In which image does the masking window appear to be more transparent?"

The experiment comprised three dot sizes {1/20, 1/40, 1/60}, two dot densities {25%, 50%} and 2 occlusion types {OC, nOC}. The resulting 12 (3x2x2) conditions therefore necessitated 66 ($_{12}C_2$) comparisons. All stimuli shown in Fig. 3 were used.

Results and discussion

Fig. 6 shows the scale value results that were calculated using Thurston's method, for (a) Ease of Perceiving Behind (EPB) on top and (b) Apparent Transparency Scale (ATS) on the bottom, for dot size, dot density and occlusion conditions respectively. In each graph, the vertical axis shows the rating scale values, where larger values signify more agreement among participants, in units of standard normal deviates, about corresponding parameters being rated.

The top right EPB graphs indicate clearly that it was easier to perceive that the circle was behind the surface in the With Occlusion condition than in the No Occlusion condition. For the bottom graphs, on the other hand, any influence of the occlusion condition was much less evident. In contrast, the most striking effect evident from the ATS scale values is the influence of Dot Density, where we see that larger dot densities greatly increase the tendency to perceive the surface as being transparent. These results suggest that perceiving the circle behind the surface is not the same process as perceiving the masking window as transparent.

CONCLUSIONS

Part 1, a psychophysical investigation of perceived (virtual) object depth displacement relative to an intervening surface, revealed relatively low accuracy and low precision when no surface occlusion was provided. From a practical point of view, it would be easy to conclude that, as a means of improving virtual object placement performance with monitor based stereoscopic augmented reality displays, one should simply allow some of the surface elements to occlude any virtual object that goes behind the surface. The challenge with such a solution, however, would be to maintain a real-time depth map for detecting and comparing surface elements to virtual object pixels.

Part 2 involved the use of paired comparisons for rating both the ease of perceiving an object as being behind an intervening surface and the relative strength of perceived transparency of the intervening surface.

Results supported the expected importance of the occlusion cue for perceiving object location relative to a surface, as well as the importance of dot density for perceiving surface transparency. The results of this experiment have implications for the ongoing challenge of facilitating the perception of surface transparency in augmented reality.

LIMITATIONS AND FUTURE WORK

One limitation of our experiment already mentioned here is the distance we used to place the circle behind the surface in Part 2 of the experiment. Because we did not have the results of Part 1 when we planned the experiment, we did not anticipate PSE biases along the order of 0.3 to 0.4 units for the No Occlusion case. Our decision to place the circle at a distance of 0.1 units could therefore have confounded our EPB scores for the No Occlusion conditions in Part 2, since it clearly would have been more difficult (i.e. less easy) to perceive the circle as being behind the surface for a distance that, according to Part 1, corresponded to a distance that was subjectively in front of the surface. Our future experiments will certainly take this factor into account.

Another arguable limitation of our experiment was our decision to apply the surface texture to only an identifiable window within the simulated object surface, rather than to the entire surface. Although we have no predictions about what difference this might have made, this factor also merits further consideration.

Related to the above is our basic premise that, for the sake of expediency, our initial investigations of monitor based stereo augmented reality surface texture effects could be carried out using a simulated real object surface. It will definitely be necessary to confirm our findings using actual real object surfaces with real stereo video sensors.

References

[1] R.A. Akerstrom *et al*. The perception of stereoscopic transparency. *Perception & Psychophysics*, 44 (5), pp. 421-432, 1988.

[2] M. Bajura *et al*. Merging virtual objects with the real world: Seeing ultrasound imagery within the patient. *Computer Graphics,* 26(2), pp. 203-210, 1992.

[3] C. Bichlmeier *et al*. Contextual anatomic mimesis hybrid in-situ visualization method for improving multi-sensory depth perception in medical augmented reality. *In Proc. ISMAR '07*, pp. 1-10, 2007.

[4] P.J. Edwards *et al*. Clinical experience and perception in stereo augmented reality surgical navigation. In *Proc. MIAR 2004*, pp. 369-376, 2004.

[5] N. Elmqvist *et al*. Employing dynamic transparency for 3D occlusion management: Design issues and evaluation. In *Proc. INTERACT 2007*, pp. 532-545, 2007.

[6] G. Gescheider, Psychophysics: The Fundamentals. 3rd Edition. Erlbaum, 1997.

[7] S. Nicolau *et al*. Augmented reality in laparoscopic surgical oncology. *Surgical Oncology*, 20 (3), pp. 189-201, 2011.

[8] T. Sielhorst *et al*. Depth perception – A major issue in medical AR: Evaluation study by twenty surgeons. In *Proc. MICCAI 2006*, pp. 364-372, 2006.

[9] L.L. Thurstone. The method of paired comparisons for social values. *J. Abnormal & Social Psychology,* Vol. 21, pp. 384 - 400, 1927.

[10] I. Tsirlin *et al*. Perceptual artifacts in random-dot stereograms. *Perception,* 39, pp. 349-355, 2010.

[11] I. Tsirlin *et al*. Stereoscopic transparency: Constraints on the perception of multiple surfaces. *J. Vision* 8(5):5, pp. 1-10, 2008.