THE UNIVERSITY OF

WARWICK

Selective Rendering for Efficient Ray Traced Stereoscopic Images

Cheng-Hung Lo · Chih-Hsing Chu · Kurt Debattista · Alan Chalmers

Abstract Depth-related visual effects are a key feature of many virtual environments. In stereo-based systems, the depth effect can be produced by delivering frames of disparate image pairs, while in monocular environments, the viewer has to extract this depth information from a single image by examining details such as perspective and shadows. This paper investigates via a number of psychophysical experiments, whether we can reduce computational effort and still achieve perceptually high quality rendering for stereo imagery. We examined selectively rendering the image pairs by exploiting the fusing capability and depth perception underlying human stereo vision. In ray tracing based global illumination systems, a higher image resolution introduces more computation to the rendering process since many more rays need to be traced. We first investigated whether we could utilise the human binocular fusing ability and significantly reduce the resolution of one of the image pairs and yet retain a high perceptual quality under stereo viewing condition. Secondly, we evaluated subjects' performance on a specific visual task that required accurate depth perception. We found that subjects required far fewer rendered depth cues in the stereo viewing environment to perform the task well. Avoiding rendering these detailed cues saved significant computational time. In fact it was possible to achieve a better task performance in the stereo viewing condition at a combined rendering time for the image pairs less than that required for the single monocular image. The outcome of this study suggests that we can produce more efficient stereo images for depth-related visual tasks by selective rendering and exploiting inherent features of human stereo vision.

Cheng-Hung Lo · Chih-Hsing Chu
Department of Industrial Engineering and Engineering Management
National Tsing Hua University, Hsinchu, Taiwan, 30113
Tel.: +886-3-5715131
E-mail: loch@mx.nthu.edu.tw

Kurt Debattista · Alan Chalmers
Warwick Digital Lab, University of Warwick, Coventry CV4 7AL, UK

1 Introduction

The horizontal separation of our two eyes introduces disparities to the images projected onto the two retinas. The human visual system is able to decode the disparity to get a unique sense of depth - Stereoscopic Vision. Many virtual reality applications make use of this visual phenomenon. Stereoscopic-viewing environments not only provide a qualitative feature that gives the users a better sense of immersion, but also a quantitative ability that allows the users to have binocular depth information available to them when performing depth-related visual tasks. However, in order to deliver the two disparate views at an interactive rate, these stereo systems use simple shading models for computer-generated scenes. Physically-based global illumination rendering techniques, which are notorious for their time-consuming calculations, are thus are thus precluded as an option in these systems. In this paper, we attempt to address the issue by introducing the possibility to reduce the computational effort for rendering realistic stereoscopic imagery by incorporating perceptually-based rendering approaches. We conducted a series of psychophysics experiments on the images rendered with ray tracing based algorithms. Our findings from these experiments shed some light on bringing advanced rendering techniques closer to the stereoscopic viewing environments.

2 Background and Related Work

2.1 Ray traced stereoscopic imagery

To generate ray traced stereoscopic images, two distinct views of the scene needs to be created: a left-eye view and a right-eye view. The most straight-forward way to do this is to render both views, doubling the required work for a single perspective image [1]. In addition, for a larger view, or a projection that covers a larger visual field, the desire for a higher resolution requires more pixels to be drawn onto the screen and thus more rays to be traced. Furthermore, stereoscopic imagery is rendered for providing depth information of the scene layout and the spatial relationship between scene objects. Without stereoscopic depth cues, we have to interpret other cues such as occlusions, perspective, and shadows, from a single image to comprehend the spatial structure of the scene. To produce shadows in a ray tracing based renderer, an additional set of shadow rays needs to be traced from the ray-surface intersection points to the light sources. Such shadow rays can lead to great computational expense. If the virtual scene is lit by multiple or complex forms of light sources, then we need to compute and test the intersections of the shadow ray for the many light source. As a result of this, shadow computations are quickly going to predominate since the major cost at each hit point is the cost of intersection testing [2,3].

In this paper, we first investigated whether we could exploit the human binocular fusing ability and significantly reduce the resolution of one of the image pairs and yet retain a high perceptual quality in the stereo viewing condition. In our second experiment, we designed a depth-related visual task that required the viewer to interpret the depth cues provided in the images. Without affecting the subject's performances, we investigated whether we could avoid rendering shadows if the information about scene depth can be obtained stereoscopically, while reducing overall rendering time.

2.2 Psychophysical investigations

Advanced global illumination algorithms, which aim to simulate the behaviour of light in the real world, are nonetheless computationally expensive. In the recent years, a new trend in the research of realistic graphics has emerged by taking human visual perception into account. Many researchers are convinced that extensive previous work done in visual perception for identifying perceptual criteria for perceiving realism can be used to help develop more efficient and accurate perceptually guided rendering algorithms, for example [4–6]. The perceptual criteria identify features of the human visual system and thus provide a means for determining whether a detailed feature in an image is visible, or salient, to the eye. We can then exploit these features to avoid unnecessary computations initiated by the physically-based rendering approaches. Several studies have attempted to establish the framework for balancing the physical computations with the evaluations of perceptual responses such as image comparison [7,8] and visual attention factors [9]. Part of the work in this paper is to further extend such a framework by considering whether a particular piece of visual information contained in the synthetic image is essential when the viewer is performing a depth related visual task.

2.3 Related work

Numerous studies have been done to evaluate the task performance of depth perception when viewing renderings of three-dimensional virtual scenes. Kjelldahl and Prime [10] studied a different set of depth cues and found illumination and object placement were the most effective cues for conveying relative depth information. In Wanger et al. [11], six different depth cues were compared for their effectiveness on three visual tasks. They found that for each task a different combination of depth cues was important. However, both of these studies did not compare their depth cues to stereopsis. Servos et al. [12] examined stereoscopic depth cues and found stereo vision to be significantly important for distance judgments for a grasping task. Hubona et al. [13] compared the effectiveness of stereo, cast shadows, and background images for distance judgments for 3D inferencing. However, these studies examined depth cues for tasks without considering the context of the visual realism and rendering effort. Although the findings show that the set of depth cues that are essential for each task varies depending on the type of task involved. It is still unclear if any of these previous findings will hold when considering the visual realism of a virtual 3D scene.

Wanger [14] studied how the rendering quality of shadows affect the perception of spatial relationships in computer generated imagery. The results from his experiments indicated that rendering less physically accurate hard edged shadows could help the viewer make more accurate judgement of an object's shape as compared to the performance with soft edged shadows. Hu et al. [15] studied how shadows and other visual cues affected subjects' ability to discriminate properties such as object orientation or proximity. Horvitz and Lengyel [16] measured subjects' response on various settings of image quality, to optimize the computation in order to develop an efficient renderer. However, none of these perceptually-based studies has taken binocular visual factors into account.

Perkins [17] introduced a mixed resolution coding scheme for compressing stereoscopic images. Other researchers adapted this idea and developed techniques for transmitting stereoscopic video sequences [18,19]. The findings in these studies indicated

that the perceptual quality of the stereoscopic imagery can be maintained when the resolution of one view is degraded. However, these image degradations were often implemented at the post-processing stage instead of being used as a means to reduce computation at the rendering stage. Some studies used the image re-projection technique for rapidly generating ray traced stereoscopic images [20]. The drawbacks of the re-projection techniques is the difficulty of accurately rendering view-dependent components such as specular highlights.

In this paper, we produced different rendering qualities, which were reflected in the total time to render different level of resolutions of the stereoscopic images, to see whether a difference in quality is perceivable. We also aim to show that, since stereopsis is a strong cue for depth, it is possible to take advantage of this and compute images for depth related visual tasks more efficiently.

3 Experiments

We investigated via a number of psychophysical experiments, whether we could significantly reduce the computational effort of realistic rendering by selectively rendering images, exploiting the fusing capability and depth perception underlying human stereo vision. Our first experiment investigated whether the viewer still demands as high a resolution setting under stereo viewing conditions compared with the monocular condition. We used *eDimensional*'s [21] shutter glasses and the line-interlaced technique for displaying the visual stimuli. The images were rendered at different resolutions for the pair of stereo images. We investigated whether one image of the pair rendered with higher resolution could compensate for the other one rendered with a lower setting. The second experiment focused on the rendering effort required for the monocular depth cues. We evaluated subjects' performance on a specific visual task that required accurate depth perception. Two depth-judgment tests were carried out. In this experiment, we assumed that the subject could obtain depth information without looking for the details of these monocular cues when binocular cues are present in the scene. The images used for the visual stimulus were created using the ray tracing based, advanced lighting simulation system, Radiance [22].

3.1 Resolution experiment: Design

As shown in figure 1, the image for the resolution test consisted of a simple scene lit by a single light source. The image was rendered at ten different levels of resolution. Starting from the maximum setting of 800×800 pixels, we gradually degraded the resolution by a 10% interval to produce subsequent images, e.g. 720×720, 640×640, 560×560 ...until 80×80.

For generating the visual stimuli to be viewed stereoscopically, we rendered two views, separated by a small horizontal difference, of the scene for each level of resolution. The left and right-view images were then processed in the software, *3DCombine* [23], to obtain the scanline-interlaced stereoscopic images. Two sets of stereoscopic images, with different combinations of the level of resolution were created. We combined the left and right views with same resolution setting and same level of degradations for one set, whilst for the other set, we produced the stereoscopic pairs with the right view being always at the highest resolution, 800×800, and the left view being gradually

(a) Standard resolution 800×800

(b) One test resolution 80×80

(c) Interlaced stereoscopic image, 160×160, both view degraded.

(d) Interlaced stereoscopic image, 160×160 + 800×800, one view degraded.

Fig. 1 Scene rendered with different levels of resolutions

degraded. It should be noted here that instead of reducing the image resolution at the post-processing stage as discussed in the related work, the degradation was done by rendering the scene in lower resolution, which means that less rays were traced and less rendering time was spent.

A total of 45 subjects participated in the experiment. They were divided into three groups for three experimental conditions:

– **Monocular viewing condition:** One group was under the condition of viewing without the stereo glasses. The images presented to this group of participants were single-view with the ten levels of resolution.
– **Stereoscopic viewing condition 1:** The second group was presented with the stereoscopic images combined with the same level of resolution, e.g. 80×80(left) and 80×80 (right), 160×160(left) and 160×160 (right), and 200×200(left) and 200×200 (right).
– **Stereoscopic viewing condition 2:** The third group viewed the stereoscopic images combined with different levels of resolution, e.g. 80×80 (left) and 800×800

(right), 160×160 (left) and 800×800 (right), and 200×200 (left) and 800×800 (right).

The participants in the second and third groups were asked to wear the shutter glasses when carrying out the experiment. All of the participants have normal or corrected-to-normal vision. The experiment was carried out with a standard Pentium IV PC running *Presentation* [24], a Windows-based program for stimulus delivery and experiment control. The visual stimuli were displayed on a 19-inch CRT monitor with refresh rate set to 120 MHz to reduce the flicker effects caused by the shutter glasses as much as possible.

For all three groups, the experimental procedure remained the same. The visual stimuli were presented in pairs in random order, showing the standard image of the highest resolution, 800×800, and another whose resolution was varied according to the experimental conditions. After each pair presentation of the visual stimulus, the subject responded with pre-programmed mouse key to indicate whether he/she perceived differences in the resolutions between the two images.

3.2 Resolution experiment: Results

The results for the first experiment are illustrated in figure 2. The horizontal axis in the bar chart indicates the different levels of resolution settings. The vertical axis denotes the detectability of the changes of resolutions, which were normalized by dividing it by the number of subjects in each group. The bars are clustered with three viewing groups for easy comparisons. The black bars represent the subjects' performances under normal viewing condition. The other two bars represent the subjects' performances when wearing shutter glasses. The grey bars indicate the detectability of the resolution changes with both views being degraded. The hatched bar shows the detectability of the resolution when only one view of the stereoscopic pair is degraded.

The results show that under monocular viewing conditions without wearing the shutter glasses, about 2/3 of the subjects were able to tell the difference between the resolutions of 800×800 and 720×720. Almost all of the subjects in this group were able to detect the changes when the resolution was reduced to 480×480 and below. On the other hand, in the stereoscopic viewing group when both left and right views were degraded in resolution, the perceivable differences are more likely to appear when the resolution changes from 800×800 to 480×480 or to 400×400. Moreover, when the resolution for one view was held consistent and the other is degraded gradually, the difference is expanded further to a range of 800×800 to 320×320 or 240×240 to be detectable. The results show that when generating a stereoscopic image, it is not essential to compute left and right view images with the same level of resolution quality. The human visual system is simply able to fuse both views, within a certain range of difference in resolution quality, into one perceptually-equivalent image.

Fig. 2 Results of the resolution experiment

3.3 Depth judgment experiment: Design

Two tests were carried out to measure the subject's performance on two similar visual tasks that require accurate depth perception. The virtual scene for the first test consisted of a large room, a chess piece (the knight) floating in the air, and a square attached to the ground, Figure 3(a)-(d). A series of images with the knight's position being varied in depth were rendered. The images were also rendered with the addition of four different kinds of depth cues: Stereo, Shadow, Relative size, and Texture gradient:

– **Stereo:** Through the shutter-glasses displaying the ray-traced two views simultaneously, the participant obtained stereoscopic depth information to judge the floating piece's position.
– **Shadow:** The shadow cue was produced by the global illumination algorithm in Radiance. The shadow was physically based and thus in accordance to the position of the light sources. To provide a fair comparison to the other depth cues, we placed the light sources at the center above the chess piece.
– **Relative size:** The relative size cue here was depicted by adding two extra pieces to the settings. The participants were expected to compare the size of the target piece with the two extra ones to obtain depth information.
– **Texture gradient:** The final depth cue considered was texture gradient. It was simply implemented by adding textures to the ground and back-wall of the room.

(a) Interlaced stereoscopic image

(b) Image rendered with shadows

(c) Two extra knights added

(d) Texture-mapped background

Fig. 3 Scene conditions for the adjusting test

The Radiance renderer was specifically modified to avoid tracing shadow rays for generating images for the three categories other than the one with shadow as the depth cue. For the first test, we recruited 75 participants. They were divided into five groups, 15 people in each, for the above four viewing conditions and one control group viewing
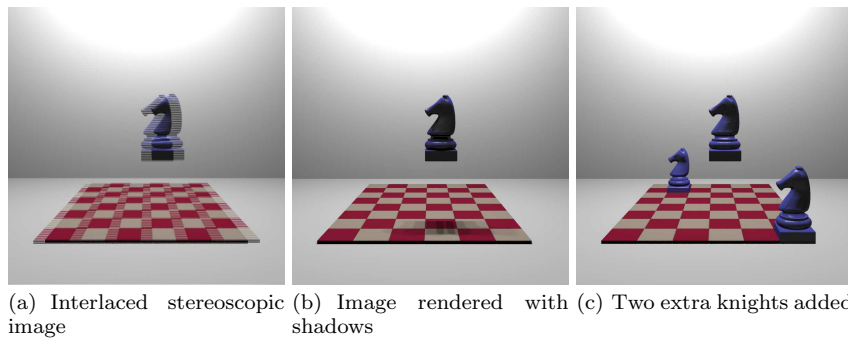

(a) Interlaced stereoscopic image

(b) Image rendered with shadows

(c) Two extra knights added

Fig. 4 Scene conditions for the matching test

the scene without shadows and the other depth cues. All of the subjects were, however, asked to wear shutter glasses for isolating the depth cue factors from other conditions, for example, better contrast, that might influence their task performance.

For the first test, the task was to adjust the position of the knight in the horizontal plane and make it align with the square on the ground. Each subject was told that the size of the knight's base and the square are physically the same. For the group viewing the images with the depth cue of relative size, the subjects were also told that the extra two knights were physically the same size. We used the same hardware and display settings as described in the resolution experiment to run the test.

For the second test, we replaced the square with a chess board as shown in figure 4(a)-(c). There were a total of 80 subjects taking part in this second test. The subjects were divided into four viewing groups, for the three viewing conditions as illustrated in figure 4(a)-(c) and one control group, and all wore shutter glasses during the test. For each category of visual stimulus, the floating knight was placed at six different locations, but with same height. The task here for the subject was to find a matching square on the chessboard according to the knight's positions. That is, the subject needed to make judgement over which square the knight is floating.

3.4 Depth judgement experiment: Results

Figure 5(a)-(b) shows the results from the second experiment. In the two bar charts, the viewing groups are laid out on the horizontal axis while the vertical axis represents the error rate. As can been seen in both bar charts, the subjects in the stereoscopic viewing group performed the task with less error and thus higher accuracy. We further analyzed the data obtained from both experiments using One-way Anova with post-hoc analysis. For the adjusting test, we found that except for the group viewing shadowed scene, there are significant differences between the performance of the stereoscopic viewing group and the other three groups, including the control group. The p-values of the stereoscopic viewing against depth cueing with shadows, no shadow, relative size, and texture were: 0.554, 0.009, 0.005, and 0.016. A significant level 0f 0.05 was assumed for the Anova test. The relative size and texture gradient as depth cues are not so reliable in this visual task, as shown by the error rates which are significantly higher in the corresponding viewing groups. We found the similar result in the matching test. The p-values of the stereoscopic viewing against depth cueing with shadows, no shadow, and relative size were: 0.796, 0.023, 0.015. The performance from the control group (no shadow) also indicated that without stereoscopic depth cues, we need to compute more monocular depth cues such as shadows, to have the same task performance.
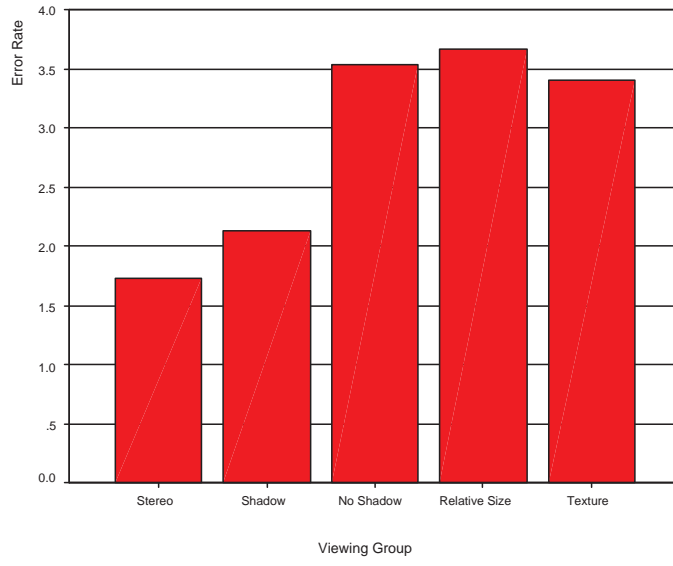
(a) The adjusting test
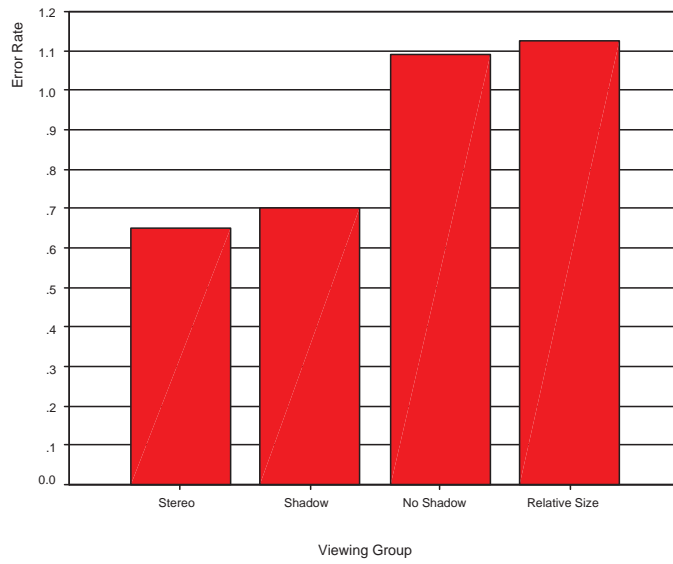
(b) The matching test

Fig. 5 Results of the depth judgement experiment

4 Discussion

4.1 Limitations of the stereoscopic display

From the results of the resolution experiment, we discovered that it makes no significant difference to produce a stereoscopic image with the resolution setting of 800×800 or

480×480, if viewed stereoscopically. Although this might help save rendering effort on producing an image with higher resolution, it also reveals the problems of the shutter glasses. The way the shutter glasses work is to synchronize the refresh rate of its LCD screen in front of each eye with the monitor. The monitor thus needs to be adjusted to a higher refresh rate to obtain a stable view for each eye. Although the flickering effect can be reduced by increasing the screen scan frequency, the images still look dimmer through the glasses because the lightness is reduced in half during the view-blocking process. As a result of this, the contrast of the image is reduced and the detectability of the resolution changes in images drops consequently. Nevertheless, with the help of the shutter glasses and the manipulation of different combinations of the resolution settings for the stereoscopic pair, we can still show that at least one of the views can be degraded in quality to some extent.

4.2 Shadow as a depth cue

In the second experiment, the virtual scene is lit by five point light sources placed at a central position on the ceiling of the room. The shadows formed by the knight occluding the ray path in between the light sources and the floor should approximately indicate where the knight is floating above. In our visual tasks, this could give the subjects some kind of guidance of the knight's position so all they had to do is to observe the locations of the shadows. As can be seen from the results of the second experiment, the difference of subjects' performances between the group viewing stereoscopic images and the group judging depth according to shadows is not as significant as compared to the depth cueing with relative size or texture gradient. However, in this case the light sources are approximately right above the target object making the task much easier. If we move the light sources to other locations or the lighting in the virtual is complex to make the environment highly-diffused, then the shadow as a depth cue may not be as reliable as shown in our experiment. On the other hand, the depth cueing from binocular disparity is invariant in different lighting conditions and should provide a more accountable source of depth information.

4.3 Saving rendering time

Figure 6 and 7(a)-(b) show the rendering time for the visual stimuli. As illustrated in figure 6, the rendering time goes up significantly as the resolution increases. If the image is to be viewed stereoscopically with shutter glasses, we can simply maintain a resolution level of 480×480, or even 560×560 for reassurance. A substantial amount of time can be saved by rendering the scene at this lower resolution.

In figures 7(a)-(b), we can see that by not tracing extra shadow rays, the total time to render both images of the stereo pair is still *less* than time to render the monocular image with shadows.
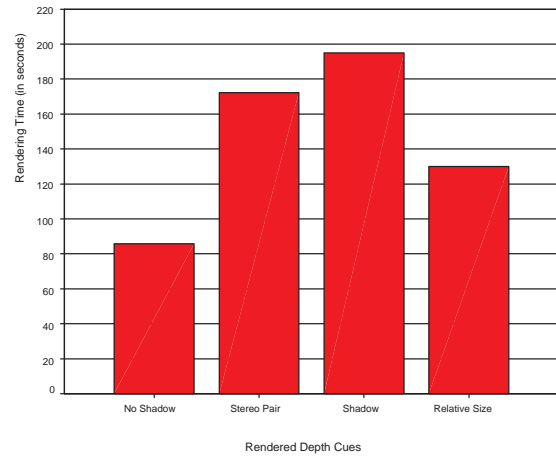
Fig. 6 Rendering time for different levels of resolution settings

(a) Rendering time of the scene for adjusting test

(b) Rendering time of the scene for matching test

Fig. 7 Rendering time for the chess scene

5 Conclusion and future work

In this paper, we investigated the possibility of reducing computational effort, both at the image level and the algorithmic level, for realistic rendering in stereoscopic virtual environments. By incorporating the findings of the psychophysical investigations into the image generation process, we can now produce physically less qualitative, but perceptually equivalent images for stereoscopic viewing. We can conclude that, with

evident support from previous work on evaluating the efficiency of depth cues, stereopsis is a strong and comparatively reliable cue for judging how far away the object is. By not rendering shadows or other depth cues, which are not so effective compared to stereopsis, we can save significant rendering time and produce images more efficiently, while improving task performance.

The resolution as a scale for indicating the quality of the image may not be sufficient. In terms of developing more efficient rendering strategies, a method for discovering the optimal number of light rays that need to be traced for a perceptually-equivalent output might be a better approach. In the future we shall exploit the binocular fusion further to investigate whether we could trace even fewer rays for either of the stereoscopic image pairs. In addition, ideally, the shadow computation should not be totally ignored, because it is widely recognized as a strong cue to increase the visual realism of the synthetic images. In our future work, we will also investigate the subtlety of shadow rendering and how we can reduce its rendering cost when performing a visual task.

References

1. S. J. Adelson, L. F. Hodges, Visible surface ray-tracing of stereoscopic images, in: ACM-SE 30: Proceedings of the 30th annual Southeast regional conference, ACM Press, 1992, pp. 148–156.
2. G. J. Ward, The radiance lighting simulation and rendering system, in: SIGGRAPH '94: Proceedings of the 21st annual conference on Computer graphics and interactive techniques, ACM Press, 1994, pp. 459–472.
3. A. Watt, F. Policarpo, The Computer Images, Addison-Wesley Longman Limited, 1999.
4. C. G. Barbour, G. W. Meyer, Visual cues and pictorial limitations for computer generated photo-realistic images, Visual Computer 9 (3) (1992) 151–165.
5. K. Myszkowski, T. Tawara, H. Akamine, H.-P. Seidel, Perception-guided global illumination solution for animation rendering, in: SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques, ACM Press, 2001, pp. 221–230.
6. H. Yee, S. Pattanaik, D. P. Greenberg, Spatiotemporal sensitivity and visual attention for efficient rendering of dynamic environments, ACM Trans. Graph. 20 (1) (2001) 39–65.
7. G. W. Meyer, H. E. Rushmeier, M. F. Cohen, D. P. Greenberg, K. E. Torrance, An experimental evaluation of computer graphics imagery, ACM Trans. Graph. 5 (1) (1986) 30–50.
8. A. McNamara, A. Chalmers, T. Troscianko, I. Gilchrist, Comparing real & synthetic scenes using human judgements of lightness, in: Proceedings of the Eurographics Workshop on Rendering Techniques 2000, Springer-Verlag, 2000, pp. 207–218.
9. K. Cater, A. Chalmers, G. Ward, Detail to attention: exploiting visual tasks for selective rendering, in: EGRW '03: Proceedings of the 14th Eurographics workshop on Rendering, Eurographics Association, 2003, pp. 270–280.
10. L. Kjelldahl, M. Prime, A study on how depth perception is affected by different presentation methods of 3d objects on a 2d display, Computer and Graphics 19 (2) (1995) 199–202.
11. L. C. Wanger, J. A. Ferwerda, D. P. Greenberg, Perceiving spatial relationships in computer-generated images, IEEE Comput. Graph. Appl. 12 (3) (1992) 44–51, 54–58.
12. P. Servos, A. Goodale, L. S. Jakobson, The role of binocular vision in prehension: A kinematic analysis, Vision Research 32 (8) (1992) 1513–1521.
13. G. S. Hubona, P. N. Wheeler, G. W. Shirah, M. Brandt, The relative contributions of stereo, lighting, and background scenes in promoting 3d depth visualization, ACM Trans. Comput.-Hum. Interact. 6 (3) (1999) 214–242.
14. L. Wanger, The effect of shadow quality on the perception of spatial relationships in computer generated imagery, in: SI3D '92: Proceedings of the 1992 symposium on Interactive 3D graphics, ACM Press, 1992, pp. 39–42.

15. H. H. Hu, A. A. Gooch, W. B. Thompson, B. E. Smits, J. J. Rieser, P. Shirley, Visual cues for imminent object contact in realistic virtual environment, in: VIS '00: Proceedings of the conference on Visualization '00, IEEE Computer Society Press, 2000, pp. 179–185.

16. E. Horvitz, J. Lengyel, Perception, attention, and resources: A decision-theoretic approach to graphics rendering, in: Proceedings of the 13th Conf. on Uncertainty in Artificial Intelligence, Morgan Kaufman:San Francisco, 1997, pp. 238–249.

17. M. G. Perkins, Data compression of stereopairs, IEEE Transactions on Communication 40 (4) (1992) 684–696.

18. M. Siegel, S. Sethuraman, J. S. McVeigh, A. Jordan, Compression and interpolation of 3d-stereoscopic and multi-view video, in: Stereoscopic Displays and Virtual Reality Systems IV, Vol. 3012, 1997, pp. 227–238.

19. S. H. Kim, M. Siegel, J.-Y. Son, Synthesis of a high resolution 3d-stereoscopic image from a high resolution monoscopic image and a low resolution depth map, in: Proceedings of the 1998 SPIE/IS&T Conference, Vol. 3295A, 1998, pp. 76–86.

20. S. Badt, Jr., Two algorithms taking advantage of temporal coherence in ray tracing, Visual Computer 4 (1) (1988) 123–132.

21. eDimensional, Wireless 3D shutter glasses, eDimensional, Inc., http://www.edimensional.com/, updated 2004.

22. G. Ward Larson, R. Shakespeare, Rendering with RADIANCE: The art and science of lighting simulation, Morgan Kauffman: San Francisco, 1998.

23. 3DCombine, A stereoscopic image generation program, Shareware, http://www.3dcombine.com/.

24. Presentation, High-precision program for stimulus delivery and experimental control for behavioral and physiological experiments, Neurobehavioral Systems, Inc., http://nbs.neuro-bs.com/.