Adaptive Sampling and Bias Estimation in Path Tracing

Rasmus Tamstorf

Graphical Communication
Technical University of Denmark
DK-2800 Lyngby, Denmark
rt@acm.org

Henrik Wann Jensen

mental images GmbH & Co. KG
Fasanenstrasse 81
D-10623 Berlin, Germany
henrik@mental.com

Abstract : One of the major problems in Monte Carlo based methods for global illumination is noise. This paper investigates adaptive sampling as a method to alleviate the problem. We introduce a new refinement criterion, which takes human perception and limitations of display devices into account by incorporating the tone-operator. Our results indicate that this can lead to a significant reduction in the overall RMS-error, and even more important that noisy spots are eliminated. This leads to a very homogeneous image quality.
As most adaptive sampling techniques our method is biased. In order to investigate the importance of this problem, a nonparametric bootstrap method is presented to estimate the actual bias. This provides a technique for *bias correction* and it shows that the bias is most significant in areas with indirect illumination.

1 Introduction

Monte Carlo based global illumination methods such as path tracing [6] and bidirectional path tracing [10, 9, 17] are very powerful techniques when highly complex scenes with complex reflection models are rendered. The main disadvantage of Monte Carlo based ray tracing is the high variance of the result which is seen as noise in the rendered images. This noise can be eliminated by increasing the number of samples per pixel, but it is costly due to the slow convergence of the Monte Carlo method.

Adaptive image sampling methods tries to avoid the problem of using a fixed (high) number of samples per pixel by concentrating the samples in the difficult parts of the image. Several methods which apply adaptive sampling to regular ray tracing have been presented. The primary aim of most of these methods is to concentrate the samples along the edges in the image to reduce aliasing effects [13, 12]. In Monte Carlo based ray tracing such as path tracing these edges represents a minor problem compared to the noise seen as spots in the image. Only a few adaptive sampling techniques capable of handling this kind of noise have been presented.

Lee et al. [11] presented a method in which each pixel is sampled based on the variance of the samples, and Purgathofer [14] presented a technique in which confidence intervals are used. The use of confidence intervals has the advantage that error bounds can be specified explicitly for each pixel. However, as demonstrated by Kirk and Arvo [8] both these methods are biased. To avoid bias, Kirk and Arvo presented a method in which an initial sample set is used only to estimate the necessary number of samples per pixel.

An alternative to using adaptive sampling is filtering in which the noisy spots are located and removed using for example a median filter [5] or a more complex energy preserving filter [15]. These methods have the advantage that it is relatively cheap to reduce the amount of noise in the picture. The accuracy of the filtered image is however difficult to predict since the blurring of noisy pixels also introduces errors in the image.

In this paper we investigate the use of adaptive sampling with path tracing. We use a confidence interval based approach similar to Purgathofer's, but with a modified formula for computing the confidence interval in which a tone-operator is included. This has the advantage that samples are concentrated in those regions where they contribute most to the final appearance of the image. Like other adaptive sampling techniques our method is biased. In order to see whether this bias is significant we estimate it using a statistical technique known as non-parametric bootstrapping.

2 Adaptive sampling

In this section we will consider adaptive sampling using confidence intervals. In addition to [14] we will present a new refinement criterion which includes the tone operator.

Adaptive sampling using confidence intervals is based on an analysis of the pixel radiance, L. This is computed by integrating the image-function, $p(x, y)$, which represents the radiance towards a viewer for any point (x, y) in the rendered image. The radiance through a pixel, $\bar{L}$, is thus found by :

$$\bar{L} = \int_A p(x, y) f(x, y)\, dx\, dy \tag{1}$$

where $f(x, y)$ is an appropriate anti-aliasing filter[1] and A is the support area of the filter. $p(x, y)$ is unknown and $\bar{L}$ is estimated by a Monte Carlo method based on n random samples, X_i, $i = 1, \dots, n$:

$$\hat{\bar{L}} = \frac{1}{n} \sum_i X_i \tag{2}$$

where the samples are distributed according to the chosen filter kernel.

The idea in [14] is to continue sampling until the confidence that the true value, $\bar{L}$, is within a given tolerance of the estimate has reached a certain level. More formally this means that sampling continues until :

$$\mathrm{P}\{\bar{L} \in [\hat{\bar{L}} - d; \hat{\bar{L}} + d]\} = 1 - \alpha \tag{3}$$

[1]Since the final image contains one value for each pixel, the Nyquist frequency of interest for anti-aliasing is 0.5 pixel^{-1}. The attenuation of the box filter at this frequency is not very good, and it is our experience that a Gaussian filter provides visually superior images.

where $1 - \alpha$ is the confidence level, d is the allowed tolerance, and P denotes the probability of the specified event.

The advantage of this approach is that pixel values are estimated with the same degree of uncertainty and thus that the noise level should be the same all over the image. Considering that one of the major problems with Monte Carlo based rendering is large differences in the noise level across an image, this property is very desirable.

However, in the quest for realism, calculating accurate radiance values is not enough. Human perception must also be taken into account, and so must the limitations of current display devices. While the human eye has an input range in the order of 10^{-5} to 10^{5} cd/m^2 (cf. [16]), typical displays can only show values between 1 and 100 cd/m^2, so a transformation is obviously necessary. In order to deal with these problems, *tone operators* have been investigated in e.g. [16, 1, 18]. These operators take as input a “real world” radiance and returns a corresponding display value.

Thus, instead of displaying $\bar{L}$, the displayed value should be :

$$L = T(\bar{L}) \tag{4}$$

where $T(\cdot)$ is the tone operator.

In order to incorporate this in adaptive sampling, we propose that the refinement criteria is changed so the confidence interval concerns L and not $\bar{L}$. This means that sampling continues until the confidence of the *displayed* value is within a given tolerance. Formally this requires two values, $\hat{L}_L$ and $\hat{L}_U$, to be found, such that :

$$\mathrm{P}\{L \in [\hat{L}_L; \hat{L}_U]\} = 1 - \alpha, \quad \hat{L}_U - \hat{L}_L = 2d \tag{5}$$

where $\hat{L}_L$ and $\hat{L}_U$ are estimated from the available samples.

Compared to Eq. (3) the difference is subtle. However, the important thing is that the amount of work done depends on the displayed value, and that samples will not be wasted at obtaining accurate estimates for small differences which cannot be displayed, or differences which are simply not perceived by the viewer. If for example the display device can display values between 0 and 1 there is no point in using many samples to find out whether the radiance for a given pixel is 6.5 or 7.5 – even though the absolute difference between these values is large.

The drawback of this approach is obviously that an image rendered for one display type should not be used with a different display type. For many practical purposes though, this restriction is not severe, and it is already seen in other renders which include gamma correction.

3 Evaluation of confidence intervals

While a refinement criterion based on confidence intervals is relatively easy to understand in words, the transformation into a criterion which can be computed is not trivial. This section points out why it is not trivial, presents an empirical method to justify the necessary approximation, and presents a the new stopping condition corresponding to Eq. (5).

The reason that the desired confidence intervals are not trivial to estimate, is that the marginal distribution of the samples used to estimate the value of a pixel is unknown. And except for artificially constructed examples it will always be so. Even for the simple mean considered in Eq. (3) this means that approximations must be made. These approximations have not been investigated previously, so before the tone operator is included in the refinement criterion we will give a brief overview.

Given n samples, an approximate $1-\alpha$ confidence interval for $\bar{L}$ is given by (cf. [3]) :

$$[\hat{\bar{L}}_L; \hat{\bar{L}}_U] = \left[\hat{\bar{L}} - t(n-1)_{1-\alpha/2}\sqrt{s^2/n};\ \ \hat{\bar{L}} + t(n-1)_{1-\alpha/2}\sqrt{s^2/n}\right] \tag{6}$$

where $t(n-1)_{1-\alpha/2}$ is the $1-\alpha/2$ quantile of the t–distribution with $n-1$ degrees of freedom, and s^2 is the sample variance of X_i, $i = 1, \ldots, n$.

Unfortunately, when n is finite the coverage probability, $1-\alpha$, will only be correct for $d \rightarrow 0$ and under the assumption that the distribution of X_i is symmetric with respect to the median (cf. [3]). In applications d is typically 5% (cf. [14, 15]) or smaller which is reasonably close to 0, but there is no guarantee that the distribution of X_i is symmetric. As an example, the distribution of samples from a two colored region where one color is more frequent than the other, will not be symmetric with respect to the median. Hence, even as an approximation Eq. (6) may fail. To see whether this is a problem assume that n is large. Since each of the samples are independent and identically distributed, it follows from the central limit theorem that :

$$\hat{\bar{L}} \underset{\text{approx.}}{\in} N(\bar{L}, \sigma^2/n) \tag{7}$$

where σ^2 is the variance of the distribution of X_i. Based on this, Eq. (6) can then easily be rederived, so the question is : How large should n be for this approximation to hold ?

There is no single answer to this question because the distribution of X_i is unknown. However, if it can be justified that $\hat{\bar{L}}$ is normal distributed for a given n, then that n must be large "enough". Consider therefore, the empirical distribution of m independent estimates made with n samples. By comparing this to a normal distribution by means of a statistical χ^2 goodness of fit test [7], the hypothesis that $\hat{\bar{L}}$ is normal distributed can be tested. If it cannot be rejected it must be reasonable to assume that – at least as a working assumption – the approximation is acceptable. An experiment based on this method is reported in section 5, and it is shown that the approximation actually is acceptable for most pixels in a test scene even with a relatively small value of n (i.e. 10 to 20).

Hence, in the following it will be assumed that :

$$\mathrm{P}\{\bar{L} \in [\hat{\bar{L}}_L; \hat{\bar{L}}_U]\} = 1-\alpha \tag{8}$$

where the values of $\hat{\bar{L}}_L$ and $\hat{\bar{L}}_U$ are given in Eq. (6). Assuming furthermore that the tone operator of choice is a non-decreasing function, we get the following stopping condition :

$$\begin{aligned}\mathrm{P}\{\bar{L} \in [\hat{\bar{L}}_L; \hat{\bar{L}}_U]\} &= \mathrm{P}\{T(\bar{L}) \in [T(\hat{\bar{L}}_L); T(\hat{\bar{L}}_U)]\} \\ &= \mathrm{P}\{L \in [T(\hat{\bar{L}}_L); T(\hat{\bar{L}}_U)]\} = 1-\alpha \end{aligned} \tag{9}$$

and combined with Eq. (5) this shows that sampling should continue until :

$$T(\hat{\bar{L}}_U) - T(\hat{\bar{L}}_L) \leq 2d \tag{10}$$

where d is the accepted tolerance and $\hat{\hat{L}}_L$ and $\hat{\hat{L}}_U$ are given by Eq. (6). It should be noted that since this condition has been derived from an approximation it is obvious that it is only an approximation itself. However, for most practical purposes it is a reasonable approximation.

4 Bootstrapping and bias estimation

One of the problems with adaptive sampling is bias (cf. [8]). This means that in average the estimated pixel values will be incorrect. Obviously this is undesirable. Whether it has any practical significance though, is not a question about biased or unbiased, but a question about how large the bias is. If for example a display device with 8 bitplanes per color is used, a bias which is significantly less than $1/256$ will seldomly change any pixel values at all.

Consequently, it is interesting to estimate the bias. To this end we will introduce nonparametric bootstrapping which is a statistical method based on resampling, [2]. For this purpose, let θ be some statistic derived from a distribution F, by the functional relation $\theta = t(F)$, and let $\hat{\theta} = s(\boldsymbol{x})$ be an estimator for θ based on a set of n samples, $\boldsymbol{x}$, from F. With adaptive sampling in mind, F can be thought of as the marginal distribution of image values in a region that contributes to a given pixel, $s(\cdot)$ can be thought of as the adaptive sampling algorithm, and θ corresponds to the radiance, L. Given these definitions, the true bias caused by adaptive sampling is :

$$\mathrm{bias}_F(\hat{\theta}, \theta) = \mathrm{E}_F\{\hat{\theta}\} - \theta = \mathrm{E}_F\{s(\boldsymbol{x})\} - t(F) \tag{11}$$

where $\mathrm{E}_F\{\cdot\}$ denotes expectation with respect to F. Thus, *if* the true value, θ, is known, it is possible to estimate the bias directly, but in most situations this is not the case. Yet, F is not completely unknown. The bootstrap technique exploits this by using the initial n samples to obtain a consistent estimate, $\hat{F}$, of F (for the nonparametric bootstrap this is the empirical distribution[2]). The idea is then to transfer all calculations from the "real world" where F is the underlying and unknown distribution, to the "bootstrap world" where $\hat{F}$ is the underlying but known distribution. Because $\hat{F}$ is known a number of statistics which cannot be calculated in the real world *can* be computed in the bootstrap world, and because $\hat{F}$ is a consistent estimate of F, these statistics will be consistent estimates of the corresponding statistics in the real world.

The actual computations are very simple : F is replaced by $\hat{F}$, such that e.g. samples are drawn from $\hat{F}$ instead of F. When $\hat{F}$ is the empirical distribution of the original samples, this means that the original data are resampled.

Letting $\boldsymbol{x}^*$ denote a bootstrap sample (i.e. a vector of samples from $\hat{F}$), the bootstrap estimate of the bias, becomes :

$$\mathrm{bias}_{\hat{F}}(\hat{\theta}^*, \hat{\theta}) = \mathrm{E}_{\hat{F}}\{s(\boldsymbol{x}^*)\} - t(\hat{F}) \tag{12}$$

where $\hat{\theta}^* = s(\boldsymbol{x}^*)$ is the bootstrap estimate of θ. Since resampling from $\hat{F}$ is inexpensive (compared to sampling from F which requires e.g. path tracing) many bootstrap estimates can easily be obtained for θ. Thus, if $\{\hat{\theta}^*_b\}_{b=1}^B$ are B such estimates, $\mathrm{E}_{\hat{F}}\{s(\boldsymbol{x}^*)\}$

[2]However, the method is not restricted to the nonparametric case.

can be approximated by :

$$\hat{\theta}^*(\cdot) = \sum_{b=1}^{B} \hat{\theta}_b^* / B \tag{13}$$

This leads to the following estimate of bias :

$$\widehat{\text{bias}}_B = \hat{\theta}^*(\cdot) - t(\hat{F}) \tag{14}$$

In a similar manner the variance of $\hat{\theta}$ can be estimated by :

$$\hat{\text{se}}_B^2 = \sum_{b=1}^{B} [\hat{\theta}_b^* - \hat{\theta}^*(\cdot)]^2 / (B-1) \tag{15}$$

The ideal bootstrap estimates are obtained for $B \to \infty$, but for most practical purposes B can be chosen between 25 and 200 for the standard error estimate, and between 400 and 1000 for the bias estimate (cf. [2]).

Bootstrap estimates of bias can theoretically be used for *bias correction*, but the price will often be a higher variance of the final estimate. In images this is seen as noise, so by using it, the advantage obtained by adaptive sampling may be lost, although it need not be the case. This will require further research. The primary goal here has been to provide a method which can be used to characterize the the bias caused by adaptive sampling empirically.

The bootstrap method can also be used to estimate confidence intervals. However, the corresponding theory is fairly involved, and we will not go into details with it. Furthermore, our initial experiments have indicated that the confidence intervals obtained this way are not significantly better than those obtained with the simpler approximation in Eq. (3), and the computational burden is quite large.

5 Results and Discusssion

We have implemented the adaptive sampling scheme presented in this paper for a pure path tracing algorithm using gamma correction ($\gamma = 2.2$ cf. [4]) as the tone operator. The tolerance has been set to $d = 1/256$ corresponding to the color resolution of the final images, and the minimal number of samples has been chosen according to the lower bound, n, given in [14]. If the refinement criteria has not been met after these n samples, another n samples have been taken and so forth. Unless otherwise stated, all images have been calculated using a Gaussian filter for anti-aliasing with 3 dB attenuation at the Nyquist frequency, $1/2$ pixel^{-1}.

The results presented here are based on two test scenes, which have been rendered with 160×120 pixels. A Cornell box model and a glass sphere on a table of wood. The Cornell box is characterized by much indirect illumination and a shiny sphere, while the primary difficulty with the glass sphere is a caustic on the table.

5.1 Approximation of confidence intervals

In order to investigate the approximation used to calcultate confidence intervals below, the χ^2 goodness of fit test mentioned in section 3 has been made for the Cornell box test scene using a box filter. Each pixel value has been estimated $m = 100$ times using 10 samples for each estimate, and based on these 100 estimates the mean, $\hat{\mu}$, and sample variance, s^2 have been calculated. Given $(\hat{\mu}, s^2)$, 8 classes have been made using the 12.5%, 25%, ... , 87.5% percentiles of the $N(\hat{\mu}, s^2)$ distribution as delimiters. The test value has then been calculated as :

$$z = \sum_{i=1}^{8} \frac{(\#\text{ observations in class } i - 12.5)^2}{12.5} \tag{16}$$

For a test at a 5% level this value should be compared to the 95% percentile of the $\chi^2(5)$ distribution, i.e. 11.07. Thus for values of z larger than 11.07 the hypothesis about the estimate of the pixel value being normal distributed is rejected. Since the level of the test is 5% the largest probability of this happening when the hypothesis is actually true is 5%.

Counting the pixels in the Cornell box scene for which z is larger than 11.07 gives approximately 1900 out of 19200 pixel corresponding to about 10% of all pixels. This is higher than the expected 5% but not much considering that only $n = 10$ samples have been used to estimate each pixel value. For larger values of n the reject rate approaches 5%, and for e.g. $n = 50$ it is has been found to be about 7% for the Cornell box scene.

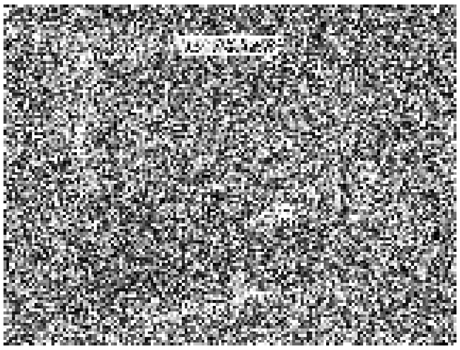

Fig. 1. The testvalue for the hypothesis that the pixel values follow a normal distribution is here visualized. Brighter values correspond to pixels where the hypothesis is less probable. The image has been histogram equalized to show the random pattern which otherwise would be very dark.

The distribution of test values is visualized in Fig. 1, and except for the border around the lightsource no pattern is apparent. The reason that the border of the light-source stands out is obviously that the distribution here is bimodal (either light or dark), so this was to be expected. As for the rest of the image, there is no pattern. This indicates that the deviations from the hypothesis are random, so it must be concluded that Eq. (6) is valid, and this conclusion remains unchanged for higher values of n.

5.2 Improvements in image quality

The results of using the adaptive sampling scheme is shown and compared to traditional sampling in Fig. 4 (see Appendix).

Table 1. The RMS error relative to a reference image based on 100, 000 samples per pixel. n indicates the average number of samples per pixel.

n	Fixed	Adaptive with tone operator	Adaptive without tone operator
40	0.062	0.049	0.069
100	0.042	0.035	0.047

Visually the difference is characterized by having eliminated virtually all noisy spots and by giving a more homogeneous image quality. The price is slightly more noise in some regions, but very importantly it should be noticed that the noise level is uniform. For the chosen confidence level all noise is *not* removed but the noise which remains is spread out evenly all over the image. Experiments have shown that this is not only true for the confidence level shown here, but also for other levels. This means that the confidence level is a good parameter for choosing the overall image quality prior to rendering.

Because of the random nature of the sampling process it is impossible to guarantee that *all* noisy spots are eliminated. As a consequence a few pixels in the caustic of the glass sphere scene are still obviously wrong. It is however evident that there is a significant improvement when using adaptive sampling instead of traditional sampling. The problem may possibly be solved by including adjacency information in the adaptive algorithm, but this will require further research.

In order to quantify the results, an analysis of the RMS error has been made for the Cornell box scene. Table 1 shows a few of these results based on a reference image with 100,000 samples per pixel. It is important to notice that the adaptive sampling without the tone operator does not perform well, because a large amount of samples are used to estimate the pixel values in the light source accurately. However, these values are all outside the display range, so in this case the accuracy is not needed. This is no longer a problem when the tone operator is included in the computation, which leads to a significant reduction of the RMS error.

All the results obtained with the Cornell box scene is shown in Fig. 2. A statistical analysis of these results shows that by using adaptive sampling with the tone operator, the reduction of the RMS error is in average 12%. Equivalently it has been found that the same RMS error can be obtained with 24% less samples when adaptive sampling is used instead of traditional sampling.

The actual value of the reduction obviously depends on the chosen test scene, so it is not possible to draw a general conclusion about the size of the reduction. However, by considering the distribution of samples in Fig. 3 is it clear that there is a tendency toward using more samples in areas with indirect illumination. Thus, it is to be expected that the largest improvements will occur in complex scenes with both direct and indirect illumination.

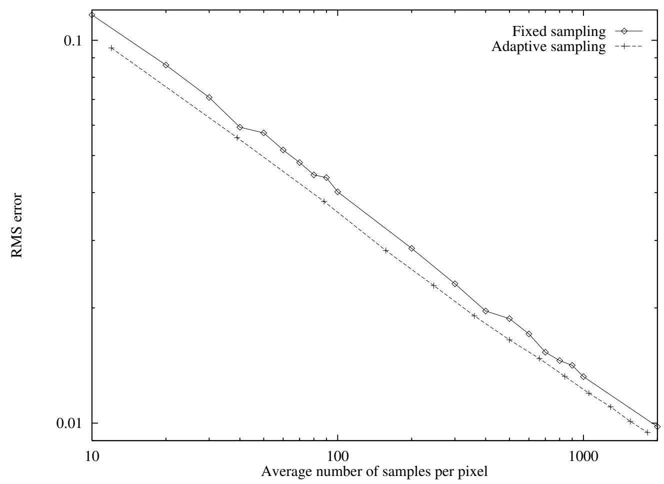

Fig. 2. The RMS error as a function of the average number of samples used per pixel. The upper curve corresponds to traditional (fixed) sampling while the lower curve corresponds to adaptive sampling with the tone operator.

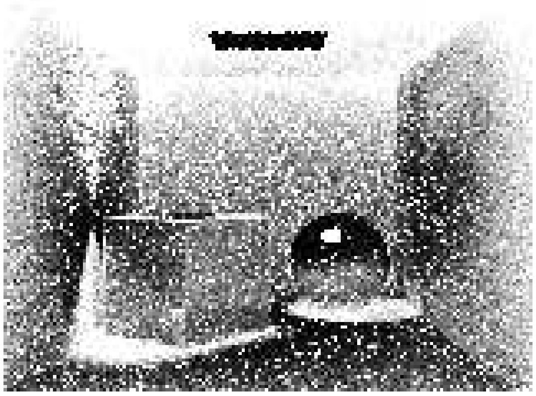

Fig. 3. The distribution of samples over the image is here visualized. Brighter values corresponds to more samples per pixel.

5.3 Bias estimation

The analysis of bias for a given pixel is made from an initial set of $10,000$ samples from the pixel, which constitute the bootstrap distribution, $\hat{F}$. Based on this the pixel value, $\hat{\theta}$, is estimated (giving the true pixel value in the bootstrap world), and the adaptive sampling algorithm is applied to obtain $B = 1000$ bootstrap estimates of the pixel value by resampling (adaptively) from $\hat{F}$. Inserting these estimates in Eq. (13) and Eq. (14) then gives the desired bias estimate.

Noticing that the bias estimate is obtained by summing 1000 independent and identically distributed random variables, we have constructed an approximate 95% confidence interval (by means of Eq. (15)) to answer the following 2 questions : Is it possible that the bias is 0, i.e. does the confidence interval contain 0 ? - And, is it possible that the bias is larger than $1/256$?

For both test scenes the number of pixels for which the bias could possibly be 0, is well below 10%, so it must be concluded that the bias is statistically significant. On the other hand the results have also shown that for about 80% of all pixels in the glass sphere scene the bias is definitely less than $1/256$, so if a display with 8 bitplanes per color is used, it is insignificant for most practical purposes. For the Cornell box the number is somewhat smaller but the bias is still quite small.

The regions with significant bias are typically those which are also sampled many times. This agrees nicely with the theoretically based observation made in [8] that bias will be highest in regions with high contrast, since these regions are sampled most because of the variability of the samples. The results are also in agreement with [8] in the sense that as the confidence level (and thus the number of samples) increases, the bias tends toward 0. Since the length of the confidence intervals approaches 0 as the number of samples is increased this means that the total error decreases monotonically toward 0. Hence, the method is consistent.

6 Conclusion

We have proposed that the tone operator is incorporated in the refinement criterion for adaptive sampling given by Purgathofer. This takes human perception and limitations of display devices into account while retaining the desirable concept of a *confidence level* for the entire image. Using this method on a pure pathtracing algorithm it has been shown that most noisy spots are eliminated and that the overall RMS error is reduced significantly. Furthermore it has been found that the overall image quality is controlled consistently by the confidence level, which makes it a useful parameter prior to rendering.

One of the problems with confidence based refinement criteria is the actual evaluation of the confidence intervals. It has been pointed out that even in the simple case where the mean of n samples is considered, a number of assumptions and approximations are necessary to calculate the confidence intervals. These assumptions cannot always be justified theoretically, but our results indicate that in most cases they cannot be rejected either. Based on this, a simple extension, which incorporates the tone operator in the calculations, has been suggested and applied succesfully.

Since the adaptive sampling scheme provided here leads to biased estimates, we have presented a method based on nonparametric bootstrapping which is used to analyze the bias. The conclusion of this analysis is that the bias is statistically significant (i.e. not zero), but that it is often insignificant in the sense that it is smaller than the color resolution of most display devices.

References

[1] K. Chiu, M. Herf, P. Shirley, S. Swamy, C. Wang, and K. Zimmerman. Spatially

non-uniform scaling functions for high contrast images. In *Proceedings of Graphics Interface 1993*, pages 245–253, Toronto, May 19–21 1993.

[2] B. Efron and R. J. Tibshirani. *An Introduction to the Bootstrap*. Monographs on Statistics and Applied Probability. Chapman & Hall, New York, 1993.

[3] J. C. Geertsema. Sequential confidence intervals based on rank tests. *The Annals of Mathematical Statistics*, 41(3):1016–1026, 1970.

[4] A. S. Glassner. *Principles of digital image synthesis*. Computer Graphics and Geometric Modeling. Morgan Kaufmann Publishers, San Francisco, 1995.

[5] H. W. Jensen and N. J. Christensen. Optimizing Path Tracing Using Noise Reduction Filters. *Proceedings of the Winter School of Computer Graphics and CAD Systems '95*, February 1995.

[6] J. T. Kajiya. The rendering equation. *Computer Graphics*, 20(4):143–150, 1986.

[7] M. G. Kendall and A. Stuart. *The Advanced Theory of Statistics*, volume 2. Charles Griffin & Co., London, 1967.

[8] D. Kirk and J. Arvo. Unbiased sampling techniques for image synthesis. *Computer Graphics*, 25(4):153–156, 1991.

[9] E. Lafortune. *Mathematical Models and Monte Carlo Algorithms for Physically Based Rendering*. Ph.D. thesis, Department of Computer Science, Katholieke Universitiet Leuven, Leuven, Belgium, February 1996.

[10] E. P. Lafortune and Y. D. Willems. Bi-directional Path Tracing. *Proceedings of Third International Conference on Computational Graphics and Visualization Techniques (Compugraphics '93)*, pages 145–153, December 1993.

[11] M. E. Lee, R. A. Redner, and S. P. Uselton. Statistically optimized sampling for distributed ray tracing. *Computer Graphics*, 19(3):61–65, 1985.

[12] D. P. Mitchell. Generating antialiased images at low sampling densities. *Computer Graphics*, 21(4):65–72, 1987.

[13] J. Painter and K. Sloan. Antialiased ray tracing by adaptive progressive refinement. *Computer Graphics*, 23(3):281–288, 1989.

[14] W. Purgathofer. A statistical mothod for adaptive stochastic sampling. *Computers & Graphics*, 11(2):157–162, 1987.

[15] H. E. Rushmeier and G. J. Ward. Energy preserving non-linear filters. *Computer Graphics*, 28(3):131–138, July 1994.

[16] J. Tumblin and H. Rushmeier. Tone reproduction for realistic images. *IEEE Computer Graphics and Applications*, 13(6):42–48, November 1993.

[17] E. Veach and L. Guibas. Bidirectional Estimators for Light Transport. *Fifth Eurographics Workshop on Rendering*, pages 147–162, June 1994.

[18] G. Ward. A contrast-based scalefactor for luminance display. In P. Heckbert, editor, *Graphics Gems IV*. Academic Press, 1994.

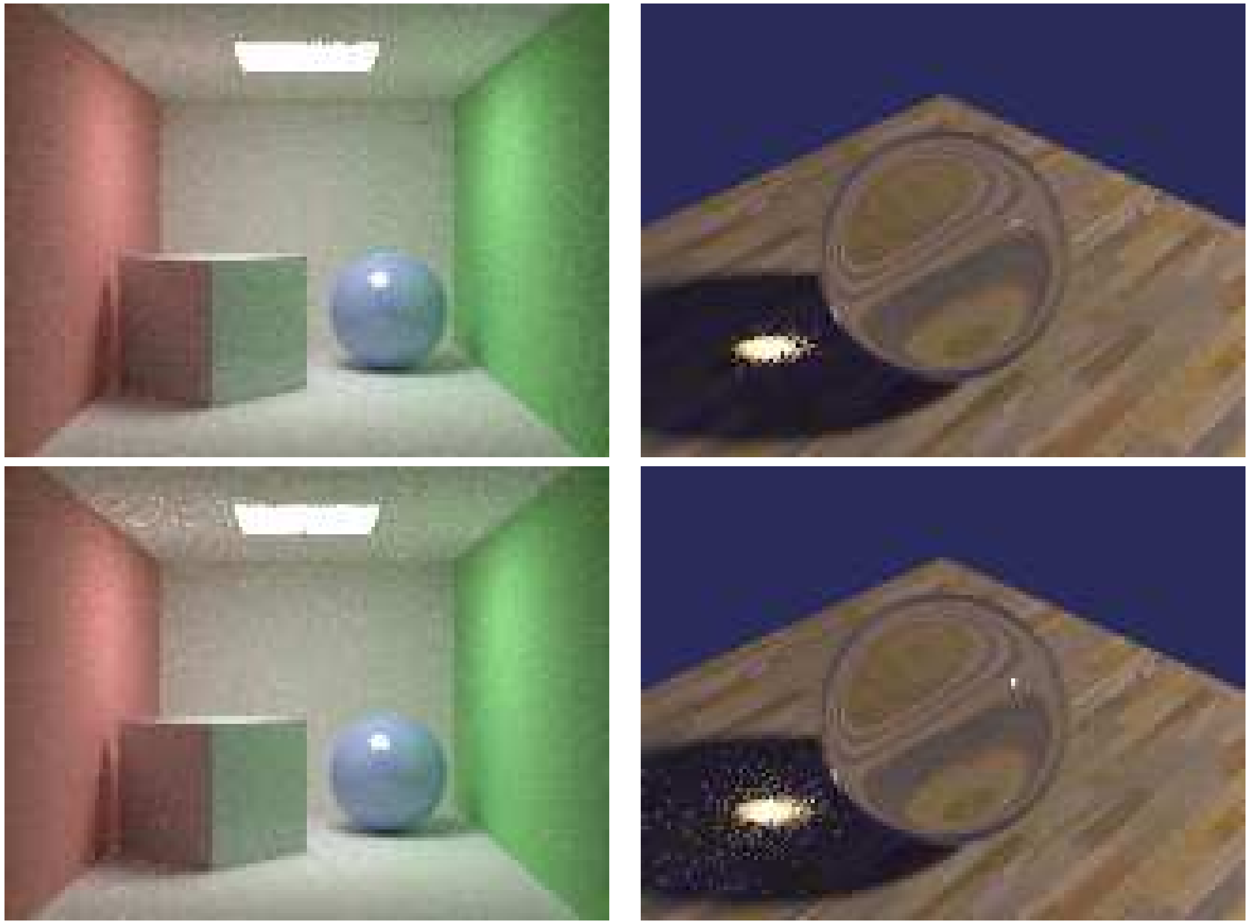

Fig. 4. The upper images have been calculated with the adaptive algorithm using 88 (left) and 314 (right) samples per pixel in average. The lower images have been calculated using 90 (left) and 300 (right) samples per pixel, but in these cases the sample rate was fixed in advance.