Two-Level Adaptive Sampling for Illumination Integrals using Bayesian Monte Carlo

R. Marques[1], C. Bouville[2], L. P. Santos[3] and K. Bouatouch[2]

[1]Universitat Pompeu Fabra (UPF), Grup de Recerca de Tecnologies Interactives (GTI), Spain
[2]Intitut de Recherche en Informatique et Systèmes Aléatoires (IRISA), France
[3]Universidade do Minho (UM), Departamento de Informática (DI), Portugal

Abstract

Bayesian Monte Carlo (BMC) is a promising integration technique which considerably broadens the theoretical tools that can be used to maximize and exploit the information produced by sampling, while keeping the fundamental property of data dimension independence of classical Monte Carlo (CMC). Moreover, BMC uses information that is ignored in the CMC method, such as the position of the samples and prior stochastic information about the integrand, which often leads to better integral estimates. Nevertheless, the use of BMC in computer graphics is still in an incipient phase and its application to more evolved and widely used rendering algorithms remains cumbersome. In this article we propose to apply BMC to a two-level adaptive sampling scheme for illumination integrals. We propose an efficient solution for the second level quadrature computation and show that the proposed method outperforms adaptive quasi-Monte Carlo in terms of image error and high frequency noise.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Raytracing

1. Introduction

Global illumination is a computationally very demanding task, whose execution time depends on the overall number of samples used to estimate the illumination integrals. Assigning a fixed number of samples to each integral evaluation is not optimal, since different image regions often exhibit different types of incident illumination and material's light scattering properties, therefore requiring different sampling rates to achieve identical quality. Alternatively, adaptive sampling allows for locally selecting the number of samples. When using adaptive sampling, a first estimate with few samples is computed. Then, if some quality criterion of the estimate is not met, a new sample set is drawn and used to refine the previous estimate. Consequently, regions of the image where the sampled signal contains higher frequencies will be more densely sampled.

We propose to apply a two-level adaptive sampling approach to the Bayesian Monte Carlo (BMC) framework for evaluating the illumination integral. BMC is particularly well suited to this problem given its flexibility to incorporate prior knowledge which, in the case of the second level integral estimate, would be the information produced by the first integral estimate. Our contribution is twofold: (i) we derive a closed form solution of the second level BMC quadrature; (ii) we propose an efficient method for the second level quadrature computation within the spherical Gaussian (SG) framework of [MBR*13a].

2. Related work

2.1. Adaptive sampling

Most of the work regarding adaptive sampling and refinement criteria has been developed for image plane sampling. The rationale is to detect regions of the image plane for which the sampling rate does not allow an adequate signal reconstruction. The sampling rate is then locally increased to reduce aliasing. Finally, the sampled signal (i.e., the incident radiance function at the image plane) is reconstructed by combining the samples using a filter whose support can be adapted to the desired resolution. Lee et al. [LRU85] and later Purgathofer [Pur87] have developed adaptive sampling algorithms which determine whether or not local oversampling is necessary based on the variance of the estimate of the pixel values. In [Mit87], Mitchell has proposed a two-level sampling method which generates point sets having spectral characteristics similar to those of blue noise point sets. As refinement criterion he uses a contrast metric based on human visual perception [Cae81], which operates per RGB channel. These works present complete solutions to the adaptive sampling problem. A complete overview of the recent advances in adaptive sampling and reconstruction for Monte Carlo rendering can be found in [ZJL*15]. Several other works focus only on the refinement criterion. Rigau et al. [RFS02] propose the use of *Shannon entropy* which can be viewed as a measure of the uncertainty associated with a random variable. The same authors, in [RFS03], further combine the samples' color and geometry entropies to de-

DOI: 10.2312/egsh.20161016

rive a contrast metric. Xu et al. [XBZ*05, XFSS07, XSFS10] later proposed using *Tsallis entropy* instead. All the above criteria rely on the sampled values to compute the uncertainty of the estimate and use it to decide on whether to refine. Estimating the illumination integral at each shading point generally requires much fewer samples than reconstructing incident radiance functions at the image plane. There are however obvious similitudes w.r.t. adaptive sampling for both processes. A refinement criterion is required to decide, based on the information brought by the samples already taken, whether a local higher sampling rate is required, an effective distribution of samples is also required, which aims at maximizing the information brought by the new samples.

2.2. Bayesian Monte Carlo

Bayesian Monte Carlo is a relatively recent method for estimating the value of complex integrals [O'H91], which can be seen as the Bayesian counterpart of the classic Monte Carlo (CMC) method. Rasmussen and Ghahramani [RG02] have shown that BMC outperforms classic Monte Carlo with importance sampling thanks to the use of prior knowledge for careful sample placement and weighting. However, the increased complexity of BMC w.r.t. CMC prevented a direct application to rendering. Brouillat et al. [BBL*09] pioneered the use of BMC for rendering diffuse surfaces. More recently, Marques et al. [MBR*13a] generalized the use of BMC for rendering by presenting a SG-based framework for BMC which supports any material whose BRDF can be modeled as a sum of SG functions. In [MBSB15], Marques et al. show that BMC outperforms the state of the art QMC techniques for hemispherical integration applied to illumination integrals. These results confirmed the potential of BMC for efficiently solving the rendering integral. However, the use of BMC in rendering is still in an incipient phase and its application to more evolved and widely used algorithms remains cumbersome. In this paper we propose an algorithm for applying BMC to a two-level adaptive sampling scheme, where the result of the first level estimate is injected as prior knowledge for a second level estimate. Furthermore, we show that our adaptive BMC algorithm outperforms adaptive QMC.

3. Adaptive Bayesian Monte Carlo

In the following we formalize the mechanism of a two-level adaptive sampling BMC estimate and propose an efficient approach for computing the resulting quadrature rule.

3.1. Background

The Bayesian approach to Monte Carlo integration is to estimate the value of an integral of the form:

$$I = \int f(\mathbf{x})p(\mathbf{x})d\mathbf{x}\,, \qquad \text{with } \mathbf{x} \in \mathbb{R}^D, \tag{1}$$

where $p(\mathbf{x})$ is the analytically known part of the integrand, and $f(\mathbf{x})$ is the unknown part of the integrand. For illumination integrals $f(\mathbf{x})$ is usually the incident radiance function, $p(\mathbf{x})$ is the product between the cosine term and the BRDF, and $\mathbf{x} = (\theta, \phi)$ is an incident direction on the unit hemisphere where θ and ϕ are the polar and azimuthal angles, respectively. In a Bayesian approach, the unknown function $f(\mathbf{x})$ is represented through a probabilistic model, usually a Gaussian Process model (GP), which is a stochastic representation of our knowledge about $f(\mathbf{x})$ before any samples are drawn. A GP is completely defined by a mean function $\bar{f}(\mathbf{x})$ and a covariance function $k(\mathbf{x},\mathbf{x}')$, which must be positive definite. The *prior* GP model of $f(\mathbf{x})$ is denoted as $f(\mathbf{x}) \sim \mathcal{GP}_1[\bar{f}_1(\mathbf{x}), k_1(\mathbf{x},\mathbf{x}')]$. $\bar{f}_1(\mathbf{x}) = \mathrm{E}[f(\mathbf{x})]$ can be thought of as a rough approximation of $f(\mathbf{x})$ before any samples are drawn, whereas $k_1(\mathbf{x},\mathbf{x}') = \mathrm{cov}[f(\mathbf{x}), f(\mathbf{x}')]$ characterizes the smoothness of the prior GP model. The sampling stage conditions $\mathcal{GP}_1$ to a set of n_1 noisy samples (i.e., observations) $\mathcal{D}_1 = \{(\mathbf{x}_1^i, Y_1^i) \,|\, i = 1,\ldots,n_1\}$ with $Y_L^i = f(\mathbf{x}_L^i) + \varepsilon_L^i$, ε_L^i being an *i.i.d.* Gaussian noise with mean 0 and variance σ_n^2. This results in a new GP model of $f(\mathbf{x})$, called *posterior*, which incorporates the information brought by the samples. It is denoted $\mathcal{GP}_2[\bar{f}_2(\mathbf{x}), k_2(\mathbf{x},\mathbf{x}')]$, where $\bar{f}_2(\cdot)$ and $k_2(\cdot,\cdot)$ are the posterior mean and covariance functions, respectively, and are given by [RG02]:

$$\begin{aligned} \bar{f}_2(\mathbf{x}) &= \mathrm{E}[f(\mathbf{x})|\mathcal{D}_1] = \bar{f}_1(\mathbf{x}) + \mathbf{k}_1^t Q_1^{-1}(\mathbf{Y}_1 - \bar{\mathbf{F}}_1) && (2)\\ k_2(\mathbf{x},\mathbf{x}') &= \mathrm{cov}[f(\mathbf{x}), f(\mathbf{x}')|\mathcal{D}_1] \\ &= k_1(\mathbf{x},\mathbf{x}') - \mathbf{k}_1^t(\mathbf{x}) Q_1^{-1} \mathbf{k}_1(\mathbf{x}'), && (3) \end{aligned}$$

where

$$\begin{aligned} \mathbf{k}_1(\mathbf{x}) &= (k_1(\mathbf{x}_1^1,\mathbf{x}),\ldots,k_1(\mathbf{x}_1^{n_1},\mathbf{x}))^t \\ Q_1 &= (K_1 + \sigma_n^2 I_n) \\ K_1^{(i,j)} &= k_1(\mathbf{x}_1^i,\mathbf{x}_1^j) \qquad \text{with } (i,j) \in [1,n_1]^2 \\ \mathbf{Y}_1 &= (f(\mathbf{x}_1^1) + \varepsilon_1^1,\ldots,f(\mathbf{x}_1^i) + \varepsilon_1^{n_1}) \\ \bar{\mathbf{F}}_1 &= (\bar{f}_1(\mathbf{x}_1^1),\ldots,\bar{f}_1(\mathbf{x}_1^{n_1}))\,. \end{aligned}$$

In practice, under a squared exponential stationary covariance function, the Bayesian estimate $\hat{I}_1$ of I amounts to computing the value of Eq. (1) by replacing the unknown function $f(\mathbf{x})$ by its *maximum a posteriori* estimate according to the posterior model $\mathcal{GP}_2$, i.e., $\bar{f}_2(\mathbf{x})$. This is achieved by computing the Bayesian quadrature equations [RG02]:

$$\hat{I}_1 = \bar{I} + \mathbf{z}_1^t Q_1^{-1}(\mathbf{Y}_1 - \bar{\mathbf{F}}_1) \tag{4}$$

where

$$\bar{I} = \int \bar{f}_1(\mathbf{x})p(\mathbf{x})d\mathbf{x} \quad \text{and} \quad \mathbf{z}_1 = \int \mathbf{k}_1(\mathbf{x})p(\mathbf{x})d\mathbf{x}\,.$$

For further details see [MBR*13a, MBSB15]. Note that $\mathbf{z}_1$ is a vector with as many integrals as the number of samples n_1. Also note that the covariance matrix Q_1 (which accounts for the relative positions between samples) has to be inverted for each integral estimate. To make BMC practical for rendering purposes, these operations must be computed efficiently. In [MBR*13a], the authors propose a SG-based solution for a non-adaptive BMC estimate. Both $p(\mathbf{x})$ and $k_1(\mathbf{x},\mathbf{x}')$ are modeled through SG functions and the location of the samples is known before rendering time, which allows Q_1^{-1} to be precomputed and each integral $\mathbf{z}_1$ to be evaluated through a simple query to a precomputed 2D table of spherical Gaussian integrals (SGI, see Eq. (22) of [MBR*13a]). In the following we maintain the same assumptions and show how the same framework can be used to compute a two-level BMC estimate.

3.2. Refining a Bayesian Monte Carlo estimate

Adaptive sampling requires that after computing $\hat{I}_1$, and if some refinement criterion is met, a second set $\mathcal{D}_2$ of n_2 noisy samples of $f(\mathbf{x})$ is collected, such that:

$$\mathcal{D}_2 = \{(\mathbf{x}_2^i, Y_2^i) \,|\, i = 1,\ldots,n_2\}, \quad \text{with} \quad Y_2^i = f(\mathbf{x}_2^i) + \varepsilon_2^i,$$

The second level BMC estimate $\hat{I}_2$ is computed using as prior the posterior $\mathcal{GP}_2$ resulting from the first level estimate, yielding:

$$\hat{I}_2 = \bar{I}_2 + \mathbf{z}_2^t Q_2^{-1}(\mathbf{Y}_2 - \bar{\mathbf{F}}_2) \tag{5}$$

where

$$\begin{aligned}
\bar{I}_2 &= \hat{I}_1 \\
\mathbf{z}_2 &= \int \mathbf{k}_2(\mathbf{x}) p(\mathbf{x}) d\mathbf{x} \\
\mathbf{k}_2(\mathbf{x}) &= (k_2(\mathbf{x}_2^1,\mathbf{x}),\ldots,k_2(\mathbf{x}_2^{n_2},\mathbf{x}))^t \\
Q_2 &= (K_2 + \sigma_n^2 I_n) \\
K_2^{(i,j)} &= k_2(\mathbf{x}_2^i,\mathbf{x}_2^j) \qquad \text{with } (i,j) \in [1,n_2]^2 \\
\bar{\mathbf{F}}_2 &= (\bar{f}_2(\mathbf{x}_2^1),\ldots,\bar{f}_2(\mathbf{x}_2^{n_2}))
\end{aligned}$$

While the terms of the Bayesian quadrature rule of level two in Eq. (5) have roughly the same role as presented in Eq. (4), the covariance function k_2 is no longer stationary and has a more complex expression as can be seen in Eq. (3). Consequently, $\mathbf{z}_2$ can no longer be computed through a simple query to a 2D table of SGI as in [MBR*13a].

3.3. Efficient quadrature computation

Given our assumption that the sample's locations are known before rendering, Q_2^{-1} can be precomputed to accelerate BMC integration. Nevertheless calculating the vector of integrals $\mathbf{z}_2$ is still the most computationally intensive operation. Using Eq. (3), the i^{th} element of the $\mathbf{z}_2$ vector from Eq. (5) can thus be expressed as:

$$z_2^i = \underbrace{\int k_1(\mathbf{x}_2^i,\mathbf{x}')\,p(\mathbf{x}')\,d\mathbf{x}'}_{\text{SGI}} - \overbrace{\int \underbrace{\mathbf{k}_1^t(\mathbf{x}_2^i)\,Q_1^{-1}}_{\text{constant}} \mathbf{k}_1(\mathbf{x}')\,p(\mathbf{x}')\,d\mathbf{x}'}^{T_2^i}. \tag{6}$$

Note that the leftmost integral of Eq. (6) is a spherical Gaussian integral (SGI, see Eq. (22) of [MBR*13a]). Its value can thus be quickly computed within the SG framework of [MBR*13a], by resorting to a compact precomputed 2D entry table. Recall that this table is only precomputed once, so there is no need to precompute it again for the adaptive BMC case. In the rightmost integral of Eq. (6) T_2^i, two terms are constant since they are not dependent on the integration variable $\mathbf{x}'$. These terms can thus be taken out of the integral and T_2^i can be rewritten as:

$$T_2^i = \mathbf{k}_1^t(\mathbf{x}_2^i)\,Q_1^{-1} \underbrace{\int \mathbf{k}_1(\mathbf{x}')\,p(\mathbf{x}')\,d\mathbf{x}'}_{\mathbf{z}_1} = \mathbf{k}_1^t(\mathbf{x}_2^i)\,Q_1^{-1}\,\mathbf{z}_1\,,$$

where $\mathbf{z}_1$ is the vector of integrals computed for the first level estimate. The term $\mathbf{k}_1(\mathbf{x}_2^i)$ is a vector containing the covariance $(k_1(\mathbf{x}_2^i,\mathbf{x}_1^1),\ldots,k_1(\mathbf{x}_2^i,\mathbf{x}_1^{n_1}))$ of the new sample $\mathbf{x}_2^i$ with each of the samples $\mathbf{x}_1^j, j \in \{1,\ldots,n_1\}$ of the first level and can be precomputed since the locations of all samples are known *a priori*. Moreover, both Q_1^{-1} and $\mathbf{z}_1$ have been computed for the previous level and can be reused for computing $\mathbf{z}_2$, hence accelerating the second level estimate.

4. Results

In this section we evaluate the efficiency of our adaptive BMC approach, by comparing it to adaptive QMC for estimating the value of the illumination integral under its hemispherical formulation. The results were generated with the Mitsuba raytracer [Jak10]. For simplicity, the values of the samples are computed from an environment map. The extension to global illumination will be addressed in future work. Reference images are computed using importance sampling. The used oversampling criterion is the samples' color entropy [RFS03].

Fig. 1 (a) to (c) shows a test case used to assess the quality of the images generated using adaptive BMC and adaptive QMC. The images shown in (b) and (c) have been generated using 40 sample directions for the first sampling level and 40 for the second level, making a total of 80 samples for the oversampled zones. The sample set is generated using a Sobol (0,2)-sequence and are the same for both methods. Fig. 1 (c) shows a sampling map. The zones selected for oversampling are shown in red, while the zones where only 40 samples are taken are shown in green. The close-up views on Fig. 1 (b) show a comparison between adaptive BMC and adaptive QMC for a zone of the image which has been oversampled. It can be seen that the BMC-based image is closer to the reference than the QMC-based image and exhibits fewer high frequency noise. The RMSE of the oversampled zone for QMC is 6.83% higher than that of BMC for the same zone. As regards the computing time, both methods rendered the full images associated with Fig. 1 (b) roughly within the same time (6.6s).

The RMSE curves shown in Fig. 1 (d) have been calculated by only taking into account the oversampled zone of the image shown in red in Fig. 1 (c). The results were generated using 10 samples for the first level of sampling and $n_2 \in \{10,30,70\}$ samples for the second level. The lines in the plot correspond to the following methods:

- adaptive BMC and QMC using a Sobol sequence [Sob67];
- adaptive BMC and QMC using spherical Fibonacci (SF) point sets [MBR*13b];
- adaptive BMC using optimized sample sets optimized independently for the number of samples of each sampling level following [MBR*13a].

Note that the sample sets used in adaptive BMC with optimal sample sets are independently optimized for each sampling level, meaning that the sample positions are not optimal as a whole since the second level sample set is produced without taking into account the sample set from the first level. To alleviate the effect of superposing samples of different levels, the sample set of level two is rotated by an angle of π w.r.t. the first level sample set. The same approach is also taken when using the spherical Fibonacci point sets. The results show that adaptive BMC outperforms adaptive QMC when using the same sample sets. For example, the RMSE of adaptive BMC is consistently smaller than that of QMC (roughly 5% lower) when both methods use a Sobol $(0,2)$-sequence. The results also show that Fibonacci point sets and optimized sample sets seem to have an advantage over a Sobol sequence, despite the hierarchical character of the latter point system. Nevertheless, the use of non-hierarchical point sets should be progressively more and more inefficient as the number of sampling levels increases. This is because the samples from high sampling levels would be more likely to overlap with samples from previous levels. The computing times are once again roughly the same for QMC and BMC, BMC being slightly slower in the case where $n_1 = 10$ and $n_2 = 70$ (6.7s for BMC and 6.6s for QMC). Note that this difference is totally neg-

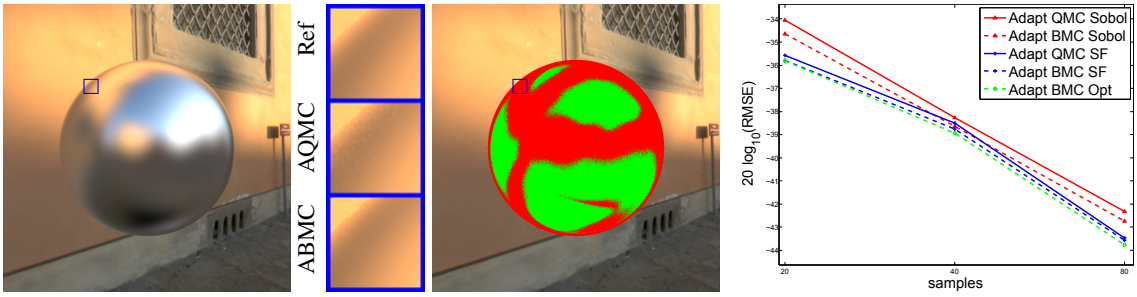

Figure 1: *Experimental results. (a) reference image. (b) close-up views comparison between adaptive BMC (ABMC) and adaptive QMC (AQMC) for a zone of the image which has been oversampled. (c) sampling map: the zones selected for oversampling are colored in red, while the zones not oversampled are in green. (d) RMSE for ABMC and AQMC using different sample sets. The value in the x-axis reflects the total number of samples for the oversampled zones.*

ligible in a real global illumination setting (such as, for example, final gathering for photon mapping) where sampling is much more costly than in our simple environment map-based test case.

5. Conclusion and future work

We proposed a two-level BMC adaptive sampling approach by deriving the second level quadrature rules and introducing an efficient quadrature computation method based on the spherical Gaussian framework of [MBR*13a]. We show with a simple test scene that adaptive BMC overcomes adaptive QMC using the same samples. This improvement is justified by the judicious use BMC makes of the information available prior to sampling. Future work will include extending the proposed approach to global illumination and explore what is the best proportion between samples allocated to the two sampling levels. Another interesting research line is to choose the samples sequentially, similarly to Huszár and Duvenaud [HD12]. Each new sample would be placed in the position of the integration domain where it decreases the most the variance of the estimate. Applied to adaptive BMC it entails starting from the optimal sample set of the first level and adding the new samples one by one until the desired number of samples for the second level is reached. The current approach's main limitation is the requirement to keep full information from all samples of previous sampling levels, which might prevent a generalization to a large number of sampling levels.

Acknowledgements This work has been supported by COMPETE: POCI-01-0145-FEDER-007043 and FCT - Fundação para a Ciência e Tecnologia within the Project Scope: UID/CEC/00319/2013.

References

[BBL*09] BROUILLAT J., BOUVILLE C., LOOS B., HANSEN C., BOUATOUCH K.: A Bayesian Monte Carlo Approach to Global Illumination. *Comp. Graph. Forum 28*, 8 (2009), 2315–2329. 2

[Cae81] CAELLI T.: *Visual Perception, Theory and Practice*. Franklin Book Co (June 1981), 1981. 1

[HD12] HUSZÁR F., DUVENAUD D.: Optimally-Weighted Herding is Bayesian Quadrature. In *Uncertainty in Artif. Intelligence.* (2012). 4

[Jak10] JAKOB W.: Mitsuba, 2010. http://www.mitsuba-renderer.org. 3

[LRU85] LEE M. E., REDNER R. A., USELTON S. P.: Statistically Optimized Sampling for Distributed Ray Tracing. *SIGGRAPH C. G. 19*, 3 (July 1985), 61–68. 1

[MBR*13a] MARQUES R., BOUVILLE C., RIBARDIÈRE M., SANTOS L. P., BOUATOUCH K.: A Spherical Gaussian Framework for Bayesian Monte Carlo Rendering of Glossy Surfaces. *IEEE Trans. on Visualization and Computer Graphics 19*, 10 (10 2013), 1619–1632. 1, 2, 3, 4

[MBR*13b] MARQUES R., BOUVILLE C., RIBARDIÈRE M., SANTOS L. P., BOUATOUCH K.: Spherical Fibonacci Point Sets for Illumination Integrals. *Computer Graphics Forum 32*, 8 (12 2013), 134–143. 3

[MBSB15] MARQUES R., BOUVILLE C., SANTOS L. P., BOUATOUCH K.: Efficient Quadrature Rules for Illumination Integrals: From Quasi Monte Carlo to Bayesian Monte Carlo. *Synthesis Lectures on Comp. Graphics and Animation 7*, 2 (2015), 1–92. 2

[Mit87] MITCHELL D. P.: Generating Antialiased Images at Low Sampling Densities. *SIGGRAPH C. G. 21*, 4 (Aug. 1987), 65–72. 1

[O'H91] O'HAGAN A.: Bayes-Hermite Quadrature. *J. Statist. Plann. Inference 29*, 3 (1991), 245–260. 2

[Pur87] PURGATHOFER W.: A Statistical Method for Adaptive Stochastic Sampling. *Comp. & Graph. 11*, 2 (1987), 157 – 162. 1

[RFS02] RIGAU J., FEIXAS M., SBERT M.: New Contrast Measures for Pixel Supersampling. In *Proceedings of CGI 02* (2002), Springer-Verlag London Limited, pp. 439–451. 1

[RFS03] RIGAU J., FEIXAS M., SBERT M.: Entropy-based Adaptive Sampling. In *Graphics Interface* (2003), pp. 149–158. 1, 3

[RG02] RASMUSSEN C. E., GHAHRAMANI Z.: Bayesian Monte Carlo. In *NIPS* (2002), MIT Press, pp. 489–496. 2

[Sob67] SOBOL I. M.: On the Distribution of Points in a Cube and the Approximate Evaluation of Integrals. *USSR Computational Math. and Math. Phys. 7* (1967), 86–112. 3

[XBZ*05] XU Q., BAO S., ZHANG R., HU R., SBERT M.: Adaptive Sampling for Monte Carlo Global Illumination Using Tsallis Entropy. In *Computational Intelligence and Security, International Conference, Proceedings, Part II* (2005), vol. 3802 of *Lecture Notes in Computer Science*, Springer, pp. 989–994. 2

[XFSS07] XU Q., FEIXAS M., SBERT M., SUN J.: A New Adaptive Sampling Technique for Monte Carlo Global Illumination. In *CAD/Graphics* (2007), pp. 191–196. 2

[XSFS10] XU Q., SBERT M., FEIXAS M., SCOPIGNO R.: A New Refinement Criterion for Adaptive Sampling in Path Tracing. In *Industr. Elect. (ISIE), 2010 IEEE Int. Symp.* (07/2010 2010), pp. 1556 –1561. 2

[ZJL*15] ZWICKER M., JAROSZ W., LEHTINEN J., MOON B., RAMAMOORTHI R., ROUSSELLE F., SEN P., SOLER C., YOON S.-E.: Recent Advances in Adaptive Sampling and Reconstruction for Monte Carlo Rendering. *Comp. Graph. Forum 34*, 2 (May 2015), 667–681. 1