PERCEPTUALLY-DRIVEN TERMINATION FOR STOCHASTIC RADIOSITY

Jan Přikryl and Werner Purgathofer
Institute of Computer Graphics
Vienna University of Technology
Karlsplatz 13/186/2
A-1040 Wien
Austria
e-mail: {prikryl,wp}@cg.tuwien.ac.at

ABSTRACT

In this paper we present a heuristic perceptually-based termination criterion for a stochastic radiosity algorithms. The proposed criterion makes it possible to terminate the iterative radiosity simulation as soon as any further changes of the radiosity solution are predicted not to be noticed by the human observer. We use tone mapping operators and perceptually uniform CIE $L^*u^*v^*$ colour space to arrive at the difference of displayed values. In order to decide when to stop the iterations of the radiosity method, we evaluate the least squares fit of a power function to the error distribution data.

Keywords: human perception, radiosity, termination prediction.

1 INTRODUCTION

Usual implementations of the radiosity method use criteria based on radiometric values to drive the computation — to decide about sufficient mesh quality or to estimate the error of the simulation process and to decide when the simulation can be safely terminated. This is absolutely correct for the case of radiometric simulation, when the user is interested in actual values of radiometric quantities. On the other hand, the radiosity method is very often just used to generate pictures for the human observer and those pictures are not required to be correct physical simulations, they just have to look the same.

The advantage of producing "just" a visually satisfying solution that is indistinguishable from the physically correct one is that such a solution is very often much cheaper to compute. Exploiting the fact that current display devices can not by large reproduce the real world range of luminances, or knowing which difference in contrast or colour can be noticed by a human under given viewing conditions, the time needed for a radiosity simulation can be reasonably decreased.

Embedding human perception-aware techniques into radiosity algorithms is still rather rare. Just recently there appeared several papers that use human perception-aware techniques to guide the computations. Gibson and Hubbold [Gibso97] presented a perceptually driven oracle function for hierarchical radiosity and used another perceptual criterion to estimate the necessary accuracy of the form-factor refinement in hierarchical radiosity. Hedley *et al.* [Hedle97] used an approach similar to that of Gibson and Hubbold to cull perceptually unimportant discontinuities in the discontinuity mesh keeping the mesh complexity reasonably low. Martin *et al.* [Marti97] developed another criterion for view potential driven hierarchical radiosity, that evaluates continuity of the Gouraud shading in the reproduced image. Myszkowski [Myszk98] used the *Visible Differences Predictor* proposed by Daly [Daly93] to compare the convergence of different radiosity approaches. He also presented a heuristic termination criterion based on image comparison.

The perceptually-based termination criterion presented by Myszkowski requires an image to be generated and a pair of images to be evaluated after every iteration step of the radiosity algorithm, which is appropriate for view-dependent computations only. When computing a view-independent radiosity solution, Myszkowski suggests using a set of representative views and generating images corresponding to every view after every iteration. The corresponding image pairs are then evaluated in order to predict the convergence. This approach has one obvious disadvantage — the process of the representative views generation is not trivial. Obviously, this set should be as small as possible, but yet producing enough information in order to predict the global perceptual convergence correctly. We feel that it may be often impossible to fulfill both these conditions at the same time.

In the following sections we will describe another heuristic approach that does not require any images to be generated and works directly on radiosity values of a stochastic radiosity method with constant elements.

The paper is further organized as follows: Tone mapping operators, image mapping methods on that perceptually-driven radiosity algorithms are based, are introduced in the Section 2. The Section 3 gives a very brief overview of the perceptual techniques that we may use to evaluate the perceptual convergence. The proposed object-space perceptually driven criterion is described in Section 4 and is is followed by analysis of our experiment in Section 5. Finally, in Section 6 we summarize the paper and present some direction for the future research.

2 TONE MAPPING OPERATORS

The physical accuracy in rendering itself does not yet guarantee that the displayed image will have a realistic visual appearance. Typical CRTs can reproduce luminances between approximately 1 and 100 $\mathrm{cd/m^2}$ and typical print reproduction materials (film slides, paper) even less, whereas the range of luminances that the human visual system is able to perceive can vary between 10^{-5} to 10^5 $\mathrm{cd/m^2}$. Having computed an image containing luminance values that appear in the real world, the last step of the image synthesis process is to display this image as perceptually correct as possible using a device that is unable to reproduce the broad real world luminance scale.

Simulating the process of human visual system adaptation to luminances observed in the real world we can achieve that the subjective brightness and contrast that occur in a real environment are reproduced on a display device, even if the luminances themselves cannot. This process is carried out by *tone mapping operators*. We will describe some of them in following subsections.

2.1 Brightness preserving mapping

Tumblin and Rushmeier's model [Tumbl93] uses results obtained by Stevens and Stevens regarding the brightness associated with a luminance at a particular adaptation level. The tone reproduction is considered for gray-scale images only, since in this environment the spectral radiance is uniform for all wavelengths and the luminance is therefore just a constant times the uniform spectral radiance. Another reason for this limitation is that for each gray input value the display device gives a unique output luminance, which is not true for colour systems where many different colours may have the same luminance.

The whole operator consists of three transformations: First, the real world luminance is transformed to the true real world brightness. Then, using inverse display observer and inverse display device transformations the value of display input that generates the desired real world response is obtained.

2.2 Contrast preserving mapping

A simpler model was proposed by Ward [Ward94] who tried to find a mapping function that would exhibit similar behaviour as the mapping of Tumblin and Rushmeier [Tumbl93], but that would be less computationally demanding. His model works with gray-scale values as well and it attempts to match luminance differences that are just noticeable on the display device with those that are just noticeable in the real world, producing a linear mapping function. This goal is accomplished by finding a constant of proportionality between display luminance and world luminance maintaining roughly the same contrast visibility on the display device and in the real scene where the image has been obtained. This means that Ward's tone mapping operator tries to preserve the perceived contrast, as opposed to the Tumblin-Rushmeier's tone mapping operator that preserves perceived brightness of an image.

2.3 Mapping of Ferwerda et al.

Ferwerda *et al.* [Ferwe96] proposed another tone reproduction technique also based on the concept of matching just noticeable luminance differences for a variety of adaptation levels. Their model captures the changes in threshold visibility, colour appearance, visual acuity and time-course of adaptation of the human visual system.

As their work is based on experimental threshold detection data, they based their model on Ward's tone mapping operator. Their extension consists of taking into account not only the cone characteristics (as Ward did), but also the rod threshold data. Moreover, their model also counts for the whole range of viewing conditions from scotopic over mesopic to photopic level. Using the adaptation dependent visual acuity function they are able to remove those image details that cannot be noticed by the human observer at a given adaptation level. Also, using the data measured for the time-course of the light and dark adaptation they can reproduce images so that they correspond to what the human observer would have seen when entering a dark or bright environment.

2.4 Visibility preserving mapping

Larson *et al.* [Larso97] present a technique that combines advantages of the previous ideas allowing us for correct display of scenes with wide dynamic luminance ranges. The mapping operator is uniform and tries to preserve the visibility of objects in the image.

The task is achieved by using more then a single view adaptation level, which is in fact coherent with the mechanism of human visual system adaptation. The authors noticed that image luminance levels tend to be clusterized rather than to be uniformly distributed throughout the whole dynamic range of the picture. They state that as long as the bright areas are displayed with higher luminances than the dim areas, the absolute difference in the luminance is not so important.

In this operator an image luminance histogram is computed first, followed by computation of the cumulative distribution function of all local adaptation luminances found in the image and clustering of similar adaptation levels. The adaptation luminances are computed for commonly used 1° field of view.

The histogram is then adjusted in order to minimize the visible contrast distortions. The clusters are mapped to the display values preserving local contrast of the cluster. The method uses also knowledge about veiling luminance (*i.e.* about glare effects), colour sensitivity and visual acuity to count for imperfections of the human vision.

2.5 Model of colour adaptation and spatial vision

The most recent work of Pattanaik *et al.* [Patta98] presents a computational model of adaptation and spatial vision for realistic colour tone reproduction. Similarly to the visibility preserving operator of Larson *et al.*, this model allows to reproduce high dynamic range images, but it is not limited to gray-scale images. The operator models the multiscale nature of the human visual perception and is able to account for changes in threshold visibility, visual acuity, color discriminability, and suprathreshold brightness, colourfulness, and apparent contrast that occur with changes in scene illumination. The displayed images match our perceptions of the scenes at both threshold and suprathreshold levels to the degree possible given a particular display device.

3 PERCEPTUALLY-DRIVEN RADIOSITY

In perceptually-driven radiosity algorithms mentioned in the Section 1, tone mapping operators are used to perform any radiosity-based decisions not on real world values, but rather on the displayed ones. This makes it possible to account only for those changes that are going to be visible when the radiosity solution will be displayed. The more detailed overview of this area may be found *e.g.* in [Při98].

3.1 Tone mapping in a radiosity algorithm

If a tone mapping operator is to be embedded into a radiosity algorithm, a severe problem arises: In order to perform the tone mapping, we have to determine the illumination level in the modelled environment that the observer's visual system will adapt to. This *real world adaption level* influences what the observer will actually see in the scene.

When tone-mapping an image, the real world adaptation level can be estimated from image pixels values. In our case, however, even if we would want to generate an image (which would make the whole algorithm view-dependent), there are no correct illumination values available at the beginning of the computation. Gibson and Hubbold [Gibso97] have solved this problem by estimating the ambient illumination of the scene. The adaptation level is then computed as a logarithmic area weighted average of estimated surface luminances, adjusted so that brightnesses are kept around the brightness constancy level of 8.4 dB.

Modern tone mapping approaches are rather complex and process usually the whole image at once in order to better model behaviour of the human visual system. This implies that not every tone mapping operator is suitable for embedding directly into a radiosity method — only mappings of Tumblin and Rushmeier [Tumbl93] and Ward [Ward94] can be used directly, the model of Ferwerda et al. [Ferwe96] can be used only partially, as the color discrimination process consists in fact of low-pass image filtering.

3.2 Assumptions of perceptually-driven computation

Perceptually-driven radiosity computations are usually based on the following three assumptions:

First, computations will be carried out in a well defined colour system (usually CIE XYZ) and not in RGB colour space, as RGB is not precisely defined and actual colour values depend on used CRT phosphors or inks.

Second, the algorithm uses a perceptually-based tone mapping operator to map the real world radiosity to display luminances before doing any further evaluation. This makes it possible to directly evaluate changes in the displayed radiosity values.

Last, displayed radiosity changes are computed using

CIE colour difference formula

$$\Delta E^*_{\mathrm{uv}} = \sqrt{(L^*_1 - L^*_2)^2 + (u^*_1 - u^*_2)^2 + (v^*_1 - v^*_2)^2} \quad (1)$$

which works in the perceptually uniform CIE $L^*u^*v^*$ colour space. The nonuniformity factor of CIE XYZ is approximately 80:1, which means that the same XYZ colour differences may be sometimes perceived as being 80 times different, dependent on a colour pair being compared. CIE $L^*u^*v^*$ perceptual nonuniformity is only about 6:1.

4 PERCEPTUALLY-DRIVEN TERMINATION CRITERION

In the scope of stochastic radiosity simulation we can use common statistic analysis criteria to analyze the convergence and variance of the solution and to stop iterations in the moment when we have reached the result at the given confidence level. However, these variance based criteria are too pessimistic for the purpose of perceptually-based radiosity solution.

During our experiments with the hierarchical well-distributed ray set radiosity algorithm [Bekae98] we have noticed that the maximum change in display luminances drops off with $O(1/k)$, where k is the number of iterations (in our case, one *iteration* means shooting the power from the set of all patches in the scene). This led us to the following heuristics:

After every iteration k we determine the maximum change of displayed luminances $\Delta E^{*(k)}_{\mathrm{uv,max}}$ using tone mapping operator and comparisons in CIE $L^*u^*v^*$ colour space as it has been discussed in Section 3. After that, the least squares functional fit to the collected data is performed in order to arrive at function $\mathcal{F}_k(i)$ that predicts the maximum change for given iteration i. Knowing the maximum unnoticeable luminance difference $\Delta E^*_{\mathrm{uv,t}}$, the last iteration number k_{t} can be predicted using $\mathcal{F}_k^{-1}(\Delta E^*_{\mathrm{uv,t}})$. This process is written down in Algorithm 1.

Algorithm 1 Termination criterion.

repeat
 Perform k-th HQMC iteration
 Determine $\Delta E^{*(k)}_{\mathrm{uv,max}}$
 Approximate $\Delta E^{*(i)}_{\mathrm{uv,max}} = \mathcal{F}_k(i), i \geq 2$
 $k_{\mathrm{t}} \leftarrow \lceil \mathcal{F}_k^{-1}(\Delta E^*_{\mathrm{uv,t}}) \rceil$
until $k > k_{\mathrm{t}}$

The value of $\Delta E^{*(k)}_{\mathrm{uv,max}}$ is chosen as the maximum over all elements of the changes in their displayed radiosities. We do not have to compute the vertex radiosities in every step in this case — when the radiosity over an element is represented using a constant basis function, the radiosity change at a vertex of the element will be an average from the radiosity changes over N elements that coincide at the vertex:

$$\Delta B_{\mathrm{v}} = \frac{1}{N}\sum_{i=1}^{N} \Delta B_i.$$

This implies that for $\Delta B_i \leq \Delta_{\max}$ we will always have $\Delta B_{\mathrm{v}} \leq \Delta_{\max}$ and we can use the maximum element radiosity change as a possibly pessimistic prediction of the maximum vertex radiosity change.

In order to compute the actual change in displayed radiosity for the given element, tone mapped radiosities in CIE XYZ colour space at iterations k and $k-1$ are computed first, obtaining the values that will be reproduced on the given display device. The display values are then converted to CIE $L^*u^*v^*$ colour space and the difference $\Delta E^{*(k)}_{\mathrm{uv,max}}$ is obtained using the colour difference formula (1). Algorithm 2 describes this process in more detail.

Algorithm 2 Determining $\Delta E^{*(k)}_{\mathrm{uv,max}}$

$\Delta E^{*(k)}_{\mathrm{uv,max}} \leftarrow 0$
for all $i \in$element_indices **do**
 $B^{(k-1)}_{\mathrm{d,xyz}} \leftarrow \mathrm{tonemap}[P_i^{(k-1)}/(\pi A_i)]$
 $B^{(k)}_{\mathrm{d,xyz}} \leftarrow \mathrm{tonemap}[P_i^{(k)}/(\pi A_i)]$
 $B^{(k-1)}_{\mathrm{d,luv}} \leftarrow \mathrm{xyz\_to\_luv}(B^{(k-1)}_{\mathrm{d}})$
 $B^{(k)}_{\mathrm{d,luv}} \leftarrow \mathrm{xyz\_to\_luv}(B^{(k)}_{\mathrm{d}})$
 $\Delta E^*_{\mathrm{uv}} \leftarrow \mathrm{luvdiff}(B^{(k)}_{\mathrm{d,luv}}, B^{(k-1)}_{\mathrm{d,luv}})$
 $\Delta E^{*(k)}_{\mathrm{uv,max}} \leftarrow \max\{\Delta E^{*(k)}_{\mathrm{uv,max}}, \Delta E^*_{\mathrm{uv}}\}$
end for

The function $\mathcal{F}_k(i)$ is approximated with the power function

$$\Delta^{*(i)}_{\max} = \mathcal{F}_k(i) = a_k \cdot i^{b_k}, \quad (2)$$

coefficients a_k and b_k are determined using least squares method:

$$\begin{aligned} b_k &= \frac{k\sum_n(\ln n \ln \Delta^{*(n)}_{\max}) - \sum_n \ln n \sum_n \ln \Delta^{*(n)}_{\max}}{k\sum_n(\ln n)^2 - \left(\sum_n \ln n\right)^2}, \\ a_k &= \frac{\sum_n \ln \Delta^{*(n)}_{\max} - b\sum_n \ln n}{k}, \end{aligned}$$

where n goes from 2 to k. The sums can be updated in every iteration without need to compute them from the scratch every time.

The iteration number k_{t} for which the approximated value $\Delta E^{*(k_{\mathrm{t}})}_{\mathrm{uv,max}}$ will stay below the given threshold $\Delta E^*_{\mathrm{uv,t}}$ can be then determined as

$$k_{\mathrm{t}} = \left\lceil \exp\left(\frac{\ln \Delta E^*_{\mathrm{uv,t}} - \ln a_k}{b_k}\right)\right\rceil. \quad (3)$$

The whole process of determining k_{t} with incremental updates of the sums is described in Algorithm 3.

Algorithm 3 Determining k_t.

if $k = 1$ **then**
  $s_k, s_\Delta, s_{kk}, s_{k\Delta} \leftarrow 0$
  $k_t \leftarrow$ invalid
else
  Compute $\Delta E^{*(k)}_{uv,max}$
  $s_k \leftarrow s_k + \ln k$
  $s_\Delta \leftarrow s_\Delta + \ln \Delta E^{*(k)}_{uv,max}$
  $s_{kk} \leftarrow s_{kk} + (\ln k)^2$
  $s_{k\Delta} \leftarrow s_{k\Delta} + \ln k \cdot \ln \Delta^{*(k)}_{max}$
  $b \leftarrow (ks_{k\Delta} - s_k s_\Delta)/(ks_{kk} - (s_k)^2)$
  $a \leftarrow \exp\left[\,(s_\Delta - bs_k)/k\,\right]$
  $k_t \leftarrow \lceil (\exp\left[\,(\ln \Delta^*_t - \ln a)/b\,\right] \rceil$
end if

5 RESULTS

In experiments presented in this section we used $\Delta E^*_{uv,t} = 2.0$, which correspons roughly to just perceivable luminance differnce of 2% measured in laboratory conditions. In the complex environment this difference is likely to be higher. Presented results have been obtained using Ward's tone mapping operator [Ward94]. Tumblin and Rushmeiers' operator [Tumbl93] produces similar results.

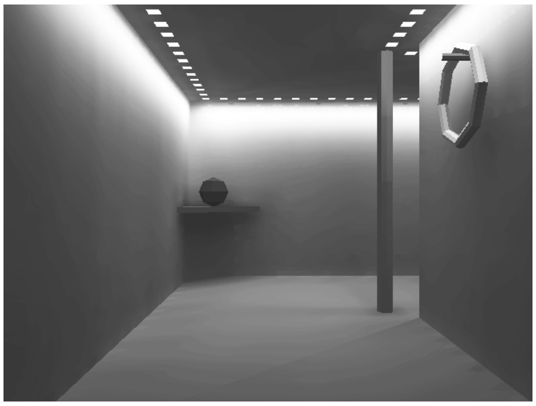

Figure 1: Corridor scene from the RenderPark distribution.

Rays per patch	$k_t(1)$	$k_t(2)$
10	577	118
100	91	17
1000	22	5

Table 1: Final iterations for two test scenes.

The results are presented for two different scenes shown in Figures 1 and 2. For both test scenes we computed three solutions with 10, 100, and 1000 rays shooted from every scene patch. In both presented cases the setup with 10 rays per patch produces solutions with high initial variance of patch radiosities, which results in high variance of $\Delta E^{*(k)}_{uv,max}$. During the

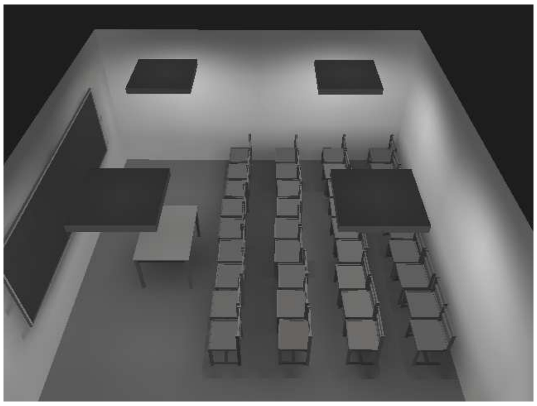

Figure 2: Scene #8 from Peter Shirley's package of test scenes for radiosity.

computation we recorded the values of k_t, $\Delta E^*_{uv,max}$, and approximation curve coefficients. The comoputation was terminated when $k > k_t$. Table 1 lists last iteration numbers for all six setups.

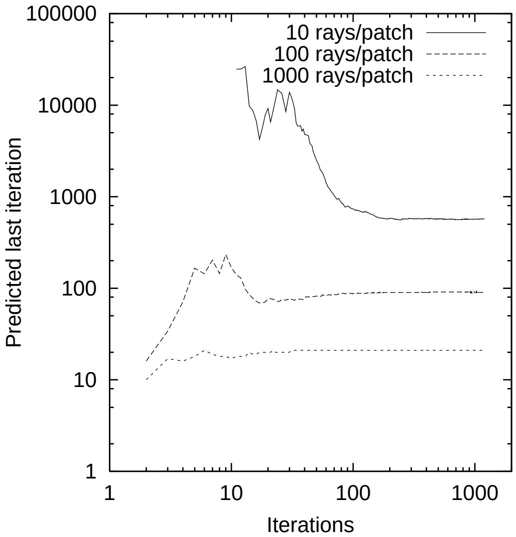

Figure 3: Termination predictions, k_t, for Scene 1.

In Figures 3 and 4 the reader can see how the prediction k_t of the final iteration evolved during the computation for fixed number of iterations (1200 for Scene 1, 200 for Scene 2), ignoring the termination criterion. We can see that after the terminating iteration as lised in Table 1 the further predictions of k_t do not differ much from the k_t value used for termination of the algorithm. The stability of the guess depends on the variance of the $\Delta E^*_{uv,max}$ data, which in turn depends on the number of rays shot from every patch. Our experience shows that even if after some further iteration the value of k_t changes, that changes are minimal (typically up to $k_t - 2$) for the case of decent variance. Unfortunately with the 10 rays per patch setup the prediction may rise up a bit in later iterations (for Scene 1 the highest predicted iteration number is 579 and not 577, for Scene 2 the computation should have

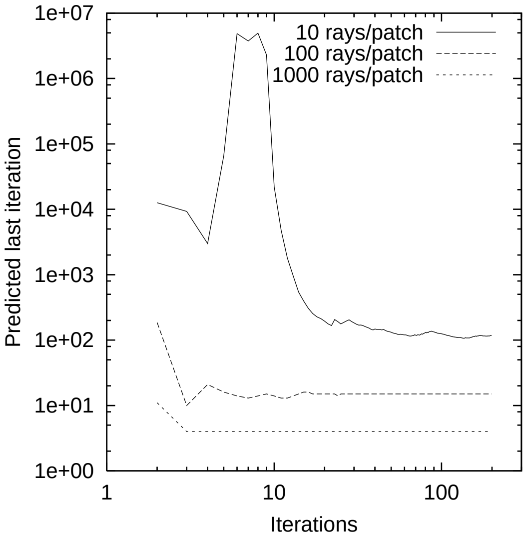

Figure 4: Termination predictions, k_t, for Scene 2.

been stopped at iteration 119 instead of 118). Figure 5 shows how variance of the error data evolved during the computations of Scene 1.

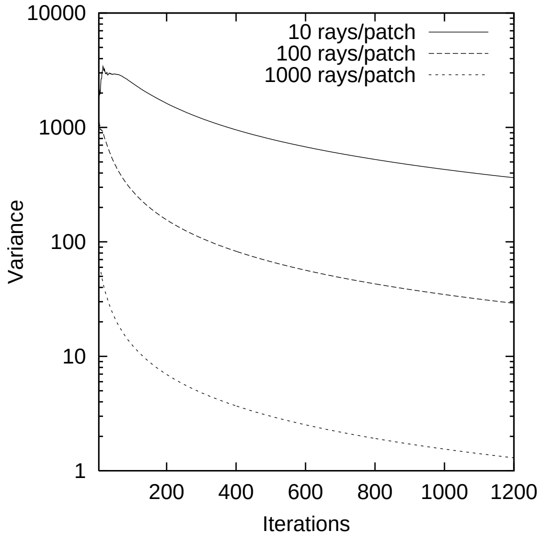

Figure 5: Variance of $\Delta E^{*(k)}_{\mathrm{uv,max}}$ for Scene 1.

Figures 6 and 7 show the measured values of $\Delta E^{*}_{\mathrm{uv,max}}$ and approximating functions $\mathcal{F}_{k_t}(i)$ for all six setups. We can see that the quality of the prediction increases as the result variation drops off (notice the logarithmic scale on the y-axis). The prediction is quite acceptable for computations with 100 and 1000 rays per patch, but for the 10 rays our method tends to overestimate the convergence speed of the algorithm.

Figure 8 shows values of parameters a_k and b_k of the approximating function $\mathcal{F}_{k_t}(i)$ dtermined after every iteration k for Scene 1.

To evaluate usability of our criterion, we performed

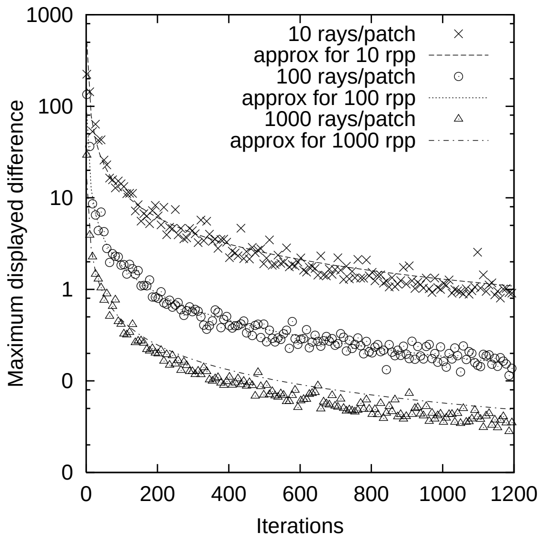

Figure 6: Perceptual convergence for Scene 1. Only every 5-th value of $\Delta E^{*}_{\mathrm{uv,max}}$ is plotted in order to maintain the graph readable.

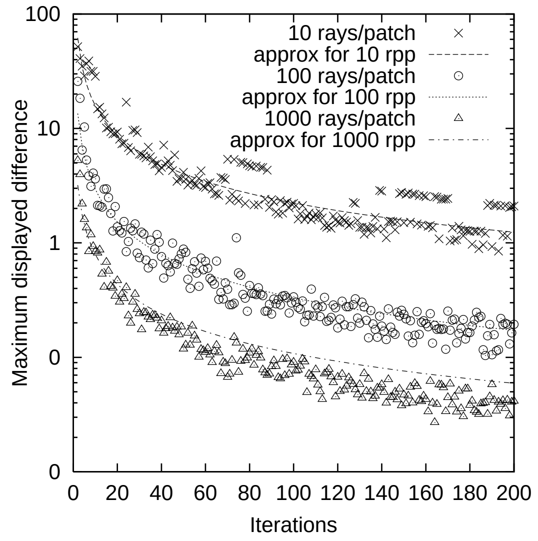

Figure 7: Perceptual convergence for Scene 2.

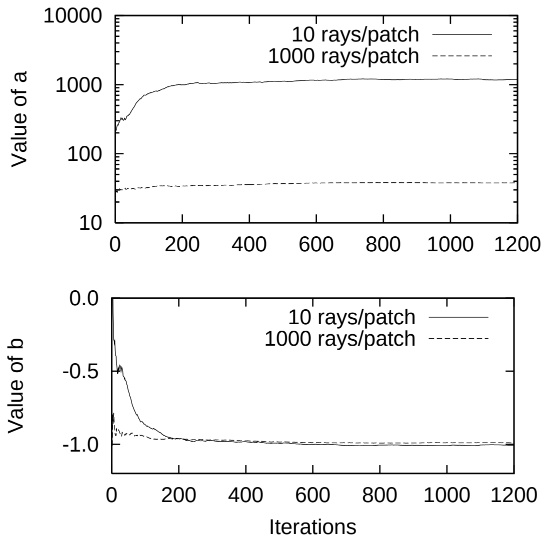

Figure 8: Values of regression curve paramters a_k and b_k for Scene 1.

a simple test using approach similar to that described by Myszkowski [Myszk98] — after every iteration of the radiosity algorithm we generated an image of the scene in display luminances. We then compared the image produced at k-th iteration with the image obtained at the terminating iteration, *i.e.* with the image that our algorithm claimed to be the perceptually converged one. Instead of using the *visible differences predictor* as Myszkowski did, we have used the perceptual colour image metrics of Neumann *et al.* [Neuma98] to evaluate the quality of the termination criterion. Two images X and Y compared by this metric will be indistinguishable, if the difference lays below the threshold $\delta_t = 1.0$. The metric has been computed for a CRT with diagonal 21 inches, resolution 1280×1024 pixels, and viewing distance 0.5 meter. Compared images were computed with the resolution 640×480 pixels.

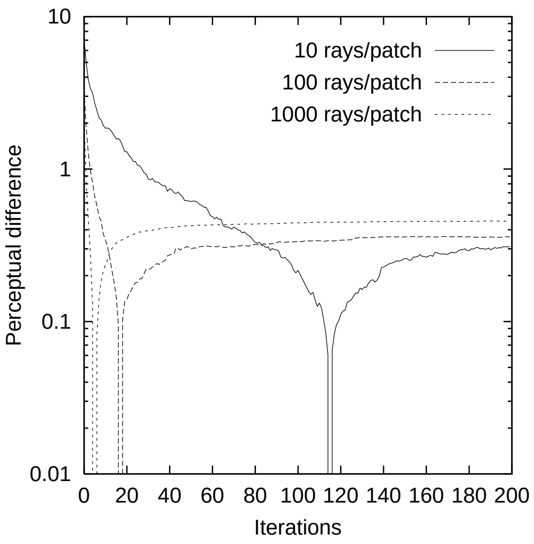

Figure 9: Differences between the results of particular iterations and preceptually converged results for Scene 2.

The graph in Figure 9 shows the reported differences between images generated after every iteration and perceptually converged images for Scene 2. For the final terminating iteration $k_{t,f}$ the graph shows that although there are differences for intermediate solutions where $k > k_{t,f}$, the differrences are always lower then the threshold δ_t and will be therefore not noticed by the human observer. The preceptually converged images have been obtained at iterations listed in Table 1. In the graph we can see corresponding zero differences for images generated at iterations $k_{t,f}$. In Table 2 the image differences between the reference image and perceptually converged images for Scene 2 have been listed. The reference image is in this case the most precise one obtained during our simulation — 200 iterations of HQMC radiosity with 1000 rays per patch. From the presented table we can see that the difference between the perceptually converged and reference image are always below the discernibility level of the human visual system.

Rays per patch	$\Delta(I_{\text{perc}}, I_{\text{ref}})$
10	0.453917
100	0.370932
1000	0.455763

Table 2: Perceptual differences between the perceptually terminated and reference solutions.

If we compare results presented in Fig. 7 with the results of our comparison experiment presented in Fig. 9, you will notice that according to the image comparison we could stop the computations much sooner than our termination criterion suggests. The reason for this phenomenon is simple: our metric evaluates the global view independent error, while the image evlauation process takes only the error in the visible part of the scene into account.

6 CONCLUSIONS AND FUTURE WORK

We have presented a heuristic perceptually-driven termination criterion for stochastic radiosity methods. The criterion uses perceptual tone mapping operator and perceptually uniform CIE $L^*u^*v^*$ colour space to compute the maximum change in displayed luminance for the given iteration step. The iteration at which the maximum displayed luminance change lays below a given threshold is then predicted using extrapolation of current data.

The criterion has certainly many drawbacks. The main of them is the fact that computations with high initial variance tend to be terminated too soon and the approximation $\mathcal{F}_{k_t}(i)$ tens to underestimate the value of $k_{t,f}$. One of the possible solutions would be to *e.g.* modify the termination threshold based on the variance of the obtained error data, another possibility would be to filter the $\Delta E^*_{uv,\max}$ values by selecting local maxima over several iterations in order to get an estimate of the upper bound of the error.

We would like to extend the termination criterion with the patch variance analysis approach, terminating the computation when variance over elements of a patch drops below a given thershold. For hierachical radiosity, this approach can be then improved with link variance analysis, possibly developing an algorithm for *perceptual error control* allowing for perceptual error limits similar to radiance error limits presented by Bekaert and Willems [Bekae96] for hierarchical radiosity.

Instead of using the comparison algorithm of Neumann *et al.*, it would be also interesting to compare results that we obtained with those produced by image comparison method based on a human visual system model, *e.g.* that of Daly [Daly93].

Acknowledgements

The above presented ideas have been implemented within the RenderPark software from K.U.Leuven [Ren]. The authors would like to thank to Phillipe Bekaert from K.U.Leuven and Robert F. Tobler from TU Vienna for worthful ideas. This work has been supported by Austrian "Fonds zur Förderung der wissenschaftlichen Forschung" under project number P11545-MAT.

References

[Bekae96] Philippe Bekaert and Yves D. Willems. Error Control for Radiosity. In *Rendering Techniques '96 (Proceedings of the Seventh Eurographics Workshop on Rendering)*, pages 153–164. Eurographics, Springer-Verlag/Wien, 1996.

[Bekae98] Philippe Bekaert, László Neumann, Attila Neumann, and Yves D. Willems. Hierarchical Monte Carlo radiosity. In *Rendering Techniques '98 (Proceedings of the Ninth Eurographics Workshop on Rendering)*, pages 259–268. Eurographics, Springer-Verlag/Wien, 1998.

[Daly93] Scott Daly. *The Visible Differences Predictor: An Algorithm for the Assessment of Image Fidelity*, chapter 14, pages 179–206. Digital Images and Human Vision. The MIT Press, 1993.

[Ferwe96] James A. Ferwerda, Sumanta N. Pattanaik, Peter Shirley, and Donald P. Greenberg. A Model of Visual Adaptation for Realistic Image Synthesis. In *Computer Graphics Proceedings*, Annual Conference Series, pages 249–258. ACM SIGGRAPH, 1996.

[Gibso97] Simon Gibson and R. J. Hubbold. Perceptually Driven Radiosity. *Computer Graphics Forum*, 16(2):129–140, 1997.

[Hedle97] David Hedley, Adam Worall, and Derek Paddon. Selective Culling of Discontinuity Lines. In *Rendering Techniques '97 (Proceedings of the Eigtht Eurographics Workshop on Rendering)*, pages 69–80. Eurographics, Springer-Verlag/Wien, 1997.

[Larso97] Gregory Ward Larson, Holly Rushmeier, and Christine Piatko. A Visibility Matching Tone Reproduction Operator for High Dynamic Range Scenes. *IEEE Transactions on Visualization and Computer Graphics*, 3(4):291–306, October 1997.

[Marti97] Ignacio Martin, Xavier Pueyo, and Dani Tost. An Image-Space Refinement Criterion for Linear Hierarchical Radiosity. In *Graphics Interface '97*, pages 26–36, May 1997.

[Myszk98] Karol Myszkowski. The Visible Differences Predictor: application to global illumination problems. In *Rendering Techniques '98 (Proceedings of the Ninth Eurographics Workshop on Rendering)*, pages 223–236. Eurographics, Springer-Verlag/Wien, 1998.

[Neuma98] László Neumann, Krešimir Matković, and Werner Purgathofer. Perception Based Color Image Difference. In *Eurographics '98 Conference Proceedings*, volume 17 of *Computer Graphics Forum*, pages C–233–C–241. Eurographics, September 1998.

[Patta98] Sumanta N. Pattanaik, James A. Ferwerda, J. A. Fairchild, and Donald P. Greenberg. A Multiscale Model of Adaptation and Spatial Vision for Realistic Image Display. In *Computer Graphics Proceedings*, Annual Conference Series, pages 287–298. ACM SIGGRAPH, 1998.

[Při98] Jan Přikryl and Werner Purgathofer. State of the Art in Perceptually-Driven Radiosity. In *Eurographics 98 State of the Art Reports*, pages 89–104, September 1998. ISSN 1017-4656.

[Ren] RenderPark system. Available from `http://www.cs.kuleuven.ac.be/cwis/research/graphics/RENDERPARK/`.

[Tumbl93] Jack Tumblin and Holly E. Rushmeier. Tone Reproduction for Realistic Images. *IEEE Computer Graphics and Applications*, 13(6):42–48, November 1993.

[Ward94] Greg Ward. *A Contrast-Based Scalefactor for Luminance Display*, volume IV of *Graphics Gems*, chapter VII.2, pages 415–421. Academic Press, 1994.