A PSYCHOPHYSICAL STUDY OF FOVEAL GRADIENT BASED SELECTIVE RENDERING

Veronica Sundstedt
Department of Computer Science, Bristol University, Bristol, United Kingdom

Keywords: Perception, user studies, attention, selective rendering, inattentional blindness, importance maps.

Abstract: High-fidelity rendering of complex scenes at interactive rates is one of the primary goals of computer graphics. Since high-fidelity rendering is computationally expensive, perceptual strategies such as visual attention have been explored to achieve this goal. *Inattentional Blindness* (IB) experiments have shown that observers conducting a task can fail to see an object, although it is located within the foveal region (2°). However, previous attention based algorithms assumed that IB would be restricted to the area outside the foveal region, selectively rendering the areas around task-related objects in high quality and the surrounding areas in lower quality. This paper describes a psychophysical forced-choice preference experiment assessing if participants, performing a task or free-viewing animations, would fail to notice rendering quality degradation within the foveal region. The effect of prior knowledge on the level of perceived quality is also studied. The study involves 64 participants in four conditions: performing a task, or free-viewing a scene, while being *naive* or *informed* about assessing rendering quality. Our results show that participants fail to notice the additional reduction in quality, decreasing the overall computation 13 times. There was also a significant difference in the results if free-viewing participants were informed.

1 INTRODUCTION

Obtaining realistic images has always been one of the major goals in computer graphics. Applications ranging from entertainment, lighting design, to archaeological reconstructions, and scientific visualisations have required realistic models of light propagation and scattering. One traditional method for speeding up global illumination calculations is to distribute the workload over several machines, each processing a part of the scene in parallel. However, even in the parallel case, rendering efforts are spent on improving details that would not be perceived by a human observer.

For many applications, rather than using more machines to accelerate the rendering, it is more efficient to reduce the number of computations needed. It is attractive to improve the efficiency of rendering by performing less work. Hence perceptually-based rendering algorithms have become an important research topic in computer graphics. The goal of perceptually-based rendering algorithms is to significantly reduce computation that is necessary to obtain an image that is perceptually indistinguishable from a fully converged solution or *gold standard*, which is a hypothetical perfect rendering (Woolley et al., 2003). Recently, models of the human visual system (HVS), in particular those based on visual attention processes, have been used in perceptually assisted renderers to make progress towards this goal.

1.1 Visual Attention Processes

Human visual perception is a selective process in which a part of the observed environment is chosen for further processing in the visual cortex of the brain. Although, the perception of an environment does not only depend on the sensory input of the observer, but also on the visual task performed (Cater et al., 2003). When free viewing a scene, the attentional processes of the observer are guided by automatic low-level vision. These are normally referred to as *bottom-up* processes (James, 1957). Low-level, bottom-up features which influence visual attention include contrast, size, shape, colour, brightness, orientation, edges, and motion (Itti et al., 1998). In contrast, when performing a visual task it is directed in the pursuit of our goals (Yarbus, 1967). These are normally referred to

Sundstedt V. (2008).
A PSYCHOPHYSICAL STUDY OF FOVEAL GRADIENT BASED SELECTIVE RENDERING.
In *Proceedings of the Third International Conference on Computer Graphics Theory and Applications*, pages 207-214
DOI: 10.5220/0001097402070214

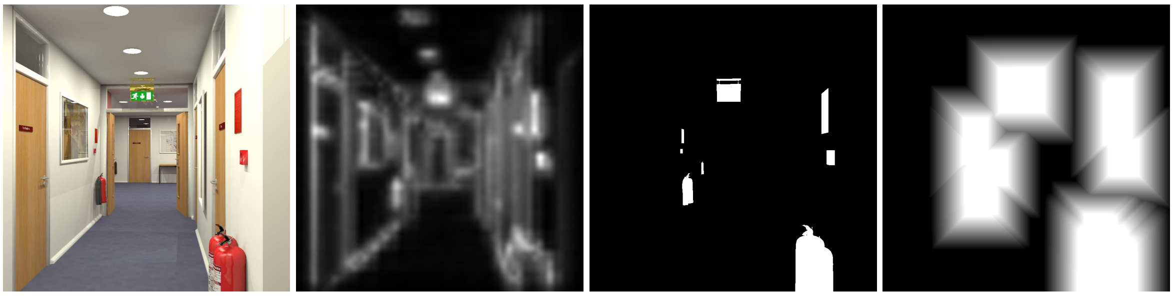

(a) (b) (c) (d)

Figure 1: Map examples from a corridor scene (Frame 1): (a) high quality rendering, (b) saliency map, (c) task objects, and (d) task map with foveal gradient angle.

as *top-down* processes (James, 1957). In both cases, the HVS focuses its attention on certain objects at the expense of other details in the scene. The *inattentional blindness* (IB) phenomenon (Mack and Rock, 1998) relates to our inability to perceive features or objects in a visual scene if we are not attending to them. It is suggested that IB occurs because the observer is focussed on performing a specific task, i.e. there is no conscious perception without attention (Mack and Rock, 1998).

1.2 Visual Attention Models

One method of reducing computation, while maintaining a result with high perceptual quality, is to adapt the rendering parameters of the image based on models of human visual attention processes. In this way areas which are perceptually more relevant will receive further improvement. This results in an image with a spatially shifting degree of accuracy, referred to as *selective rendering*. In selective rendering algorithms the bottom-up process is often modelled using a *saliency map* derived from the computational model developed in (Itti et al., 1998). This model extract salient features based on colour, intensity, and orientation. An example of a saliency map for the corridor scene is shown in Figure 1 (b).

The idea of using *task maps* to model the top-down process in selective rendering was introduced in (Cater et al., 2003). A task map is a grey-scale image, which consists of objects related to the task in white and the surrounding areas in black, as shown in Figure 1 (c). There is also an option to graduate the shade between black and white over an area that corresponds to the image registered on the foveal region (2°), mimicking the high visual acuity of the HVS (Snowden et al., 2006), as shown in Figure 1 (d).

Previous rendering algorithms using task maps have assumed that IB would be restricted to the area outside the foveal region. However, in (Mack and Rock, 1998) it was shown that participants failed to notice a critical stimulus, even though it appeared within the centre of their field of view and coincided with a fixation. This inspired the psychophysical forced-choice preference experiment presented in this paper. The psychophysical experiment investigates if previous work can be improved by using selective rendering *within* the foveal region in the presence of a high-level task focus or when free-viewing a scene. It also studies if the effect of prior knowledge (being informed about assessing rendering quality) would alter the level of perceived quality.

2 RELATED WORK

Extensive overviews of perceptually adaptive graphics techniques are given in (McNamara, 2001; O'Sullivan et al., 2004). A more recent and comprehensive summary of different selective rendering techniques are presented in (Debattista, 2006; Sundstedt, 2007). The following sections describe rendering techniques which have taken into account visual attention models.

2.1 Gaze-Contingent Techniques

Gaze-contingent displays (GCDs) (Loschky et al., 2003) track the user's attention using an eye-tracker and render information in full detail only at the observer's current focus of attention. To prevent observers from noticing the lower quality in the peripheral regions, the size of the foveal regions is based on the the extent of the user's perceptual span (2°). *Focus plus context screens* include both high and low detail by combining a wall-sized low-resolution display with an embedded high-resolution screen (Baudisch et al., 2003). The user uses the mouse to pan the display content into the high-resolution area. *Easily perceived displays* aim to direct the attention of the viewer (Baudisch et al., 2003). This has been exploited in for example art where eye-tracking information from one participant has been used to create aesthetically pleasing images.

2.2 Level of Detail

A bottom-up attention model for Level of Detail (LOD) simplification taking into account size, position, motion, and luminance was presented in (Brown et al., 2003). The mesh saliency algorithm in (Lee et al., 2005) exploits the centre-surround operation from (Itti et al., 1998) to generate more visually pleasing images by prioritising the geometry as the most important salient feature. In (Howlett et al., 2005) polygon models were simplified based on saliency identified using eye-tracking of human observers. In (Yang and Chalmers, 2005) the LOD was reduced of objects not related to the task being performed by the observer. A perceptual stategy for collision detection was proposed in (O'Sullivan, 2005). For example, a lower LOD could be used between objects that are not being focused upon.

2.3 Attention based Rendering

A first attempt to account for selective visual attention in global illumination rendering of dynamic environments was proposed in (Yee et al., 2001). Here the saliency model (Itti et al., 1998) was extended to include motion. An *Aleph map* was created that combined the saliency map with a spatiotemporal contrast sensitivity function (CSF) (Myszkowski et al., 1999). The Aleph map was used to guide the search radius accuracy for the interpolation of irradiance cache values so that perceptually important regions would be more accurate. This made the indirect lighting computations more efficient since larger error can be tolerated in less salient regions.

One of the first uses of an attention model in a real-time application was proposed in (Haber et al., 2001). One of the problems with interactive walkthroughs is to render non-diffuse objects in real-time. The non-directional light component can be efficiently pre-computed, whereas directional light from glossy surfaces, for example, can only be computed during execution. Haber *et al.* rendered a pre-computed global illumination solution using graphics hardware while updating non-diffuse objects dynamically with a ray tracer. The selection of these types of objects was based on the saliency model (Itti et al., 1998). Haber *et al.* also took top-down behaviour into account by weighting the saliency model with bias towards objects in the centre.

The idea of exploiting IB in a rendering framework was first proposed in (Cater et al., 2002). Cater *et al.* investigated if areas in a scene that normally would attract attention would do so in the presence of a task. In their experiments observers were shown two animations. One was rendered in high quality and the other one either in low quality or selectively composited. The selectively composited animation had only the pixels in the 2° foveal region centred around the location of the task object in high quality. The quality was then blended to lower quality within an angle of 4° based on findings by (McConkie and Loschky, 1997).

Half the subjects were instructed to simply watch the animations and the other half were asked to perform a task. After completion of the experiment the participants were asked if they noticed a quality difference. The results showed a significant difference for all pair-wise comparisons to high quality while free-viewing the animations. While performing a task there was only a significant difference between the high quality and low quality comparison. Cater *et al.* (Cater et al., 2003) used these results as a basis for a task-based perceptual rendering framework, which combined predetermined task maps with a spatiotemporal CSF to guide a progressive animation system.

3 SELECTIVE RENDERING

To selectively render the stimuli used for the experiment presented in this paper a *region-of-interest* (ROI) rendering system was used (Debattista, 2006; Sundstedt, 2007). For the convenience of the reader, the technical details that underpins this work is briefly reviewed. The rendering system is composed of two major processes:

Region-of-interest (ROI) guidance uses a combination of saliency and a measure of task relevance to direct the rendering computation (in a combined importance map).

Selective rendering corresponds to the traditional rendering computation (Ward, 1994). However, computational resources are focused on parts of the image which are deemed more important by the ROI guidance.

The process begins with a rapid image estimate (in the order of ms) of the scene using a quick rasterisation pass in hardware (Longhurst, 2005). This estimate can be used in two ways. Firstly for building the task map by identifying user-selected task objects, and secondly, by using it as an input to a saliency generator. During creation of the task map the program

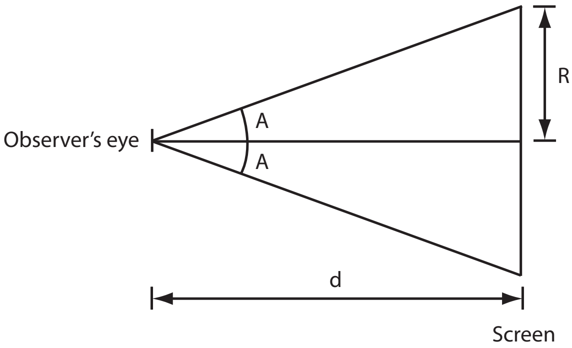

Figure 2: Relation between visual angle and pixel radius on screen.

reads in the geometry information and a list of predefined task objects. It then produces a map with task objects in white and the other geometry in black, as shown in Figure 1 (c). The task map can also take into account the area the fovea in the eye covers in the environment. The foveal region in the eye, where the visual acuity is highest, is only the central 2° of the visual field.

When an observer is watching the environment, this area corresponds to a region on an object. In the context of this paper this object is a computer monitor. Although the size of the foveal region is fixed within the eye, the area it covers on the object varies with the distance (d) between the eye and the object. As the size of regions that project an image onto the foveal region changes, for simplicity the distance on the object as an angle (A) subtended at the eye is measured, as shown in Figure 2. The radius (R) of this circular region can be measured in pixels if the *ratio* between the screen resolution and monitor size is known, as computed by equation 1:

$$R = ratio(d \cdot \tan A) \quad (1)$$

The region on the screen is a circle with radius 1°, but for simplicity squares with a width and height of 2° are used. To account for a gradient (which blends the quality from high to low), a graduated fill can be used between a square of 2° and 4°, as shown in Figure 1 (d). This graduated fill was proposed in (Cater et al., 2002) after findings described in (McConkie and Loschky, 1997).

In the creation of the saliency map the image estimate serves to locate areas where an observer will be most likely to look. The estimate of the scene contains only direct lighting, but a simple shadow and reflection calculation is also included. The saliency estimation is carried out by using the existing method proposed in (Itti et al., 1998) and is computed in 2-3 seconds per frame.

A hardware implementation can generate a saliency map in the order of tens of milliseconds (Longhurst, 2005). The two maps have previously been used separately and in combination to form an *importance map* (IM) which accounts for both the bottom-up and top-down processes (Sundstedt et al., 2005). The values in the importance map are used to direct the rendering. However, this paper is only using the task map.

3.1 Previous Visual Trial

In (Sundstedt et al., 2005) a psychophysical experiment was performed which showed that both saliency maps and task maps (including the foveal region gradient) can be used successfully to selectively render in high quality only the important areas of animations while performing a task or free-viewing a scene. The experiment investigated how models of visual attention, low-level and task-dependent on their own, and as a hybrid would work in a selective rendering framework. 160 participants took part in this trial (124 men and 36 women; age range: 18-39). There were in total ten groups with 16 participants in each group.

Each participant was shown a high quality (HQ) animation and a HQ, low quality (LQ) or selectively rendered animation of the corridor scene, shown in Figure 1 (a). The selectively rendered animations were either generated using a task map (TQ), saliency map (SQ) or a linear combination of the two (IQ). Two HQ animations were shown in one group to conclude if where performance was caused by chance, the chance performance was not based on there being a strong preference for one of the two scenes ($p > 0.05$). The HQ animations were rendered using an extended version if the *Radiance* (Ward, 1994) renderer `rpict` using 16 rays per pixel and a specular threshold value of 0.01. The LQ animations were rendered using 1 ray per pixel and a specular threshold value of 1. The extended rendering system is further described in (Debattista, 2006; Sundstedt, 2007).

Half the subjects were asked to perform a task, in this case counting the items related to fire safety, whereas the other group were only shown two animations without any previous instructions. After completion of the experiment, each participant determined which of the two animations they thought had the worse rendering quality. The results confirmed, using a one-sample chi-square test (df=1, critical value 3.84 at 0.05 level of significance), that for both participants performing a task and free-viewing in the HQ/TQ, HQ/SQ and HQ/IQ, the difference in proportions was not significant ($p > 0.05$). From this it was concluded that the two animations were perceived as a similar quality. For participants free-viewing and performing a task in the HQ/LQ condition there was a significant difference ($p < 0.05$).

As IB is described in (Mack and Rock, 1998), observers conducting a task can fail to see an object, although it is located within the foveal region. Based on this fact and the result of the previous visual trial an additional experiment was constructed. The new psychophysical experiment presented in this paper investigates whether participants fail to notice a reduction in quality within the foveal region in the presence of a high-level task focus or when free-viewing a scene. It also studies if the effect of prior knowledge (being informed about assessing rendering quality) would alter the level of perceived quality.

4 NO FOVEAL REGION STUDY

Based on the result in (Mack and Rock, 1998) it is hypothesised that in the presence of a high-level task focus participants would fail to notice quality degradations even within the foveal region. The relationship between the foveal region and an area on the screen is previously explained. This region might contain other objects unrelated to the task.

If this hypothesis is true, non-task areas within the foveal region can be rendered with lower quality than the task objects saving computation. To study if participants would fail to notice the reduction in quality a modified task map was used for the corridor scene animations. This time the fovea angle gradient was excluded, as shown in Figure 1 (c).

The categorical variables in the experiment are: *condition* (HQ/TWFQ - Task Without Fovea Quality) and *preference* (correct/incorrect). The condition is the manipulated independent variable and the preference is the dependent variable. The two hypotheses tested in the experiment are stated below:

Hypothesis 1: There is no significant difference in the level of correct/incorrect responses for participants performing a task in or free-viewing an animation rendered in HQ and one rendered selectively with only the task objects in HQ (TWFQ).

Hypothesis 2: There is no significant difference in the level of correct/incorrect responses between a participant being naive to the purpose of the experiment and informed participants.

4.1 Participants and Setup

64 participants took part in the experiment (51 men and 13 women; age range: 20-37). The participants taking part in the experiments were all undergraduate or graduate students. Subjects had a variety of experience with computer graphics, and all self-reported normal or corrected-to-normal vision. There were in total four groups with 16 participants in each group.

Two groups were performing a task, the other two were free-viewing the scene. To study the effect of prior knowledge half the participants were *naive*, or uninformed, to the forced-choice preference. The other half were aware that they would be asked about the *quality* of the animations they had seen. They were thus *informed* as to the purpose of the experiment.

All stimuli were presented on a 17” LCD monitor (1280×1024 resolution, 60 Hz refresh frequency). The effect of ambient light was minimised. The participants were seated on an adjustable chair, with their eye-level approximately level with the centre of the screen. The viewing distance from the participants to the screen was around 60 cm. All stimuli were rendered at 900×900 resolution and displayed in the centre of the screen with a black background.

4.2 Stimuli

Two different walkthroughs of a corridor scene were used, which are termed corridor A, and corridor B. This was done to avoid familiarity effects that might have influenced the scan path of the observers. Both the animations were rendered with different views and the location of the objects changed within each scene. Each animation contained the same number of task-related objects (15) but not the same number of non-task related objects. An identical type of camera path was used for both animations.

For both corridor scenes a new animation was rendered, without the foveal angle gradient added. This stimuli was termed TWFQ. Using the TWFQ maps only the task objects were rendered in HQ and the remaining regions in LQ. The same HQ animations were used as in the previous visual trial.

Figure 3 shows the timing comparison between a HQ, TQ, LQ animation used in the previous visual trial, and a new TWFQ animation. Rendering the new TWFQ frames took on average 10 minutes, which is around four times faster then the TQ renderings. Rendering the entire frame to the same detail as the task objects in the new TWFQ renderings therefore took on average 13 times longer. Computing the TWFQ condition was cheaper than all other selectively rendered stimuli presented in the previous visual trial (Sundstedt et al., 2005).

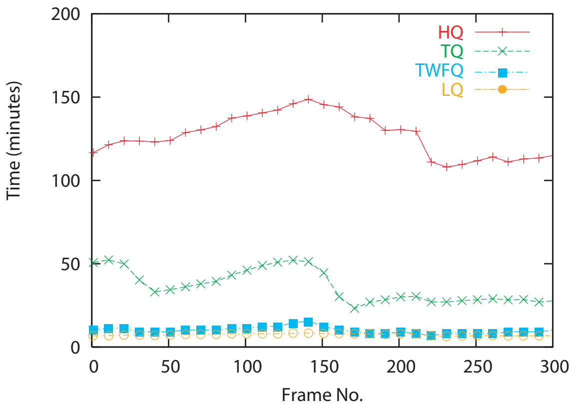

Figure 3: Timing comparison for the corridor scene between a high quality (HQ), foveal region quality (TQ), low quality (LQ) animation used in the previous visual trial, and a new no foveal region (TWFQ) animation.

4.3 Procedure

The same experimental setup and method as in the previous visual trial was used for consistency. The animations were displayed in the centre of the screen with a black background. Each animation was 17 seconds long, including a countdown before the animations started. Following a verbal introduction to the experiment, each participant was shown two animations, one from corridor A and one from corridor B. One of the animations was always HQ while the other one was a selectively rendered animation using the new TWFQ animation. Each animation was viewed only once. The order in which the subjects saw their two animations was also altered to avoid any bias.

Before beginning the experiment, half the subjects read a sheet of instructions on the procedure of the particular task they were to perform. As before, these subjects were asked to take the role of a fire security officer whereby the task was to count the total number of fire safety items in each of the two animations. Each of the participants in these groups were shown an example of what kind of fire emergency items the scene could contain. For the participants performing a task while watching their stimuli, it was also confirmed that the task was understood prior to the start of the experiment. The participants in the other half were simply shown the animations. After completion of the experiment, both groups were asked which of the two animations they thought was of worse quality. If a participant could not determine which one they thought had the worse quality, they were asked to choose either A or B (2AFC).

4.4 Results

Figure 4 shows the results of the experiment. In each pair of conditions, a result of 50% (8 out of 16 participants) correct selection in each case is the unbiased ideal. This is the statistically expected result in the absence of a preference towards one scene, and indicates that no differences between the high quality and a lower quality animation were perceived. The results for the participants performing a task are shown to the left and the free-viewing results are shown to the right. The left bar in each graph shows the result from the naive participants, while the bar to the right shows the result from the informed participants.

The results show that 62.5% of the naive participants performing a task reported a correct result in the HQ/TWFQ condition. When the participants were informed about assessing rendering quality, 56.25% reported a correct result in the same condition while performing a task.

The percentages for the naive and informed participants free-viewing the same condition were 56.25% and 93.75% respectively. While the correct classification frequency increased for the informed participants free-viewing the scene, the level when performing a task while being informed was not altered greatly.

4.5 Statistical Analysis and Discussion

The results were analysed statistically using the Chi-square test to determine any significance. This test allows us to determine if what is observed in a distribution of reported frequencies (correct/incorrect) would be what is expected to occur by chance. The reported frequencies were compared to an expected 50/50 data to ascertain whether the participants perceived the difference in quality.

The statistical analysis of the results confirmed that for naive participants performing a task in the HQ/TWFQ condition, the difference in proportions was not significant $\chi^2(1, N = 16) = 1, p = 0.32$. The results were also not significant for the informed participants that performed a task and the naive participants free-viewing the two animations, both $\chi^2(1, N = 16) = 0.25, p = 0.62$. This indicates that the two different animations were perceived as a similar quality in these three conditions.

The results for the participants free-viewing the two animations while being informed differed from the other three conditions. In this case, almost all participants managed to perceive the higher quality animation, $\chi^2(1, N = 16) = 6.25, p = 0.01$. This could potentially be caused by the fact that they were not just free-viewing the scene anymore.

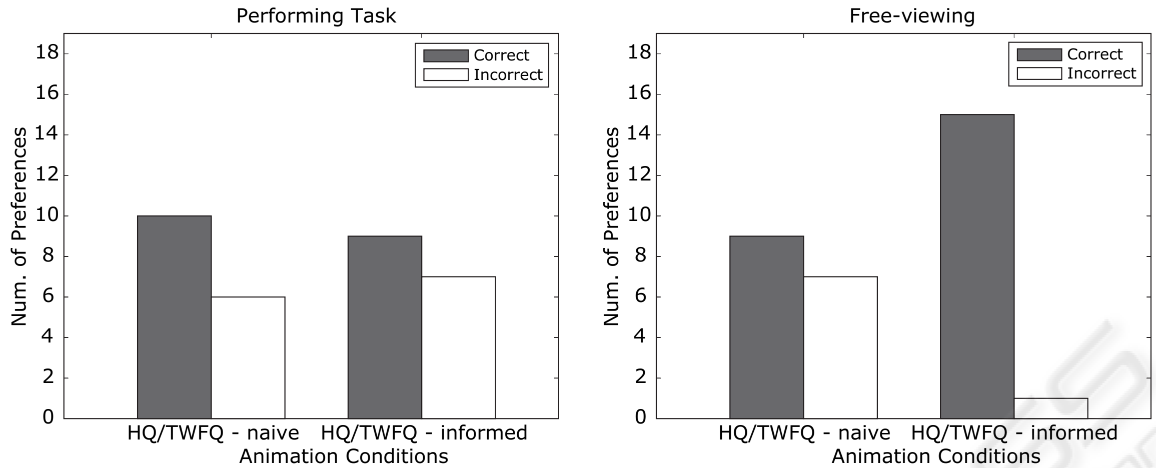

Figure 4: Results from the no foveal region experiment for the two conditions: (left) performing a task (counting fire safety items) vs. (right) watching the animations.

Asking people to watch the animations, while being informed about assessing rendering quality, might have constituted an implicit task in itself and potentially made them focus on parts of the scene where quality differences were most likely to occur, for example edges.

Mack and Rock (Mack and Rock, 1998) observed that conscious perception requires attention, which agrees with the obtained results. The participants in the new experiment failed to notice the low rendering quality within the foveal region of the animations. This indicates that rendering quality can be reduced within the foveal region and that IB can in fact be exploited in selective rendering.

5 CONCLUSIONS

This paper presented a psychophysical forced-choice preference experiment assessing if participants, performing a task or free-viewing animations, would fail to notice rendering quality degradation within the foveal region. Previous work in task-related rendering assumed that IB would be restricted to the area outside the foveal region. However, IB studies have shown that observers conducting a task can fail to see an object although it coincides with a fixation. The experiment presented in this paper showed for the first time that IB can be exploited in selective rendering, by using a sharp cut-off around the task objects, while maintaining a perceptually high quality result.

The only observers to show a statistically significant ability to detect the quality difference were those who were free-viewing the HQ/TWFQ pairs after being informed that they would be asked to judge the quality of the animations. This is an interesting result which agrees with the findings in (Mack and Rock, 1998) which propose that there is no conscious perception without attention.

The experiment indicated that reductions in quality can be achieved through analysis of high-level visual processing, and is not bound to low-level biological processes, such as knowledge of the foveal region. This is an improvement on previous work which added a foveal region around the task objects. By being able to reduce the quality further than previously proposed ROI rendering techniques, additional computation savings can be made.

Overall, the experiment shows that there is an opportunity in ROI rendering, but that more research must be done to better understand how rendering quality affects human perception. It could also be possible that using participants which are familiar with artifacts can lead to an overly strict criteria. Hence, as future work it would be interesting to study if there are significant differences between participants being familiar with the concept of rendering quality and those who are not. As future work it would also be interesting to study the eye movements of the participants while they make the quality judgements.

ACKNOWLEDGEMENTS

The author would like to thank Kurt Debattista for implementation of ROI rendering in *Radiance*, Peter Longhurst for the modified *Snapshot* program and Sumaya Ahmed for collaboration on the task map plug-in. Thanks also to Andrew Moss, Alan Chalmers, everyone that participated in the experiments, and Patrick Ledda for the original corridor model. This research was partly done under the sponsorship of the Rendering on Demand (RoD) project within the 3C Research programme and by the European Union within the CROSSMOD project (EU IST-014891-2).

REFERENCES

Baudisch, P., DeCarlo, D., Duchowski, A. T., and Geisler, W. S. (2003). Focusing on the essential: considering attention in display design. *Commun. ACM*, 46(3):60–66.

Brown, R., Cooper, L., and Pham, B. (2003). Visual attention-based polygon level of detail management. In *GRAPHITE '03: Proceedings of the 1st international conference on Computer graphics and interactive techniques in Australasia and South East Asia*, pages 55–62, New York, NY, USA. ACM Press.

Cater, K., Chalmers, A., and Ledda, P. (2002). Selective quality rendering by exploiting human inattentional blindness: looking but not seeing. In *VRST '02: Proceedings of the ACM symposium on Virtual reality software and technology*, pages 17–24, New York, NY, USA. ACM Press.

Cater, K., Chalmers, A., and Ward, G. (2003). Detail to attention: exploiting visual tasks for selective rendering. In *EGRW '03: Proceedings of the 14th Eurographics workshop on Rendering*, pages 270–280, Aire-la-Ville, Switzerland, Switzerland. Eurographics Association.

Debattista, K. (2006). *Selective Rendering for High-Fidelity Graphics*. PhD thesis, University of Bristol.

Haber, J., Myszkowski, K., Yamauchi, H., and Seidel, H.-P. (2001). Perceptually guided corrective splatting. *Computer Graphics Forum*, 20(3):142–152.

Howlett, S., Hamill, J., and O'Sullivan, C. (2005). Predicting and Evaluating Saliency for Simplified Polygonal Models. *ACM Trans. Appl. Percept.*, 2(3):286–308.

Itti, L., Koch, C., and Niebur, E. (1998). A Model of Saliency-Based Visual Attention for Rapid Scene Analysis. *IEEE Trans. Pattern Anal. Mach. Intell.*, 20(11):1254–1259.

James, W. (1957). *The Principles of Psychology*. Dover Publications Inc.

Lee, C. H., Varshney, A., and Jacobs, D. W. (2005). Mesh saliency. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers*, pages 659–666, New York, NY, USA. ACM Press.

Longhurst, P. (2005). *Rapid Saliency Identification for Selectively Rendering High Fidelity Graphics*. PhD thesis, University of Bristol.

Loschky, L. C., Mcconkie, G. W., Reingold, E. M., and Stampe, D. M. (2003). Gaze-Contingent Multiresolutional Displays: An Integrative Review. *Human Factors*, 45(2):307–328.

Mack, A. and Rock, I. (1998). *Inattentional Blindness*. MIT Press.

McConkie, G. W. and Loschky, L. C. (1997). Human Performance with a Gaze-Linked Multi-Resolutional Display. In *Advanced Displays and Interactive Displays First Annual Symposium*, pages 25–34.

McNamara, A. (2001). Visual Perception in Realistic Image Synthesis. *Computer Graphics Forum*, 20(4):211–224.

Myszkowski, K., Przemyslaw, R., and Tawara, T. (1999). Perceptually-informed Accelerated Rendering of High Quality Walktrough Sequences. In *Eurographics Workshop on Rendering*, pages 13–26. The Eurographics Association.

O'Sullivan, C. (2005). Collisions and Attention. *ACM Trans. Appl. Percept.*, 2(3):309–321.

O'Sullivan, C., Howlett, S., Morvan, Y., McDonnell, R., and O'Conor, K. (2004). Perceptually Adaptive Graphics. In *Eurographics 2004, STAR*, pages 141–164. The Eurographics Association.

Snowden, R., Thompson, P., and Troscianko, T. (2006). *Basic Vision: an introduction to visual perception*. Oxford University Press.

Sundstedt, V. (2007). *Rendering and Validation of Graphics using Region-of-Interest*. PhD thesis, University of Bristol.

Sundstedt, V., Debattista, K., Longhurst, P., Chalmers, A., and Troscianko, T. (2005). Visual attention for efficient high-fidelity graphics. In *SCCG '05: Proceedings of the 21st spring conference on Computer graphics*, pages 169–175, New York, NY, USA. ACM Press.

Ward, G. J. (1994). The RADIANCE lighting simulation and rendering system. In *SIGGRAPH '94: Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 459–472, New York, NY, USA. ACM Press.

Woolley, C., Luebke, D., Watson, B., and Dayal, A. (2003). Interruptible rendering. In *SI3D '03: Proceedings of the 2003 symposium on Interactive 3D graphics*, pages 143–151, New York, NY, USA. ACM Press.

Yang, X. and Chalmers, A. (2005). Perceptually Driven Level of Detail for Efficient Ray Tracing of Complex Scenes. In *Proceedings of Theory and Practice of Computer Graphics 2005*, pages 91–96. The Eurographics Association.

Yarbus, A. L. (1967). Eye movements during perception of complex objects. In *Eye Movements and Vision*, pages 171–196, New York. Plenum Press.

Yee, H., Pattanaik, S., and Greenberg, D. P. (2001). Spatiotemporal sensitivity and visual attention for efficient rendering of dynamic environments. *ACM Trans. Graph.*, 20(1):39–65.