

An Effective Stratified Sampling Scheme for Environment Maps with Median Cut Method

Xing Mei¹, Marc Jaeger^{1,2}, Baogang Hu¹

¹LIAMA/NLPR, Institute of Automation, Chinese Academy of Sciences, Beijing, China

²Laboratoire AMAP, CIRAD, Montpellier, France

xmei@nlpr.ia.ac.cn, jaeger@liama.ia.ac.cn, hubg@nlpr.ia.ac.cn

Abstract

Environment maps are extensively used as natural light sources in realistic rendering. We propose a stratified sampling scheme for environment maps by first stratifying the maps into a set of rectangular regions with Median Cut method, then estimating the contribution of the regions with Monte Carlo integration techniques. In this way, illumination, surface reflectance and spatial distribution are all taken into consideration for the generation of the light samples. Compared with the existing biased lighting techniques, the presented scheme produces unbiased rendering results with less noise and better shadow boundaries, particularly at low sampling rates. The proposed spatial distribution of the samples also helps to overcome the sample-clumping problem in traditional illumination-based importance sampling method. Experimental results indicate that the scheme is fast, simple to implement and effective.

1. Introduction

Realistic rendering synthetic objects illuminated by environment maps is a key issue in image synthesis. Following early seminar works in [2] [10], Debevec's study on capturing natural illumination with high dynamic range images [7] has enabled many new image-based lighting techniques for direct light computation, which can be classified into two categories: biased and unbiased.

Most of existing methods, falling into biased methods, try to stratify the environment maps into a set of regions and approximate these regions as directional lights. Different rules have been proposed for the stratification of the maps, such as k-means clustering[1][5], Lloyd's Relaxation method[11], and Penrose-based method[13]. With enough directional lights 'mimicking' the light environment, biased techniques produce nice noise-free images. However, if the number of the lights is relatively low, these techniques may

bring visual artifacts. Approximating each region as a directional light fails to consider the variation of the surface reflectance and orientation within the region. If the surface is highly glossy, inaccurate shading and high light will affect rendering quality. And unwanted banding near shadow boundaries appears when some directional lights are suddenly occluded by the scene geometry. This banding effect can be decreased by jittering the direction of the lights, but additional noise and bias are introduced into the rendering results. Moreover, the stratification process with biased methods is usually expensive.

Illumination-based importance sampling is another widely used technique [14]. This method gives unbiased rendering results with Monte Carlo estimators. Although noise is inevitable, it can be combined with Multiple Importance Sampling [16] and Bidirectional Importance Sampling [3] [15] for further noise reduction. Importance sampling techniques may suffer from the sample-clumping problem at low sampling rates (see section 3).

Our concern is to produce better shadow boundaries and more accurate high lights on surfaces than existing biased methods, and to decrease the effect of the sample-clumping problem in traditional importance sampling. We therefore propose an effective stratified sampling scheme for environment maps by combining the two kinds of techniques: first stratify the maps into regions with Median Cut method [8], and then apply sampling techniques on these regions. Basic concepts about Monte Carlo lighting techniques are given in Section 2. We describe the stratified sampling scheme in Section 3, and show the experimental results in Section 4. Conclusions are given in Section 5.

2. Monte Carlo Direct Lighting

Monte Carlo integration techniques [9] have been shown to be effective and robust for light computation. The direct radiance, L_o , at a surface point p with an outgoing direction

$\vec{\omega}_o$ can be computed by the rendering equation:

$$L_o(p, \vec{\omega}_o) = \int_{\Omega_{2\pi}} L_i(\vec{\omega}_i) f(p, \vec{\omega}_i, \vec{\omega}_o) (\vec{\omega}_i \cdot \vec{n}) V(p, \vec{\omega}_i) d\vec{\omega}_i \quad (1)$$

where $\Omega_{2\pi}$ is the hemisphere of directions above the surface point p , L_i is the incident radiance from direction $\vec{\omega}_i$, f is the Bidirectional Reflectance Distribution Function (BRDF), \vec{n} is the surface normal at p , and V is a visibility term. In image-based lighting, the incident radiance from all directions around p is given by the environment maps.

To evaluate a specific integral $I = \int_{\Omega} g(x) dx$, Monte Carlo estimator gives an approximation of I as $\hat{I} = \frac{1}{N} \sum_{i=1}^N \frac{g(x_i)}{p(x_i)}$, where x_1, \dots, x_N are N samples drawn from a Probability Density Function (PDF) $p(x)$ defined in Ω . The accuracy of the estimator (estimated by its theoretic variance) is closely related to the choice of the PDF $p(x)$. To minimize the variance, importance sampling technique suggests samples be drawn from a PDF proportional to the integrand $g(x)$. However, in the case of equation (1), it is difficult to sample directly from the integrand. Many techniques have been developed to sample from the L_i component (environment map) or the BRDF component.

Illumination-based importance sampling is an example of sampling from L_i . An environment map in latitude-longitude format can be seen as a two-dimensional illumination PDF. For a given map $W \times H$, each pixel (i, j) is mapped to an incident direction $\vec{\omega}_i = (\phi, \theta)$ (spherical coordinates) by a scaling factor $(2\pi/W, \pi/H)$. To draw a sample from this 2D distribution, a 2D Cumulative Distribution Function (CDF) is built from the PDF, and inverse sampling method can then be applied on the CDF. The probability of selecting pixel (i, j) (i.e. the corresponding incident direction $\vec{\omega}_i$) can be expressed as [14]:

$$p(\vec{\omega}_i) = p(i, j) \frac{WH}{2\pi^2 \sin\theta} \quad (2)$$

where $p(i, j) = \frac{I_P}{\sum_P I_P}$, I_P - light intensity for pixel (i, j) .

Sampling from only one component, illumination or the BRDF, may be inefficient if L_i or f does not fit the integrand well. For example, sampling from the environment maps introduces much noise on highly glossy surfaces. Many techniques have been proposed to reduce this variance. Multiple Importance Sampling (MIS), introduced by Veach and Guibas [16], combines the light samples and BRDF samples linearly with heuristic weights. Recent techniques sample from the product of illumination and BRDF functions, such as Bidirectional Importance Sampling (BIS) [3] [15] and Wavelet Importance Sampling (WIS) [4]. We focus here on BIS, since WIS needs an expensive preprocess to convert environment maps and measure BRDF data

into wavelet representation. To evaluate equation (1), BIS proceeds in two stages:

1. Draw M_0 'proposal' samples $X = \{x_1, \dots, x_{M_0}\}$ from L_i . The probability of choosing sample $x_i, p(x_i)$, can be evaluated from equation (2)
2. Draw M_1 ($M_1 \leq M_0$) 'final' samples $Y = \{y_1, \dots, y_{M_1}\}$ from X with probability proportional to the ratio $\frac{q(x_i)}{p(x_i)}$, $q(x) = L_i(x) f(p, x, \vec{\omega}_o) (x \cdot \vec{n})$.

When M_0 approaches infinity, the 'final' M_1 samples from Y can be seen as drawn from the PDF defined by $q(x)$. Monte Carlo estimator for equation (1) can then be adopted as [3]:

$$\hat{L}_o(p, \vec{\omega}_o) \approx \frac{\hat{L}_{ns}}{M_1} \sum_{i=1}^{M_1} V(p, y_i) \quad (3)$$

where \hat{L}_{ns} is an estimation of $L_{ns} := \int_{\Omega_{2\pi}} q(x) dx$, and V is the visibility term. When $M_1 = M_0 = 1$, BIS is equivalent to illumination-based importance sampling.

Note that importance sampling and BIS can be successfully applied on any rectangular sub-region of environment maps. And equations (2) (3) still hold for sub-regions.

3 Stratified Sampling Scheme

One potential problem with importance sampling is that, if the environment map contains many small bright lights, most of the light samples will clump in these areas due to their high sampling probability, with few or even none in dimmer areas. Therefore, the poor spatial distribution of the samples will cause observable noise at shading points where those bright lights are occluded by the geometry of the scene. Samples need to be distributed according to the light intensity, and also need to define a uniform spatial distribution over the whole environment map, especially when scene geometry is complex. We intend to constrain the position of the samples by a stratification process. And we also require the stratified regions being rectangles so that importance sampling techniques can be easily applied on them. In our stratified sampling scheme, Median Cut method satisfies these requirements.

3.1 Median Cut Method

Debevec [8] proposed a fast method called 'Median Cut Algorithm' to subdivide the environment map into a set of N rectangular regions with equal energy:

1. Push whole map into the region list as a single region.
2. Subdivide each region in the list along the longer dimension into two halves with equal energy.

3. Repeat Step 2 until number of the regions reaches N .

Each region's energy is defined as the sum of pixel intensity in that region. Computing the energy and finding the splitting position for each region can be accelerated by using a summed-area table [6]. Figure 1 shows the splitting result on Grace Cathedral environment map (<http://athens.ict.usc.edu/Probes/>). Compared with precited stratification methods, this method is fast and efficient, as we will show in Section 4.

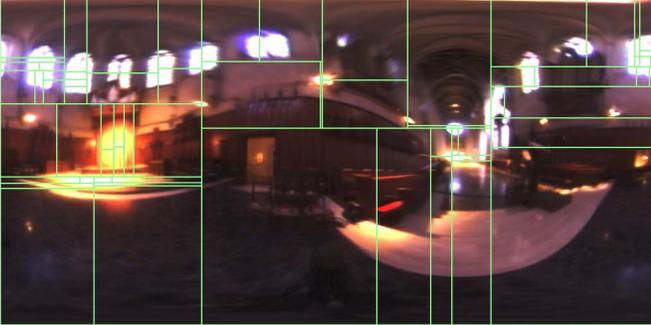


Figure 1. 64 regions on Grace Cathedral environment map by Median Cut method.

3.2 Stratified Sampling Scheme

After the stratification, we get N regions (N area lights), noted as $\Omega_1, \dots, \Omega_N$. Rendering equation (1) can then be expressed as the sum of the contribution from these lights:

$$\begin{aligned} L_o(p, \vec{\omega}_o) &= \sum_{i=1}^N L_{o_i} \\ &= \sum_{i=1}^N \int_{\Omega_i} L_i(\vec{\omega}_i) f(p, \vec{\omega}_i, \vec{\omega}_o) (\vec{\omega}_i \cdot \vec{n}) V(p, \vec{\omega}_i) d\vec{\omega}_i \end{aligned} \quad (4)$$

Approximating these regions into directional lights by pre-integrating the illumination has been proved to be an efficient biased method [8]. However, the rectangularity of the 2D regions make it quite straight to apply importance sampling techniques directly on these regions. This idea is quite similar with 'stratified sampling' variance reduction technique in traditional Monte Carlo methods [12]. Illumination-based sampling and corresponding equation (2) still hold for each region. After the stratification, the light samples from each region can not be linearly combined with BRDF samples because they may not be in the same domain, so MIS is not an effective choice for further variance reduction, while BIS's sampling-resampling framework still works for each region.

We propose the *stratified sampling scheme* as follows:

1. Preprocessing:

- Subdivide the environment map into N rectangular regions $\Omega_1, \dots, \Omega_N$ with Median Cut method.
- Build 2D CDFs for each region.

2. Sampling: For each region $\Omega_i (1 \leq i \leq N)$,

- Draw M 'proposal' samples $X_i = \{x_{i1}, \dots, x_{iM}\}$ on Ω_i with illumination-based importance sampling.
- Draw one 'final' sample y_i from X_i with BIS.

For the sampling step, totally $N(M + 1)$ proposal samples are generated for each shading point, but only N samples $Y = \{y_1, \dots, y_N\}$ are used for final visibility test. Combining equation (4) and (3), the outgoing radiance L_o at shading point p is expressed as:

$$L_o(p, \vec{\omega}_o) = \sum_{i=1}^N \hat{L}_{nsi} V(p, y_i) \quad (5)$$

where \hat{L}_{nsi} is a Monte Carlo estimator for

$$L_{nsi} := \int_{\Omega_i} L_i(\vec{\omega}_i) f(p, \vec{\omega}_i, \vec{\omega}_o) (\vec{\omega}_i \cdot \vec{n}) d\vec{\omega}_i \quad (6)$$

using M illumination-based samples X_i from Ω_i . Since BIS is applied on each region, $M = 1$ means $M_0 = M_1 = 1$, i.e. directly importance sampling on each region.

Stratified sampling scheme is more effective than traditional sampling techniques. It samples on each stratified region and not on the whole map. This can be seen as a 'restriction' on sampling positions to avoid sample-clumping problem. Median Cut subdivision also promises that more regions will be put in brighter areas. Therefore one sample per region strategy guarantees the samples' good distribution both on illumination and spatiality. BIS can further optimize the selection of the samples on each region by taking the BRDF component and surface orientation into account. Moreover, the stratification based on median cut algorithm and the CDF computation are fast enough to be included in the rendering process.

4 Experimental Results and Discussions

We implement our stratified sampling scheme and other lighting methods on PBRT, a convenient open source Monte Carlo ray tracer designed by M. Pharr and G. Humphreys [14]. We use the 1024×512 Grace Cathedral environment map, presented in Figure 2, with the Utah

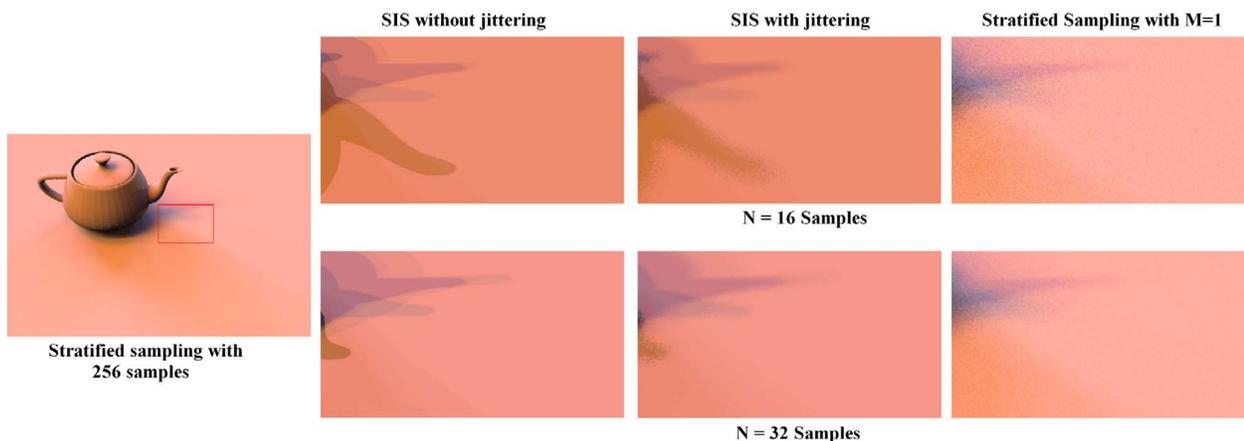


Figure 2. A teapot in Grace Cathedral light environment with SIS and stratified sampling. The image on the left is computed by stratified sampling with 256 samples. We render the region in the red square for the other six images. The top and bottom rows are computed with 16 and 32 samples per shading point respectively. Left Column: SIS without jittering. Middle Column: SIS with jittering. Right Column: Stratified Sampling with $M = 1$.

Teapot model (containing more than 6000 polygons) and the Stanford Happy Buddha model (containing more than 1 million polygons). All tests are performed on a 2.8GHz P4 PC. Image output resolution is 600×400 .

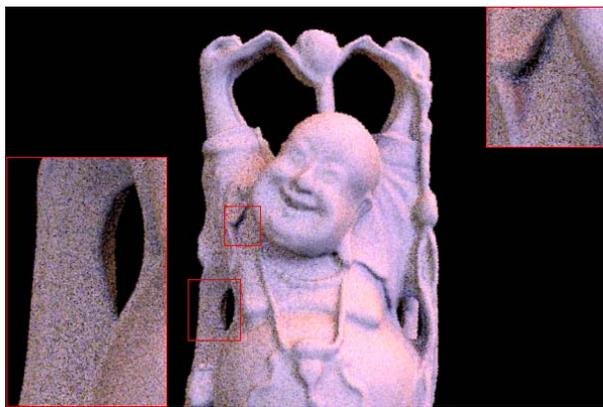
First, we use our sampling scheme on the teapot scene and compare it with Structured Importance Sampling (SIS), which is reported as an effective biased method [1]. For SIS, each region is approximated as a directional light with pre-integration, and both the rendering results with or without random 'jittering' (in light directions) are presented for comparison. For stratified sampling, we use a single light sample on each region ($M = 1$), since the simple teapot model carries a diffuse surface. As shown in Figure 2, SIS produces visible banding effect near shadow boundaries when sampling rate is relatively low (16 or 32). Although random jittering helps to 'blur' this banding effect, it still fails to retrieve the correct shadow boundaries. Stratified sampling gives much better soft shadow boundaries without introducing much noise even at 16 samples per shading point. Moreover, the time cost for the preprocess in stratified sampling scheme (32 regions) is drastically lower; it takes about 0.10 seconds, while SIS preprocess takes about 20 seconds.

The second test on the Buddha scene compares our stratified sampling ($M = 1$) with traditional illumination-based importance sampling. We use low-discrepancy sequences as described in [14] to generate random samples in importance sampling. Rendering results with 16 samples per shading point are presented in Figure 3. The Buddha model has complex scene geometry and surface orienta-

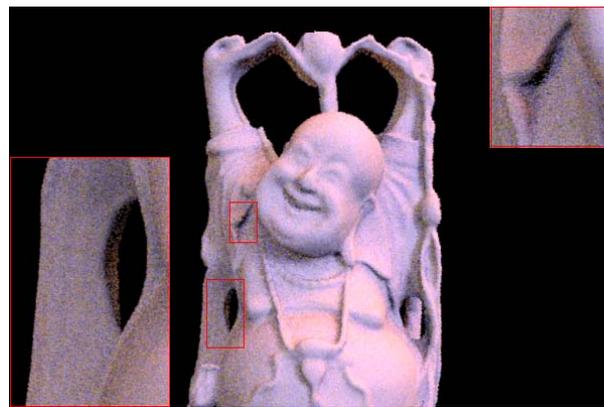
tions. Illumination-based sampling suffers from the sample-clumping problem and introduces significant noise at surface points where bright light sources are occluded by the geometry of the model (marked by red rectangles). Stratified sampling, with its better spatial distribution of the samples, reduces the error in such regions. For this test, the rendering time with importance sampling and stratified sampling is 161.4 seconds and 145.1 seconds respectively. On the one hand, for a light sample, averagely speaking, importance sampling needs to be computed on the whole map, while our method ($M = 1$) samples only on a stratified region. But on the other hand, stratified sampling needs context changes (loading new CDF, for instance) when moving to a new region. Implementation was not optimized on this point so far, so better performance could be gained in future work.

We also test the scheme on the Buddha scene at higher sampling rates. Figure 4 compares the *RMS* image error of the three sampling techniques (SIS included). The reference image was rendered by standard importance sampling with 10,000 samples per shading point. Stratified sampling presents the lowest *RMS* error at all sampling rates, and this improvement is more effective at low sampling rates.

We finally render the Buddha model with a high glossy BRDF (Cook-Torrance material with a specular exponent of 100), to test the effect of applying BIS in the scheme. As shown in Figure 5, SIS with 32 samples gives nice static rendering result, but fails to catch part of the high light. See for example the hands and the arms. Importance sampling on each region ($M = 1$) introduces significant noise on



(a) Illumination-based importance sampling with $N=16$



(b) Stratified sampling with $N = 16$ and $M = 1$

Figure 3. A Buddha model in Grace Cathedral light environment with illumination-based importance sampling and stratified sampling. Both methods use 16 samples per shading point. For surface points where most bright light sources are occluded (marked by red rectangles), stratified sampling significantly reduces the noise.

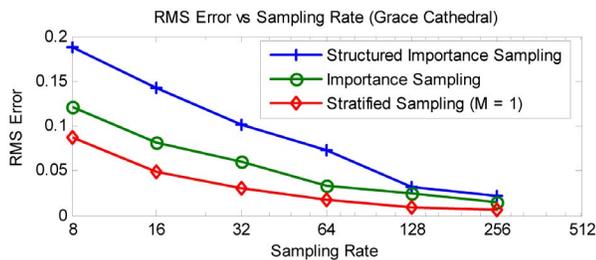


Figure 4. RMS Error vs Sampling Rate for the Buddha scene in Grace Cathedral with three methods. Stratified sampling gives the lowest error at all sampling rates.

such high glossy surfaces. By applying BIS on each region ($M = 50$), stratified sampling greatly reduces the noise and catches more high light, since surface reflectance and orientation are both considered for sample selection. As stated in [3], BIS is a practical variance reduction method for sampling techniques when the scene geometry is complex. On such scenes, ray-object intersection tests (visibility tests) are usually more expensive than sample generation. The Buddha model contains more than 1 million faces. In our experiment, the average time ratio of a visibility test (accelerated by a kd-tree data structure) to sample generation reaches 7.36 : 1. In such case, increasing the value of M might be more effective than increasing the number of the shadow rays for each shading point (N).

5 Conclusions

This paper proposes a novel combination of the lighting methods with environment maps to produce better shadow boundaries and more accurate high light than existing biased methods, and to decrease the effect of the sample-clumping problem in classical importance sampling.

We present an effective stratified sampling scheme by combining a fast Median Cut stratification method with Monte Carlo integration techniques. Importance sampling and BIS can be easily applied on the stratified rectangular sub-regions of the environment maps. Experimental tests show the following benefits gained from both techniques: 1. Importance sampling on each region produces better yet realistic soft shadow boundaries than existing biased methods. 2. Median Cut stratification helps to overcome the sampling-clumping problem in traditional importance sampling. Compared with structured importance sampling and importance sampling, our scheme presents the lowest RMS error at all sampling rates. 3. By applying BIS (increasing the number of the proposal samples) on each region, stratified sampling catches more high light than biased methods, and greatly reduces the rendering noise, since surface reflectance and orientation are both considered for sample selection. 4. The stratification process is fast and simple to implement.

The proposed scheme is slightly faster than importance sampling (with one proposal sample per region). Computational performance improvements are still possible; for instance, region context changes are not optimized so far. For complex glossy models, the balance between the number



(a) SIS with 32 samples



(b) Stratified sampling with $N = 32, M = 1$



(c) Stratified sampling with $N = 32, M = 50$

Figure 5. A glossy Buddha in Grace Cathedral light environment. SIS with 32 samples gives nice static rendering result, but fails to catch some of the high light. See for example the hands and arms. Importance sampling on each region ($M = 1$) in stratified sampling brings significant noise for such high glossy surfaces. By applying BIS on each region ($M = 50$), stratified sampling greatly reduces the noise and catches more high light at 32 final samples per shading point, since the surface orientation and the glossy BRDF component are both considered for sample selection.

of the regions and the number of proposal samples on each region is an interesting topic to be further explored.

Acknowledgements

This work is supported by LIAMA project 98-01, and NSFC project #60073007.

References

- [1] S. Agarwal, R. Ramamoorthi, S. Belongie, and H. W. Jensen. Structured importance sampling of environment maps. *ACM Trans. Graph.*, 22(3):605–612, 2003.
- [2] J. Blinn and M. Newell. Texture and reflection in computer generated images. *Communication of the ACM*, 19:542–546, 1976.
- [3] D. Burke, A. Ghosh, and W. Heidrich. Bidirectional importance sampling for direct illumination. In *EGSR '05: Eurographics Symposium on Rendering 2005*, pages 147–156, 2005.
- [4] P. Clarberg, W. Jarosz, T. Akenine-Möller, and H. W. Jensen. Wavelet importance sampling: efficiently evaluating products of complex functions. *ACM Trans. Graph.*, 24(3):1166–1175, 2005.
- [5] J. Cohen and P. Debevec. Lightgen hdrshop plugin. <http://www.ict.usc.edu/graphics/HDRShop/lightgen/>, 2001.
- [6] F. C. Crow. Summed-area tables for texture mapping. In *SIGGRAPH '84: Proceedings of the 11th annual conference on Computer graphics and interactive techniques*, pages 207–212, 1984.
- [7] P. Debevec. Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *SIGGRAPH '98: Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, pages 189–198, 1998.
- [8] P. Debevec. A median-cut algorithm for light-probe sampling. In *SIGGRAPH '05 Poster*, 2005.
- [9] P. Dutré, H. W. Jensen, J. Arvo, K. Bala, P. Bekaert, S. Marschner, and M. Pharr. State of the art in monte carlo global illumination. In *SIGGRAPH 2004 course notes*, 2004.
- [10] N. Greene. Environment mapping and other applications of world projections. *IEEE Computer Graphics & Applications*, 6(11):21–29, 1986.
- [11] T. Kollig and A. Keller. Efficient illumination by high dynamic range images. In *EGRW '03: Proceedings of the 14th Eurographics workshop on Rendering*, pages 45–50, 2003.
- [12] J. S. Liu. *Monte Carlo strategies in scientific computing*. Springer, 2001.
- [13] V. Ostromoukhov, C. Donohue, and P. M. Jodoin. Fast hierarchical importance sampling with blue noise properties. *ACM Trans. Graph.*, 23(3):488–495, 2004.
- [14] M. Pharr and G. Humpherys. *Physically-based rendering: from theory to implementation*. Morgan Kaufmann, 2004.
- [15] J. Talbot, D. Cline, and P. Egbert. Importance resampling for global illumination. In *EGSR '05: Eurographics Symposium on Rendering 2005*, pages 139–146, 2005.
- [16] E. Veach and L. J. Guibas. Optimally combining sampling techniques for monte carlo rendering. In *SIGGRAPH '95: Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 419–428, 1995.