Entropy-based Adaptive Sampling

Jaume Rigau Miquel Feixas Mateu Sbert

Institut d'Informàtica i Aplicacions
Universitat de Girona, Catalonia, Spain

Abstract

Ray tracing techniques need supersampling to reduce aliasing and/or noise in the final image. Since not all the pixels in the image require the same number of rays, supersampling can be implemented by adaptive subdivision of the sampling region, resulting in a refinement tree. In this paper we present a theoretically sound adaptive sampling method based on entropy, the classical measure of information. Our algorithm is orthogonal to the method used for sampling the pixel or for obtaining the radiance of the hitpoint in the scene. Results will be shown for our implementation within the context of stochastic ray tracing and path tracing. We demonstrate that our approach compares well to the ones obtained by using classic strategies based on contrast and variance.

Key words: Adaptive sampling, antialiasing, contrast, entropy, pixel colour, ray tracing, stochastic sampling.

1 Introduction

Ray tracing [28] is a point-sampling-based technique for image synthesis. Rays are traced from the eye through a pixel to sample radiance at the hitpoint in the scene, where radiance is usually computed by a random walk method [25]. Since a finite set of samples is used, some of the information in the scene is lost. Thus, aliasing errors are unavoidable [8].

These errors can be reduced using extra sampling in regions where the sample values vary most. In order to obtain reliable data, the edge of an object, the contour of a shadow, or a high illumination gradient area, would need a more intensive treatment than a region with almost uniform illumination. This method of sampling is called *adaptive sampling* [8, 17]: A pixel is first sampled at a relatively low density. From the initial sample values, a refinement criterion is used to decide whether more sampling is required or not. Finally, all the samples are used to obtain the final pixel colour values [15].

Adaptive sampling can be implemented by adaptive subdivision of the sampling region. This subdivision generally corresponds to a binary tree or a quadtree [28, 10, 17]. Subdivision is triggered by the result of a refinement test based on a given error measure. New samples are then added to the newly created subregions. We can also trade aliasing for noise using *stochastic* ray tracing, as the human visual system is more sensitive to structured aliasing artifacts than to noise [17, 9].

In this paper we introduce a new refinement scheme for adaptive sampling, complementary to the one defined in [20], with the important feature that it is based on the recursive expression of the Shannon entropy, i.e. its grouping property [5]. The Shannon entropy is the classical measure of *information* [22], where information is simply the outcome of a selection from among a finite number of possibilities. In our context, entropy is interpreted as a measure of the degree of homogeneity of a pixel or subpixel. The idea behind the new scheme is to obtain sufficient *information* (homogeneity) in the refinement tree which results from the recursive decomposition of a pixel into subpixels.

One of the main features of this approach is that it uses a sound theoretical framework, namely *information theory*, to obtain the refinement process. We will show that the recursive decomposition of entropy provides us with a natural method to deal with a refinement tree. Our refinement scheme, valid for any pixel sampling and ray tracing method, will be applied to stochastic ray tracing and compared with a contrast-based technique [2, 15, 9] and a variance-based criterion [13, 19, 26].

The organization of this paper is as follows: in section 2 we present some previous work, in section 3 we introduce an adaptive sampling algorithm based on entropy, in section 4 we discuss our results, comparing them with the ones obtained by classic measures, and, finally, in section 6 we present our conclusions.

2 Previous Work

In this section we present previous work on the areas of supersampling refinement criteria, information theory and entropy-based contrast measures.

2.1 Supersampling Refinement Criteria

Three principal subproblems make up the process of obtaining a good quality image: efficient sample generation, adaptive control of the sampling rate, and filtering for image reconstruction [17]. Many approaches are to be found to deal with them:

1. Different pixel sampling methods have been introduced, among them: jittered sampling [4, 8], Poisson disk sampling [8, 15, 14], hierarchical sampling [10], complete stratification at each refinement level [21], importance sampling [23], and quasi-Monte Carlo sampling [11, 16].

2. Diverse refinement criteria for adaptive sampling, based on colour intensities and/or scene geometry, can be found to control the sampling rate: Dippé and Wold [8] present an error estimator based on the RMS signal to noise ratio and also consider its variance as a function of the number of samples; Mitchell [15] proposes a contrast [2] based on the characteristics of the human eye; Lee et al. [13], Purgathofer [19], and Tamstorf and Jensen [26] develop different methods based on the variance of the samples with their respective confidence intervals.

3. Samples are filtered to produce the final pixel values. Different filter shapes have been used in image reconstruction: box filter, triangular filter, Gaussian filter, multi-stage filter, etc. (see [9]).

For the purpose of this paper we review three commonly used refinement criteria: contrast, depth difference, and variance of the samples.

Mitchell, in [15], uses a contrast measure [2] for each RGB channel defined by

$$C = \frac{I_{\max} - I_{\min}}{I_{\max} + I_{\min}}, \tag{1}$$

where $I_{\min}$ and $I_{\max}$ are, respectively, the minimum and maximum light intensities of the channel. Supersampling is done if any contrast is higher than a given threshold. Mitchell proposes RGB threshold values (0.4, 0.3 and 0.6, respectively) based on the relative sensitivity of the visual system.

Simmons and Séquin [24], within an interactive rendering context, use a *colour priority* value based on contrast and perception [15, 9] combined with a geometric measure for refinement, the *depth difference*, given by $1 - \frac{d_{\min}}{d_{\max}}$ where $d_{\max}$ and $d_{\min}$ represent maximum and minimum distance.

The basic idea of variance-based methods [13, 19, 26] is to continue sampling until the confidence level or probability that the true value L is within a given tolerance d of the estimate value $\widehat{L}$ is $1-\alpha$:

$$\Pr[L \in (\widehat{L} - d, \widehat{L} + d)] = 1 - \alpha. \tag{2}$$

Mitchell considers that variance is a poor measure of visual perception of local variation [15]. Kirk and Arvo showed that these methods are biased and proposed a simple correction scheme [12].

Refinement criteria have also recently been applied in the image-based rendering field to weight pixel colour for reconstruction purposes [18] and adaptive sampling strategies [6, 7]. Also Bolin and Meyer [1] have developed a perceptually-based approach using statistical and vision models.

2.2 Information Theory

The Shannon entropy $H(X)$ of a discrete random variable X with values in the set $\mathcal{X} = \{x_1, \ldots, x_n\}$ is defined [22] as

$$H(X) = -\sum_{i=1}^{n} p_i \log p_i, \tag{3}$$

where $n = |\mathcal{X}|$, $p_i = \Pr[X = x_i]$ for $i \in \{1 \ldots n\}$, the logarithm is taken in base 2 (in this case, entropy is expressed in bits), and also the convention that $0 \log 0 = 0$ is used by continuity. As $-\log p_i$ represents the *information* associated with the result x_i, the entropy gives the average information or *uncertainty* of a random variable.

Some relevant properties [22, 5] of the entropy are:

- $0 \leq H(X) \leq \log n$.
- If we equalize the probabilities, entropy increases.
- Grouping:

$$\begin{aligned} H(p_1, \ldots, p_n) = & H(p_1 + p_2, p_3, \ldots, p_n) \\ & + (p_1 + p_2) H\left(\frac{p_1}{p_1 + p_2}, \frac{p_2}{p_1 + p_2}\right). \end{aligned} \tag{4}$$

It is worth mentioning the case $n = 2$, with $p_1 = p$ and $p_2 = 1 - p$. The entropy of this probability distribution is called *binary entropy* (Figure 1) and is given by

$$H(X) = -p \log p - (1-p) \log(1-p). \tag{5}$$

2.3 Entropy-based Contrast Measures

In this section we summarize the previous work on entropy-based contrast measures done by Rigau et al. [20].

The *pixel channel entropy* was defined by

$$H^c = -\sum_{i=1}^{N_s} p_i \log p_i, \tag{6}$$

where $p_i = \frac{c_i}{\sum_{i=1}^{N_s} c_i}$ represents the channel colour fraction of ray i with respect to the sum of the colours of the

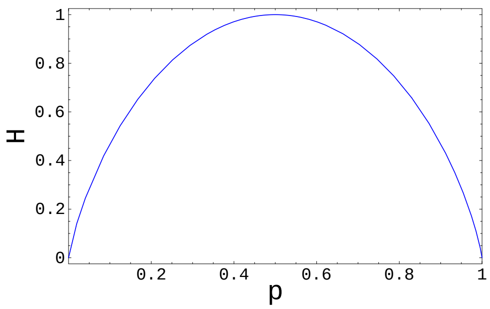

Figure 1: Binary entropy corresponding to the probability distribution $\{p, 1-p\}$ of random variable X. The maximum value, $H(X) = 1$, is obtained when $p = \frac{1}{2}$ and the minimum value, $H(X) = 0$, when $p = 0$ or $p = 1$.

same channel of all the rays passing through the pixel, and N_s is the number of rays traversing the pixel. Pixel channel entropy was interpreted as the channel colour homogeneity of the rays passing through the pixel. It can also be considered as a measure of the pixel colour quality.

In order to give a pixel contrast measure between 0 and 1, the pixel channel entropy is normalized with $\log N_s$. Thus, the *pixel channel contrast* was defined by

$$C^c = 1 - \frac{H^c}{\log N_s} \tag{7}$$

and represents the channel colour *inhomogeneity* of a pixel. When considering all the colour channels (N_c), the global *pixel colour contrast* [20] was given by

$$\mathbf{C}^c = \frac{\sum_{i=1}^{N_c} \omega_i C_i^c \overline{c}_i}{\sum_{i=1}^{N_c} \omega_i}, \tag{8}$$

where the channel contrasts are weighted by perceptual coefficients ω_i and $\overline{c}_i = \frac{1}{N_s}\sum_{i=1}^{N_s} c_i$, the colour average of channel i of all the pixel rays (channel *importance*).

Similar to (6), the *pixel geometric entropy* H^g was defined by

$$H^g = -\sum_{i=1}^{N_s} p_i \log p_i, \tag{9}$$

where now $p_i = \frac{\cos\theta_i/d_i^2}{\sum_{i=1}^{N_s}\cos\theta_i/d_i^2}$ represents the geometric fraction of ray i with respect to the sum of the geometric factors of all the rays traversing a pixel. The geometric information of each ray is given by the angle θ_i which the normal forms at the hitpoint with the ray, and also by the distance d_i between this point and the eye. Similar to the case of colour, the geometric entropy represents the pixel geometric homogeneity. Analogous to (7), the *pixel geometric contrast* C^g was defined by

$$C^g = 1 - \frac{H^g}{\log N_s}, \tag{10}$$

which represents the geometric inhomogeneity of a pixel.

It is also possible to obtain alternative colour and geometric contrast measures by substituting the pixel entropy for the binary pixel entropy, which is computed by only considering the maximum and minimum values captured by the pixel (in formula (5), the probability distribution would be $\{\frac{\min}{\min+\max}, \frac{\max}{\min+\max}\}$).

A combination of colour and geometric contrasts was considered. This combination enables the influence of both measures to be graduated with a coefficient δ between 0 and 1:

$$\mathcal{C} = \delta\mathbf{C}^c + (1-\delta)C^g. \tag{11}$$

3 Adaptive Sampling Algorithm-based on Entropy

In this paper, our attention focuses on obtaining an adaptive algorithm centred mainly on the refinement phase. The approach to be used in refinement will be to evaluate the similarity or homogeneity of the *information* provided by the set of samples in a given region. If the information obtained from this region is heterogeneous we will refine it until each subregion is uniform. This process is a naturally recursive process, giving rise to a refinement tree.

3.1 Recursive Entropy Tree

Generalizing the grouping property (4), the entropy can be recursively decomposed in the following way: Let X be a discrete random variable over the set $\mathcal{X} = \{x_1, \ldots, x_n\}$ with probability distribution $p = \{p_1, \ldots, p_n\}$ where $p_i = \Pr[X = x_i]$. Let us consider a partition of the set $\mathcal{X}$ in m-disjoint sets $\mathcal{G} = \{\mathcal{G}_1, \ldots, \mathcal{G}_m\}$ where $|\mathcal{G}_j| = n_j$. Let us associate the discrete random variable Y to $\mathcal{G}$ with probability distribution $q = \{q_1, \ldots, q_m\}$ where $q_j = \sum_{k=1}^{n_j} p_{j_k}$ ($j_k \in \{1, \ldots, n\}$), and a new discrete random variable Y_j to each set $\mathcal{G}_j$ with probability distribution $r_j = \{r_{j_1}, \ldots, r_{j_{n_j}}\}$ where $r_{j_k} = \frac{p_{j_k}}{q_j}$. Then

$$H(X) = \sum_{j=1}^{m} q_j H(Y_j) - \sum_{j=1}^{m} q_j \log q_j. \tag{12}$$

This formula can be written as $H(X) = H_{in} + H_{out}$ where $H_{in} = \sum_{j=1}^{m} q_j H(Y_j)$ and $H_{out} = H(Y) = -\sum_{j=1}^{m} q_j \log q_j$ represent, respectively, the hidden information (pending to be discovered) and the information already acquired in the descent of the tree (see Figure 2).

In our case, formula (12) can also be interpreted (for one colour channel (6)) in the following way:

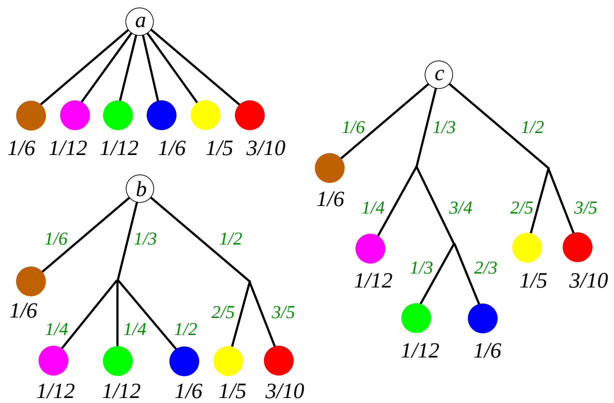

Figure 2: Grouping property of entropy. The entropy of probability distributions of ⓐ *is* $H(\frac{1}{6},\frac{1}{12},\frac{1}{12},\frac{1}{6},\frac{1}{5},\frac{3}{10})$, *of* ⓑ *is* $H(\frac{1}{6},\frac{1}{3},\frac{1}{2})+\frac{1}{3}H(\frac{1}{4},\frac{1}{4},\frac{1}{2})+\frac{1}{2}H(\frac{2}{5},\frac{3}{5})$, *and of* ⓒ *is* $H(\frac{1}{6},\frac{1}{3},\frac{1}{2})+\frac{1}{3}(H(\frac{1}{4},\frac{3}{4})+\frac{3}{4}H(\frac{1}{3},\frac{2}{3}))+\frac{1}{2}H(\frac{2}{5},\frac{3}{5})$. *Accordingly to (12), all have the same value:* $H(a) = H(b) = H(c) = 2.445$.

- $H(X)$ represents the entropy of the whole image.
- $H(Y_j)$ represents the entropy of each root pixel.
- Probability q_j is the colour of pixel j divided by the sum of the colours of all pixels. It can be considered as the "importance" of pixel j.

The decomposition of entropy (12) can be recursively extended to the subpixels. This interpretation can also be applied to geometric entropy (9).

In our approach, probabilities are obtained by stochastic sampling. From the definition of entropy, we can see that when the number of samples tends to infinity, entropy also goes to infinity. In fact, we can consider that the original continuous scene contains infinite information. The following sampling algorithm will *capture* or *extract* more information from the regions with more sample variation.

3.2 Algorithm

In this section we show how a practical adaptive sampling algorithm can be obtained from the entropy tree. For the sake of simplicity, in the following analysis we only consider the colour information of a channel, although the final algorithm will take the combination of colour and geometric contrasts (11) into account, as in [24].

A general description of our algorithm is as follows: On the image plane we sample each pixel to capture the colour of hitpoints and thus evaluate the information content (entropy) from the colour probability distribution. If the information of a pixel is high enough, i.e. the rays give us sufficient colour homogeneity on that pixel, refinement is not made, and the colour reconstruction of this pixel is done. When it is not high enough, this pixel is subdivided into regions and we proceed in the same way for each region (subpixel).

This recursive process defines a tree with two well-separated phases for a pixel: refinement (tree descent) and colour computation (tree ascent). The descent in the refinement tree can be interpreted as a progressive information gain. The information acquired at each level is added together so that, at the end of the refinement process, the total information from the tree is the sum of the information obtained over all the branches (see formula (11)).

Before introducing the algorithm we will give the definitions of the data used in it. Concerning the tree data structure, the root (level $n = 0$) is the image, level $n = 1$ corresponds to the N_p pixels of the image, and levels $n > 1$ to the subpixels. Each new n-node (i.e., node of $n > 0$ level), is sampled N_s times and it can potentially be subdivided in N_r regions or subpixels of equal size ($N_s \in N_r\mathbb{N}^+$). Other data referred to in the refinement phase are described in Table 1. To compute the final colour of a pixel, we follow a path through the tree (see Figure 3). In the analysis below, we focus our attention on the tree-path k of length N (see Table 1) going from pixel k_0 to subpixel k_{N-1}. In this path, p_n represents the probability of the tree-branch at level n and q_n the *importance* of the n-node. In our algorithm, this quantity appears in a natural way due to recursive decomposition of the entropy (see (12) and Figures 2 and 3). The value of importance is given by

$$q_n = \begin{cases} 1, & n = 0, \\ p_0 \cdots p_{n-1} = \frac{\overline{c}_{0,k_0}}{\sum_{i \in R_0} \overline{c}_{0,i}} \prod_{\ell=1}^{n-1} p_\ell, & n > 0. \end{cases} \tag{13}$$

For our purposes, q_n needs not to be normalized, thus we omit normalization constant $\sum_{i \in R_0} \overline{c}_{0,i}$ and we take $q_n = \overline{c}_{0,k_0} \prod_{\ell=1}^{n-1} p_\ell$. The computation of q_n can then be simplified to (see proof in Appendix):

$$q_n = \frac{\overline{c}_n}{{N_r}^{n-1}}. \tag{14}$$

Now we proceed to explain the algorithm. In the descent phase we sample an n-node and compute the contrast using expression (11). In (8) we must substitute the channel importance $\overline{c}$ by q_n and we take RGB perceptual coefficients [27] $\omega_r = 0.213$, $\omega_g = 0.715$ and $\omega_b = 0.072$ which capture the sensitivity of human colour perception.

Thus, for each n-node, the colour contrast (8) converts

name	*description*	*relations*
R_n	Set of regions of an n-node	$\|R_0\| = N_p$, $\forall_{n>0}\|R_n\| = N_r$
k	Path: k_n is the region taken at level n	$k = (k_0, k_1, \ldots, k_{N-2}, k_{N-1})$, $N > 0$, $\forall_{n<N} k_n \in R_n$
S_n	Set of samples of a n-node	$\|S_0\| = N_s N_p$, $\forall_{n>0}\|S_n\| = N_s$
$S_{n,i}$	Set of samples of a n-node region $i \in R_n$	$\|S_{n,i}\| = \frac{\|S_n\|}{\|R_n\|}$, $S_n = \bigcup_{i\in R_n} S_{n,i}$
$c(s)$	Colour obtained with sample s	*RGB* value
$\overline{c}_n$	Average colour in the n-node	$\overline{c}_n = \frac{1}{\|S_n\|}\sum_{s\in S_n} c(s)$
$\overline{c}_{n,i}$	Average colour in the n-node region $i \in R_n$	$\overline{c}_{n,i} = \frac{1}{\|S_{n,i}\|}\sum_{s\in S_{n,i}} c(s)$, $\overline{c}_n = \frac{1}{\|R_n\|}\sum_{i\in R_n} \overline{c}_{n,i}$
p_n	Probability of region k_n of n-node	$p_n = \frac{\sum_{s\in S_{n,k_n}} c(s)}{\sum_{s\in S_n} c(s)} = \frac{\overline{c}_{n,k_n}}{\sum_{i\in R_n} \overline{c}_{n,i}}$
q_n	Probability of the n-node	$q_n = \prod_{\ell=0}^{n-1} p_\ell$

Table 1: Description of the data in the refinement phase. The constants are: N_p, the number of pixels in the image, N_r, the number of equal area regions of an n-node, and N_s, the number of samples cast in an n-node ($N_s \in N_r\mathbb{N}^+$).

into

$$\mathbf{C}^{c_n} = \sum_{i=1}^{N_c} \omega_i C_i^{c_n} q_{n\,i} \tag{15}$$

and the colour and geometric combination (11) will be

$$\mathcal{C}_n = \delta \mathbf{C}^{c_n} + (1-\delta) C^{g_n}. \tag{16}$$

Note that this expression could be calculated from the respective binary versions of colour and geometric contrasts (see section 2.3).

In the algorithm, we subdivide when the contrast for inhomogeneity of n-node is greater than a given threshold ($\mathcal{C}_n > \epsilon$). Thus, the ascent phase begins when the test fails ($\mathcal{C}_n \leq \epsilon$). This happens because either the contrast (which represents the colour inhomogeneity) or the importance ($q_n \to 0$ for growing n) are low. In the colour reconstruction process, each n-node in the path provides its colour estimation $\widehat{c}_n$ computed from S_n where each colour $c(s)$ is filtered.

The final colour of an n-node is given by

$$c_n = \begin{cases} \widehat{c}_n, & \text{if } \mathcal{C}_n \leq \epsilon, \\ \sum_{i\in R_n} c_{n,i}, & \text{otherwise,} \end{cases} \tag{17}$$

where $c_{n,i}$ is the final colour of i-region of the n-node. Finally, we get c_1 for the colour of the pixels (or equivalently c_{0,k_0} in the path considered). An example of the process is shown in Figure 3.

Observe that importance sampling is naturally integrated in the algorithm. Following importance sampling criterion a function should be sampled proportionally to its value which is what we get with our adaptive descent.

4 Empirical Results

In Figure 5 we present comparative results with different techniques for the test scene (Figure 4). We compare the following methods:

a) *Classic contrast*: A recursive adaptive sampling scheme based on contrast by channel (1) (with thresholds proportional to the visual system) weighted by their respective channel colour average [9, 24]. The maximum recursive level has been limited to 4 (Figure 5.*a*).

b) *Importance-weighted contrast*: The same as in **a** but each channel contrast is weighted with the respective importance q (14), as in our approach (Figure 5.*b*).

c) *Variance-based contrast*: Statistical approach (2) [26] (Figure 5.*c*).

d) *Entropy-based contrast*: Our approach (17) taking only colour contrast, $\delta = 1$ in (16) (Figure 5.*d*).

All methods have been implemented on the *RenderPark* [3] software (`www.renderpark.be`). Observe that our approach can be easily implemented on any standard hierarchical algorithm using importance (14) and the new refinement criterion (16), with negligible additional cost.

In **a**, **b**, and **d**, the number of subdivisions, N_r, is 4 and 8 rays, N_s, were cast in a stratified way at each n-node (pixel or subpixel) to compute the contrast measures for the refinement decision. These rays were re-used at the next levels in the tree. In **c**, gropus of 8 rays were added in a stratified way until meeting the condition of the criterion (2) with $\alpha = 0.1$ and $d = 0.025$. An implementation of classic path-tracing with next event estimator [3] was used to compute all images. The parameters were tuned so that all four test images were obtained with a similar average number of rays per pixel (60) and computation

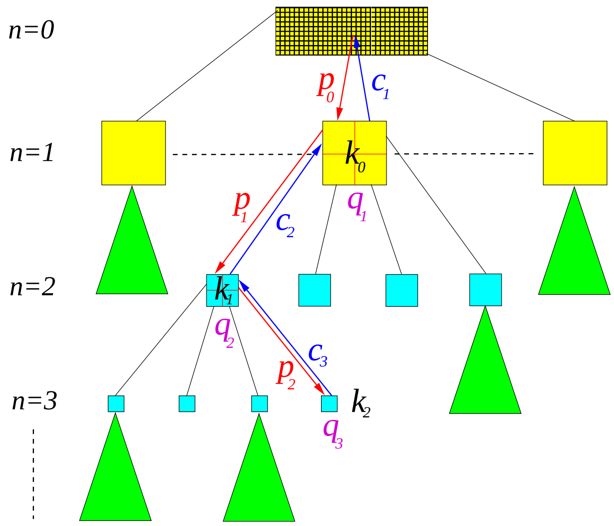

Figure 3: A tree-path $k = (k_0, k_1, k_2)$ of length $N = 3$. The number of regions of an n-node is $N_r = 4$. We show the computation of the k_0-pixel colour: $c_{0,k_0} = c_1$ from the refinement (red) and reconstruction (blue) phases. The probabilities p_n and importances q_n (Table 1) are computed in the refinement phase to evaluate the entropy contrast (16).

cost. A constant box filter was used in the reconstruction phase for all the methods.

The resulting images are shown in column *i* of Figure 5 with close-ups in column *iii*. Sampling maps are given in column *ii* (warm colours correspond to the highest sampling rate and the cold colours to the lowest).

The overall aspect of the images in Figure 5.*i* shows that our supersampling scheme performs best. Observe, for instance, the reduced noise in the shadows cast by the objects. This is further checked in the close-up images in Figure 5.*iii*. Observe also the detail of the sphere shadow reflected on the pyramid. It must be noted that we managed to improve the classic contrast approach in **a** greatly by including the importance used in our scheme (compare results in Figure 5.*a* with Figure 5.*b*). Comparison of the sampling temperature maps in Figure 5.*ii* shows a better discrimination of complex regions of the scene in the entropy case against the classic contrast case. This explains the better results obtained by our approach. Moreover, the variance-based approach **c** (Figure 5.*c*) also performs better than the classic contrast-based methods **a** (Figure 5.*a*) and **b** (Figure 5.*b*). Its sampling map also explains why it performs better. However, it is unable to render the reflected shadows under the mirrored pyramid and sphere with precision (see close-up in Figure 5.*c.iii*).

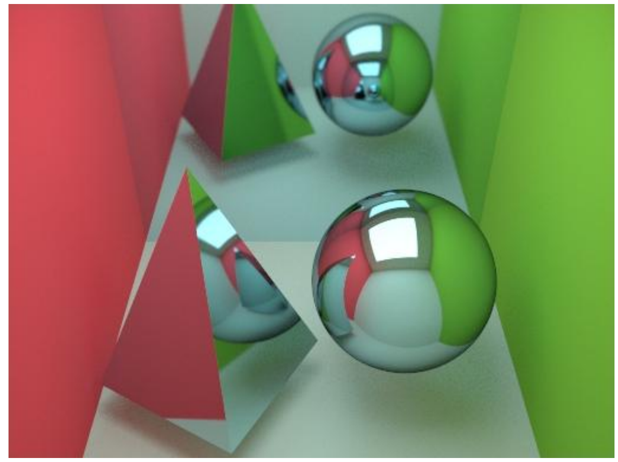

Figure 4: Reference image used in the test in Figure 5 with 1024 rays per pixel.

In Figure 6.*a* we show another scene obtained with our approach using an average of 200 rays per pixel and $\delta = 0.95$ (16). Observe, in Figure 6.*b*, how well the sampling map works out both the geometric and colour details, as in the shadow contours on the walls.

5 Conclusions

We have presented a new adaptive sampling algorithm for ray tracing based on the recursive decomposition of the entropy of a pixel, computed from the sampled radiances through the pixel. Entropy is shown to be a natural measure for the criterion used in the refinement tree. Thus, we use a sound theoretical framework (information theory) in order to establish the refinement criterion.

The results obtained show that the new refinement algorithm offer a substantial improvement over the classical techniques, both contrast and variance-based. From this, it could be deduced that entropy captures better the inhomogeneity of a region. Future work will address the problem of finding automatic criteria for the threshold used in the refinement test and the analysis of the bias incurred by our algorithm, in the sense of Kirk and Arvo [12] and Tamstorf and Jensen[26].

Acknowledgements

This project has been funded in part with grant numbers TIC-2001-2416-C03-01 of the Ministry of Science and Technology (Spanish Government) and 2001-SGR-00296 and ACI2002-29 of the Ministry of Universities, Research and the Information Society (Catalan Government).

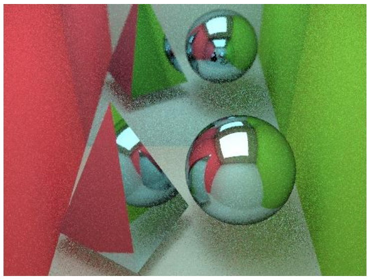

($a.i$) Classic contrast

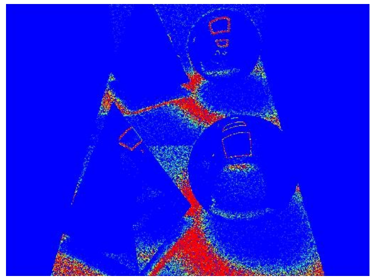

($a.ii$) Sampling map of ($a.i$)

($a.iii$) Close-up of ($a.i$)

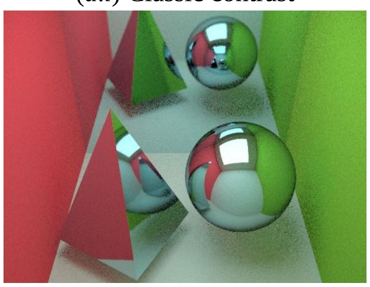

($b.i$) Importance-weighted contrast

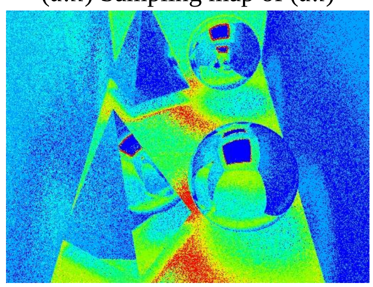

($b.ii$) Sampling map of ($b.i$)

($b.iii$) Close-up of ($b.i$)

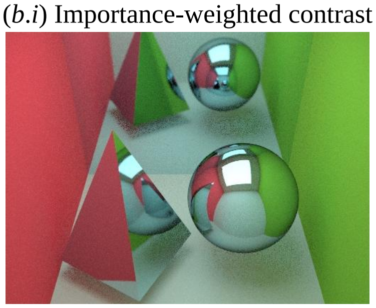

($c.i$) Variance-based contrast

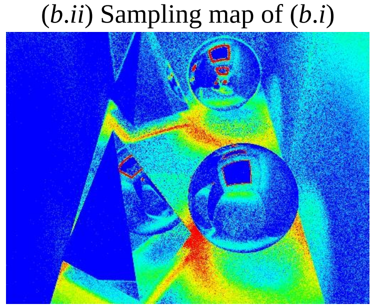

($c.ii$) Sampling map of ($c.i$)

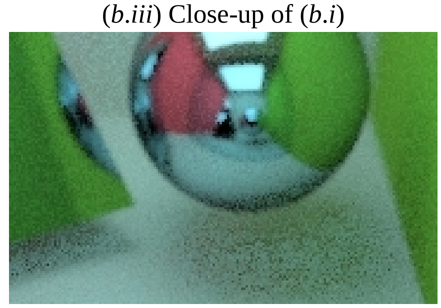

($c.iii$) Close-up of ($c.i$)

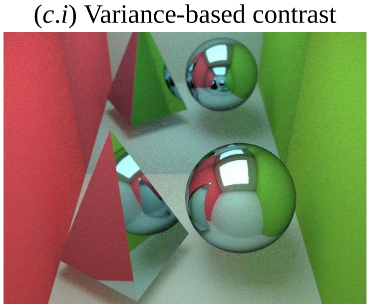

($d.i$) Entropy-based contrast

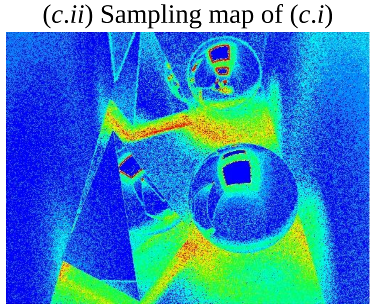

($d.ii$) Sampling map of ($d.i$)

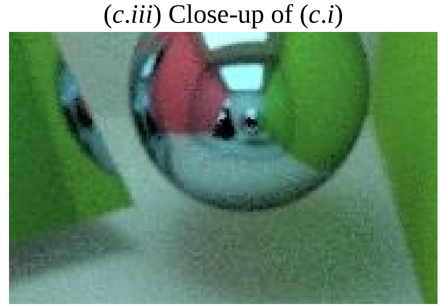

($d.iii$) Close-up of ($d.i$)

Figure 5: Results of comparisons: (a) adaptive sampling scheme based on classic contrast, (b) importance-weighted contrast, same as in (a) but weighting with importance q (14), (c) variance-based method, and (d) entropy-based method with only colour contrast ($\delta = 1$). Column (i) shows the resulting images, (ii) the colour temperature sampling map of (i), and (iii) a close-up region of (i). Average number of rays per pixel is 60 in all methods, with a similar computation cost.

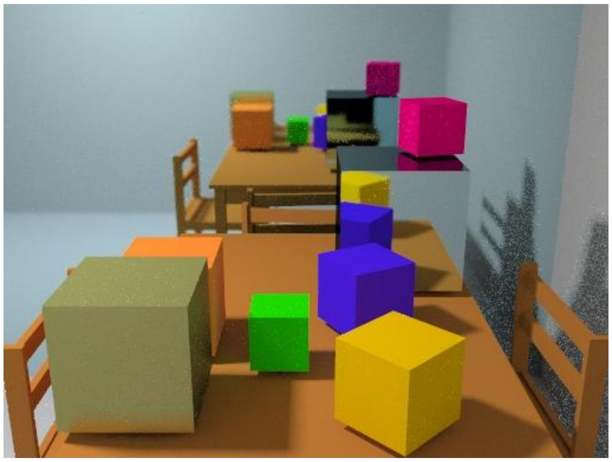

(a) Entropy-based contrast

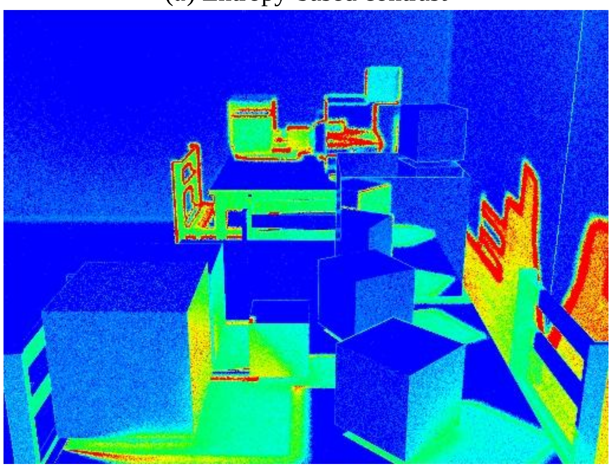

(b) Sampling map of (a)

Figure 6: (a) An image obtained by our approach with 200 rays per pixel on the average and $\delta = 0.95$. (b) Colour temperature map of the sampling to obtain (a).

Appendix

Observe first that for a given path and $n > 0$, the colour $\overline{c}_n$ of an n-node is more accurate than the colour average of its respective region, k_{n-1}, in the preceding level. Thus, the accuracy of p_n, and at the same time of q_n, can be increased by substituting $\overline{c}_{n-1,k_{n-1}}$ for $\overline{c}_n$. Let us prove now (14) by induction:

Proof. For $n = 1$,

$$q_n = \overline{c}_{0,k_0} \approx \overline{c}_1 = \frac{\overline{c}_1}{{N_r}^0} = \frac{\overline{c}_n}{{N_r}^{n-1}}.$$

Hypothesis: $\forall_{0<\ell<n} q_\ell = \frac{\overline{c}_\ell}{{N_r}^{\ell-1}}$. Then, for $n > 1$

$$\begin{aligned} q_n &= \overline{c}_{0,k_0} \prod_{\ell=1}^{n-1} p_\ell = q_{n-1} p_{n-1} \\ &= \frac{\overline{c}_{n-1}}{{N_r}^{n-2}} \frac{\overline{c}_{n-1,k_{n-1}}}{\sum_{i \in R_{n-1}} \overline{c}_{n-1,i}} \\ &\approx \frac{\overline{c}_{n-1}}{{N_r}^{n-2}} \frac{\overline{c}_n}{\overline{c}_{n-1} N_r} = \frac{\overline{c}_n}{{N_r}^{n-1}}. \quad \square \end{aligned}$$

References

[1] Mark R. Bolin and Gary W. Meyer. A perceptually based adaptive sampling algorithm. *SIGGRAPH'98 Conference Proceedings*, pages 299–309, July 1998. Held in Orlando (FL), USA.

[2] Terrence M. Caelli. *Visual Perception: Theory and Practice*. Pergamon Press, Oxford, UK, 1981.

[3] Computer Graphics Research Group. *RenderPark: A Photorealistic Rendering Tool*. Katholieke Universiteit Leuven, Leuven, Belgium, November 2000.

[4] Robert L. Cook, Thomas Porter, and Loren Carpenter. Distributed ray tracing. *Computer Graphics (Proceedings of SIGGRAPH'84)*, 18(3):137–145, July 1984. Held in Minneapolis (MN), USA.

[5] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley Series in Telecommunications, 1991.

[6] Lucia Darsa and Bruno Costa. Multi-resolution representation and reconstruction of adaptively sampled images. In *Proceedings of IX Brazilian Symposium on Computer Graphics and Image Processing (SIBGRAPI'96)*, pages 321–328, October 1996.

[7] Lucia Darsa, Bruno Costa, and Amitabh Varshney. Navigating static environments using image-space simplification and morphing. In Michael Cohen and David Zeltzer, editors, *1997 Symposium on Interactive 3D Graphics*, pages 25–34. ACM SIGGRAPH, April 1997.

[8] Mark A.Z. Dippé and Erling H. Wold. Antialiasing through stochastic sampling. *Computer Graphics (Proceedings of SIGGRAPH'85)*, 19(3):69–78, July 1985. Held in San Francisco (CA), USA.

[9] Andrew S. Glassner. *Principles of Digital Image Synthesis*. Morgan Kaufmann Publishers, San Francisco (CA), USA, 1995.

[10] James T. Kajiya. The rendering equation. *Computer Graphics (Proceedings of SIGGRAPH'86)*, 20(4):143–150, August 1986. Held in Dallas (TX), USA.

[11] Alexander Keller and Stefan Heinrich. Quasi-Monte Carlo methods in computer graphics, Part I: The QMC-buffer. Technical Report 242/94, University of Kaiserslautern, Kaiserslautern, Germany, 1994.

[12] David Kirk and James Arvo. Unbiased variance reduction for global illumination. In *Proceedings of the 2nd Eurographics Workshop on Rendering*, pages 153–156, May 1991. Held in Barcelona, Spain.

[13] Mark E. Lee, Richard A. Redner, and Samuel P. Uselton. Statiscally optimized sampling for distributed ray tracing. *Computer Graphics (Proceedings of SIGGRAPH'85)*, 19(3):61–67, July 1985. Held in San Francisco (CA), USA.

[14] Michael McCool and Eugene Fiume. Hierarchical Poisson disk sampling distributions. In *Proceedings of Graphics Interface'92*, pages 94–105, May 1992. Held in Vancouver, Canada.

[15] Don P. Mitchell. Generating antialiased images at low sampling densities. *Computer Graphics (Proceedings of SIGGRAPH'87)*, 21(4):65–72, July 1987. Held in Anaheim (CA), USA.

[16] Ryutarou Ohbuchi and Masaki Aono. Quasi–Monte Carlo rendering with adaptive sampling. Technical report, Tokyo Research Laboratory, IBM Japan Ltd., Tokyo, Japan, 1996.

[17] James Painter and Kenneth Sloan. Antialiased ray tracing by adaptive progressive refinement. *Computer Graphics (Proceedings of SIGGRAPH'89)*, 23(3):281–288, July 1989. Held in Boston (MA), USA.

[18] Kari Pulli, Michael F. Cohen, Tom Duchamp, Hugues Hoppe, Linda Shapiro, and Werner Stuetzle. View-based rendering: Visualizing real objects from scanned range and color data. In Julie Dorsey and Philipp Slusallek, editors, *Rendering Techniques'97 (Proceedings of the 8th Eurographics Workshop on Rendering)*, pages 23–34, New York (NY), USA, June 1997. Springer-Verlag Vienna-New York. Held in St. Etienne, France.

[19] Werner Purgathofer. A statistical method for adaptive stochastic sampling. *Eurographics'86: Proceedings of the European Conference and Exhibition*, 11(2):157–162, August 1986. Held in Lisbon, Portugal.

[20] Jaume Rigau, Miquel Feixas, and Mateu Sbert. *New Contrast Measures for Pixel Supersampling*, pages 439–451. Springer-Verlag London Limited, London, UK, July 2002. Proceedings of CGI'02, Bradford, UK.

[21] Christophe Schlick. An adaptive sampling technique for multidimensional ray tracing. In *Proceedings of the 2nd Eurographics Workshop on Rendering*, pages 48–56, May 1991. Held in Barcelona, Spain.

[22] Claude E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423, 623–656, July, October 1948.

[23] Peter Shirley. A ray tracing method for illumination calculation in diffuse–specular scenes. In *Proceedings of Graphics Interface'90*, pages 205–212, May 1990. Held in Halifax, Canada.

[24] Maryann Simmons and Carlo H. Séquin. Tapestry: A dynamic mesh-based display representation for interactive rendering. In Bernard Péroche and Holly Rushmeier, editors, *Rendering Techniques 2000 (Proceedings of the 11th Eurographics Workshop on Rendering)*, pages 329–340, New York (NY), USA, June 2000. Springer-Verlag Vienna-New York.

[25] László Szirmay-Kalos. *Monte Carlo Methods in Global Illumination*. Institute of Computer Graphics, Vienna University of Technology, Vienna, Austria, 1999.

[26] Rasmus Tamstorf and Henrik W. Jensen. Adaptive sampling and bias estimation in path tracing. In Julie Dorsey and Philipp Slusallek, editors, *Rendering Techniques'97 (Proceedings of the 8th Eurographics Workshop on Rendering)*, pages 285–295, New York (NY), USA, June 1997. Springer-Verlag Vienna-New York. Held in St. Etienne, France.

[27] International Telecommunication Union. Basic parameter values for the HDTV standard for the studio and for international programme exchange. In *ITU-R Recommendation BT.709*. ITU, Geneva, Switzerland, 1990. Formerly Comité Consultatif International Radio.

[28] Turner Whitted. An improved illumination model for shaded display. *Communications of the ACM*, 23(6):343–349, June 1980.