EG UK Theory and Practice of Computer Graphics (2005)
L. Lever, M. McDerby (Editors)

Perceptual Level of Detail for Efficient Ray Tracing of Complex Scenes

X. Yang† and A. Chalmers

Department of Computer Science , University of Bristol, United Kingdom

Abstract

Rendering complex scenes in real time remains one of the major challenges in computer graphics. Recent research in perceptual rendering algorithms and level of detail techniques have shown that, by exploiting knowledge of the human visual system, significant computation time can be saved by only rendering in high quality those parts of the image that the user will see. This paper describes how a level of detail approach can be used to reduce overall computation time when users are undertaking a task within a virtual environment.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism;

1. Introduction

There is an increasing demand from the computer graphics industry for more realistic rendering of complex scenes at interactive rates, however, satisfactory solutions for this computationally intensive process do not yet exist. Performance increases in graphics hardware have significantly improved rendering capabilities, but concurrent increases in scene complexity and desired quality of the resultant images mean that computational demands continue to outstrip hardware resources and will do so for the foreseeable future.

Level of detail (LOD) methods are one of the recognised approaches that seek to lower overall computational time by reducing the number of polygons that need to be sent through the graphics pipeline. The majority of these approaches take advantage of geometric visibility to speed up the rendering process. Recently, a number of researches have taken advantage of the limitations of the human visual system to also reduce overall computational time by avoiding rendering those parts of a scene that are imperceptible to the human viewer.

Previous LOD research has concentrated on the rasterization pipeline to implement the LOD techniques. In this paper we will consider LODs in a ray tracer; the Maya rendering system. Furthermore, our LOD approach will extend present methods of LOD management with knowledge of the tasks that a user will perform in the virtual environment together with the saliency of objects within the scene to reduce the polygon complexity of several objects in the environment without altering the user's overall perception of the scene. The key concept that we exploit is *Inattentional Blindness*, that is the failure of a human to notice detail in objects to which he/she is not attending [MR98]. The perceptual importance of objects in the environment are used to modulate the LOD assigned to those objects. A series of psychophysical experiments were run to investigate the effectiveness of our approach. The rest of the paper is structured as follows. Section 2 describes the previous work in both the LOD and visual perception fields. Section 3 postulates an approach that exploits the knowledge of visual perception to optimise existing LOD techniques. Section 4 analyses the experimental outcome of this approach statistically using the Chi Square test. Section 5 validates the results further through a visual difference prediction algorithm VDP [Dal93]. Finally, section 6 concludes with a discussion of achievements and future work.

† Email: yang@cs.bris.ac.uk

2. Previous Work

2.1. Visual Attention

Visual attention theory seeks to explain the way humans focus their conscious attention. It is commonly accepted that attention is affected by certain visual attractors within the environment as well as predefined tasks that the user is performing [Jam90]. These concepts are typically labelled *bottom-up* and *top-down* processes respectively [Jam90].

2.2. The bottom-up process

Bottom-up describes the visual searching process in which the eyes are involuntarily drawn to salient features. Many researchers such as Itti et al. [IKN98], Yantis [Yan96] and Koch et al. showed that the visual system is highly sensitive to features such as edges, abrupt changes in colour, and sudden movement. For example, a big red object is visually conspicuous and will catch a person's attention straight away when it is viewed amongst many grey objects. The *saliency map* was developed as a two dimensional grey scale image which identifies the salient objects, with white objects being the most salient.

Saliency maps have been applied to computer graphics to accelerate global illumination. In particular, Yee et al. [YPG01] presented a framework for dynamic environments to guide where less rendering effort should be spent. This combined a spatiotemporal error tolerance map with the saliency map. Although saliency maps are being used more frequently, Marmitt and Duchowski showed that such maps do not always predict attention regions in a reliable manner in a virtual environment [MD02].

2.3. The top-down process

The top-down processes is the visual phenomenon that volitional control selects attentional field. For example, people will focus on the road ahead when they play driving games. Visual psychology researchers Mack and Rock [MR98] introduced the concept of Inattentional Blindness, showing that people fail to notice significant objects in a scene if they are not being attended to. This top-down process has been exploited in computer graphics by Cater et al. [CCL02, CCW03] and combined with the bottom-up process by Sundstedt et al. [SDC04] to produce, what they termed a task-importance map. The task-importance map created a ranking in which pixels should be rendered in order of their visual importance. Selective rendering is then used to render this ranked list of pixels at the different quality levels, with high priority pixels being rendered with the highest fidelity. The results from their psychophysical experiments confirmed that participants were unable to notice these different quality levels within the selectively rendered animations.

2.4. Perception related LOD techniques

Perceptually based LOD techniques consider the human visual system as an important factor in LOD management/selection. A number of researchers have exploited human vision knowledge to develop level of detail frameworks to remove imperceptible details from models and speed up the rendering process. Reddy [Red97] was amongst the first to attempt a LOD system directly based on the Contrast Sensitivity Function (CSF) model. At run-time the system, guided by the CSF model, selected the highest resolution LOD such that the spatial change it would induce was below the user's threshold frequency for the specific velocity and eccentricity. Reddy recently redeveloped his earlier perceptual LOD system, describing a view-dependent real time LOD system for rendering dense terrain meshes using these perceptual models [Red01].

Scoggins et al. [SMM00] applied CSF as a modulation function after transforming a pre-rendered reference image to frequency space. In addition they decided which LOD is appropriate according to mean-square error results. Brown et al. [BCP03] described a method of visual attention-biased LOD for geometry meshes which extended the polygon LOD management techniques with the bottom-up visual process models for real-time rendering systems.

Luebke et al. [LH01] developed a view-dependent LOD system augmented with the CSF perceptual model and with silhouette preservation. They evaluated vertex merge operations according to the worst-case contrast with the spatial frequency. Merge operations were only performed when the resulting image was predicted to be imperceptible from the original. In addition, this system also took silhouette testing into account during the merge operation, preserving the silhouettes.

Howlett et al. [HHO04] determined the saliency of objects through an eye-tracker device and preserved the prominent features using a modified QSlim software. Their results showed that saliency based simplifications work for non-familiar natural objects as well as familiar ones, but not for man-made artifacts.

This previous research has concentrated on the LOD management of objects individually. In our work LOD techniques are considered in the concept of entire scene.

3. Our approach

We investigated the benefits that could be gained from combining knowledge of the human top-down visual process into LOD management techniques. Previous research strongly suggested that performing a visual task does indeed have a more significant effect on the viewer's gaze pattern across the scene than when performing a non-visual task [CCL02]. A precomputed animation was used for our psychophysical investigation. This comprised a fly-through of four rooms, and the visual task required subjects to count the number of pencils in the pot in each room. The animations were modelled and rendered in Maya. A single frame from the animation is shown in Figure 1 with the task region highlighted.

Figure 1: *The visual task has been highlighted by red circles.*

3.1. Pilot study

A pilot study was run to ascertain which non-task objects were most suitable for our LOD reduction study. The visual importance of the objects in the scene were obtained from the saliency map of the animation, Figure 2. Subsequently, a validation experiment asked 10 participants to free view the animation and recall the objects they remembered. This experiment was able to confirm whether the saliency map predicted the visually important objects correctly. Table 1 summarizes this validation study. As the table shows, the pot and paintings, including the nearby light on the wall, are the most visually important objects; whilst the chair and teapot on the corner of the room showed similar visual importance as the objects on the table. This coincided with the saliency map.

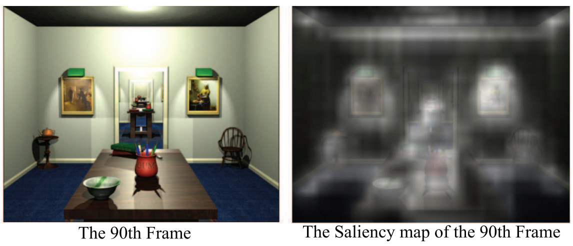

Figure 2: *The figure illustrates the 90th frame of our animation and the saliency map of this frame. The brighter area on the map indicates the regions of higher visual importance. As shown above, the more salient areas of this scene are on the table and pictures on the wall.*

Using the results of this pilot study, we reduced the level of detail on those less salient non-task objects, i.e. teapots, chairs and objects on the table apart from the pencils and associated pots.

Objects	Number of subjects who noticed
Paintings	8
Light/Teapots/Chair/Books	4
Pots/Jars	4
Mug(task)	6

Table 1: *Validation experiment results*

3.2. Psychophysical Experiments

The main experiment used four animations identical apart from different LOD on the non-task objects. The LOD on these non-task objects was reduced from Full LOD to 50% LOD and to 20% LOD respectively, Figure 3. Here, the term 50% LOD means that the number of polygons of the objects is 50% of the Full LOD ones. For example, if the Full LOD of a teapot has 992 polygons 50% LOD will contain just 496 polygons. This was achieved using the standard LOD command in Maya. We deliberately did not consider more perceptually based LOD techniques, such as Luebke's [LH01], as we wished to investigate whether, even with such a crude LOD method, participants would still fail to notice the quality difference.

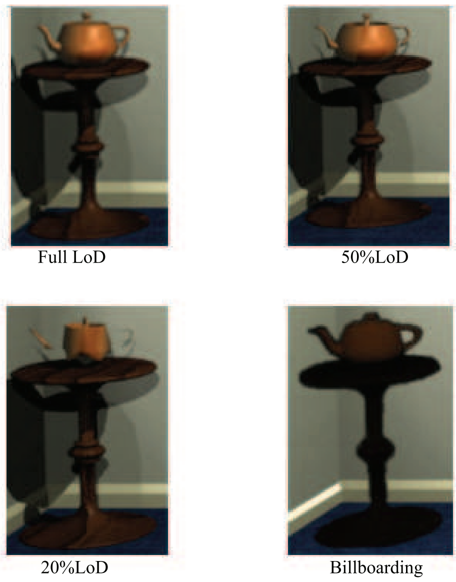

Figure 3: *Shows the definition of different LOD.*

In addition, we also considered billboarding for the non-task related objects as a drastic means of reducing the number of polygons by merely considering the objects' representation using a simple plane. Here we used the images for the billboard acquired from Snapshot, a rapid preview of scene content [LDC05], shown in Figure 4. Each frame for the Full

LOD animation, which comprised 800 frames, took on average 11 minutes 6 seconds to render in Maya on an Intel Pentium IV 1 GHz process, while the frames for the 50% LOD animation and 20% LOD animation were each rendered on average in 8 minutes 58 seconds and 7 minutes 46 seconds respectively.

Figure 4: *Snapshot at the 90th frame.*

Animations with different LOD were tested with 60 participants in total and all participants exhibited at least average corrected eyesight. To reduce the bias of the experiment running conditions, 10 experiments were run in the morning and 10 experiments in the afternoon. The experiments were spread over 3 days.

Before beginning the experiment the subjects read a sheet of instructions on the procedure of the task they were to perform. After the participant had read the instructions they were asked to clarify that they had understood the task. They then rest their head on a chin rest that was located 60cm away from a 17-inch monitor. The chin rest was positioned so that their eye level was approximately level with the centre of the screen. The animation was displayed at a resolution of 1280x1024.

To investigate the effect of inattentional blindness on our LOD related experiments, the animations were grouped into 3 pairs: Full LOD with 50% LOD, Full LOD with 20% LOD and Full LOD with billboard LOD. The subjects were asked to watch both animations and to count the number of pencils in the pot in each room. A countdown was shown before the start of each animation. On completion of the experiment, each participant was asked to fill in a detailed questionnaire. This questionnaire asked whether participants perceived any visual differences between the two animations. They were also shown two pictures of the rendered image of the 90th frame, and were asked to identify which animation the picture belonged to. The subjects were required to support their choice with a detailed explanation.

3.3. Results

To obtain a baseline with which to compare the experimental results, initially we tested 10 subjects with 10 pairs of Full LOD vs. Full LOD animations. This was designed to detect whether they were able to perceive any differences over two identical animations. As expected, no difference was shown in the results, that is, 100% of the subjects did not perceive any difference.

Along with visual task related experiments, we also conducted free view experiments in which participants were required to simply watch two animations. These animation pairs were the same as those used for the visual task related experiments. To minimize the experimental bias, different subjects were chosen for the free viewing from those who performed the visual task related experiments. These results were necessary when analyzing the effect of inattentional blindness.

Figure 5 shows the results of the experiments. During the counting pencils task, 95% of subjects failed to spot differences between 50% LOD and Full LOD, 60% failed to notice between 20% LOD and Full LOD, while only 40% failed to notice when billboarding was used. When simply watching the animations, the free viewing experiment, 50% failed to notice the difference between the 50% LOD and the Full LOD, 15% between 20% LOD and Full LOD, while 30% failed to notice the difference between the animation using full LOD and that using the billboarding technique.

A detailed statistical analysis was also run on the data to determine whether the experimental results do indeed show that Inattentional Blindness significantly affects a viewer's ability to detect LOD differences.

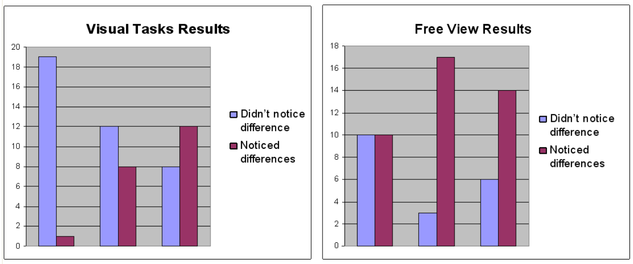

Figure 5: *Visual Task results and Free View results.*

Figure 6: *Shadows not present in the Billboarding LOD*

level of detail	$p-value$
50% LOD	0.02
20% LOD	0.04
Billboarding	0.37

Table 2: *Statistical Analysis*

3.4. Statistical Analysis

The Table 2 shows the statistical analysis of the results when comparing the task related experiment with the free viewing for each animation pair. As the experimental data was categorical, a Chi-square analysis was used. The analysis shows that the results are statistically significant in both the 20% LOD (p-value 0.02) and 50% LOD (p-value 0.04) cases as both values are smaller than the α-value 0.05. Thus we can conclude visual attention can indeed be used together with a LOD management system. For the billboarding experiment, the results were not statistically significant, with a p-value of 0.37. From the questionaires we discovered that the use of billboard, although significantly reducing polygon counts and not altering the perceived shape and size of the objects, does have a major affect on the overall lighting within the environment. Failure to include the objects' geometry results in a lack of shadows which was very noticeable to the participants.

Furthermore, the questionaires also revealed that for the 50% LOD and 20% LOD cases, shadows, shape of objects and lighting changes in the scene were the major reasons that any difference was perceived. Future work needs to consider not only the perception of objects with reduced LOD, but also the effect of the reduction of LOD on the overall lighting, including shadows, within the environment.

4. VDP Validation

To further investigate the perception of the images with reduced LOD, Visual Difference Predictor (VDP) comparisons [Dal93, MPT99] were used to compare the Full LOD images with the 50%, 20% LOD and billboarding images. Figure 7 shows the comparison taken for the 90th frame. The brighter the areas of the image, the larger the perceptual difference. The VDP comparisons confirmed that the degree of visual differences increased as the LOD decreased. In addition, the lack of shadows is perceptually very important.

5. Conclusion and Future Work

The results of our psychophysical experiments show that visual perception, and in particular, inattentional blindness can be used together with a LOD management system. This allows us to substantially reduce the number of polygons, even for those objects which are salient within the environment, without these reductions being noticed by a significant number of viewers. However, care must be taken to ensure the overall lighting within the environment is not adversely affected significantly as this increases the viewer's awareness of quality differences, as shown by the billboarding results.

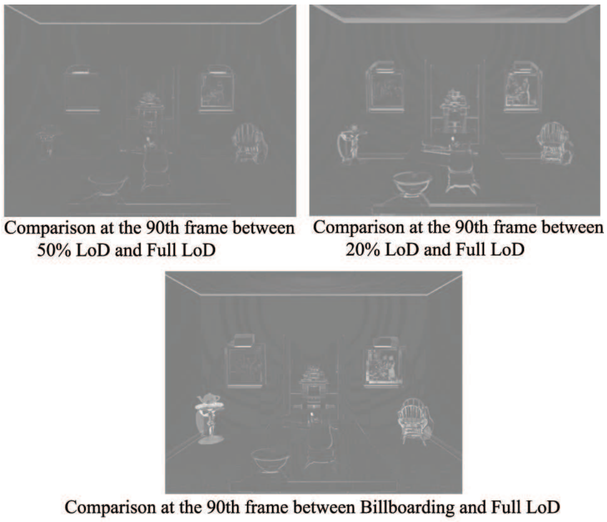

Figure 7: *The VDP comparisons between Full LOD and 50, 20% LOD and Billboarding. The bright area represents the visual differences between the pairs.The greater the differences, the brighter areas appear.*

The next step is to develop an empirical framework of this LOD management technique using task-importance maps in real time in a selective renderer. Although our experiments were rendered with Maya, our primary interest lies with high fidelity rendering. Our future work will incorporate our LOD techniques within the Radiance lighting simulations system [LS98].

LOD saves computation in a ray tracing environment by the lower LODs occupying less voxels in the octree, thus resulting in less ray-object intersection tests during rendering. The computational time saved with our LOD approach in Maya's ray tracer, 11 minutes 6 seconds for the Full LOD compared with 7 minutes 46 seconds for the 20% LOD, was not that significant; a speedup of 1.5. We expect this performance to be improved when considering a full global illumination solution with its significantly higher number of ray object intersection tests. In our experiment, we rendered all the pixels at a high quality. The computational time we saved is thus in addition to any performance improvements we will get from selectively rendering our animations in the future.

6. Acknowledgements

We would like to thank Kurt Debattista for his help and everyone who did the experiment.

References

[BCP03] BROWN R., COOPER L., PHAM B.: 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* (2003), ACM Press, pp. 55–ff.

[CCL02] CATER K., CHALMERS A., LEDDA P.: 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* (2002), ACM Press, pp. 17–24.

[CCW03] CATER K., CHALMERS A., WARD G.: Detail to attention: exploiting visual tasks for selective rendering. In *EGRW '03: Proceedings of the 14th Eurographics workshop on Rendering* (2003), Eurographics Association, pp. 270–280.

[Dal93] DALY S.: The visible differences predictor: An algorithm for the assessment of image fidelity, 1993.

[HHO04] HOWLETT S., HAMILL J., O'SULLIVAN C.: An experimental approach to predicting saliency for simplified polygonal models. In *APGV '04: Proceedings of the 1st Symposium on Applied perception in graphics and visualization* (2004), ACM Press, pp. 57–64.

[IKN98] ITTI L., KOCH C., NIEBUR E.: A model of saliency-based visual attention for rapid scene analysis. In *Pattern Analysis and Machine Intelligence* (1998), vol. 20, pp. 1254–1259.

[Jam90] JAMES W.: A saliency-based search mechanism for overt and covert shifts of visual attention. In *Principles of Psychology* (1890).

[LDC05] LONGHURST P., DEBATTISTA K., CHALMERS A.: Snapshot: A rapid technique for driving a selective global illumination renderer. In *WSCG 2005 SHORT papers proceedings* (2005).

[LH01] LUEBKE D. P., HALLEN B.: Perceptually-driven simplification for interactive rendering. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques* (2001), Springer-Verlag, pp. 223–234.

[LS98] LARSON G. W., SHAKESPEARE R.: *Rendering with radiance: the art and science of lighting visualization*. Morgan Kaufmann Publishers Inc., 1998.

[MD02] MARMITT G., DUCHOWSKI A.: Modeling Visual Attention in VR: Measuring the Accuracy of Predicted Scanpaths. In *Eurographics 2002, Short Presentations* (2002), pp. 217–226.

[MPT99] MYSZKOWSKI K., PRZEMYSLAW R., TAWARA T.: Perceptually-informed Accelerated Rendering of High Quality Walktrough Sequences. In *Eurographics Workshop on Rendering* (1999), pp. 13–26.

[MR98] MACK A., ROCK I.: *Inattentional Blindness*. MIT Press, 1998.

[Red97] REDDY M.: *Perceptually Modulated Levle of Detail for Virtual Enviroments*. PhD thesis, University of Edinburgh,Edinburgh,Scotland, 1997.

[Red01] REDDY M.: Perceptually optimized 3d graphics. *IEEE Comput. Graph. Appl. 21*, 5 (2001), 68–75.

[SDC04] SUNDSTEDT V., DEBATTISTA K., CHALMERS A.: Selective rendering using task-importance maps. In *APGV '04: Proceedings of the 1st Symposium on Applied perception in graphics and visualization* (2004), ACM Press, pp. 175–175.

[SMM00] SCOGGINS R. K., MOORHEAD R. J., MACHIRAJU R.: Enabling level-of-detail matching for exterior scene synthesis. In *VIS '00: Proceedings of the conference on Visualization '00* (2000), IEEE Computer Society Press, pp. 171–178.

[Yan96] YANTIS S.: Attentional Capture in Vision. In *Converging operations in the study of selective visual attention* (1996), Washington, DC: American Psychological Association, pp. 45–76.

[YPG01] YEE H., PATTANAIK S., GREENBERG D. P.: Spatiotemporal sensitivity and visual attention for efficient rendering of dynamic environments. *ACM Trans. Graph. 20*, 1 (2001), 39–65.