

Second-Order Approximation for Variance Reduction in Multiple Importance Sampling

H. Lu^{†1}

Inria - Univ. Bordeaux/LaBRI


R. Pacanowski^{‡1}

CNRS/LP2N

X. Granier^{§1}

IOGS/LP2N

1. MANAO: Inria Bordeaux Sud-Ouest - LP2N (Univ. Bordeaux, IOGS, CNRS) - LaBRI (Univ. Bordeaux, CNRS)



Low to high glossy materials with five diffuse area light sources Glossy materials with high-frequency environment map lighting
Figure 1: Our per-pixel second-order approximation of the variance leads to a new and automatic approach for balancing the number of samples between two different sampling strategies. Except for light sources, the inset images show the sample distribution for each pixel. The yellow corresponds to the default balance heuristic strategy [Vea98]. Compared to the balance heuristic, the variance is reduced by (Left) 26% and (Right) 20% in average (14% and 11% for the standard deviation).

Abstract

Monte Carlo Techniques are widely used in Computer Graphics to generate realistic images. Multiple Importance Sampling reduces the impact of choosing a dedicated strategy by balancing the number of samples between different strategies. However, an automatic choice of the optimal balancing remains a difficult problem. Without any scene characteristics knowledge, the default choice is to select the same number of samples from different strategies and to use them with heuristic techniques (e.g., balance, power or maximum). In this paper, we introduce a second-order approximation of variance for balance heuristic. Based on this approximation, we introduce an automatic distribution of samples for direct lighting without any prior knowledge of the scene characteristics. We demonstrate that for all our test scenes (with different types of materials, light sources and visibility complexity), our method actually reduces variance in average. We also propose an implementation with low overhead for offline and GPU applications. We hope that this approach will help developing new balancing strategies.

1 Motivation and Previous Work

The computation of physically realistic shading [Kaj86] for a position \mathbf{p} viewed from direction \mathbf{o} relies mostly on the estimation of the following integral:

$$L(\mathbf{p} \rightarrow \mathbf{o}) = \int_{\Omega} \rho(\mathbf{o}, \omega) \cos \theta V(\mathbf{p}, \omega) L(\omega \rightarrow \mathbf{p}) d\omega$$

where Ω is the unit hemisphere, $\rho(\mathbf{o}, \omega)$ is the reflectance function, $\cos \theta$ is the cosine factor, and $L(\omega \rightarrow \mathbf{p})$ represents

the radiance coming from direction ω . For the sake of simplicity, we denote the whole integrand by

$$f(\mathbf{p}, \mathbf{o}, \omega) = \rho(\mathbf{o}, \omega) \cos \theta V(\mathbf{p}, \omega) L(\omega \rightarrow \mathbf{p}).$$

The main trend to compute such integral is to use Monte Carlo methods where the estimation of a reflected radiance $L(\mathbf{p} \rightarrow \mathbf{o})$ is computed from a set of N samples:

$$L(\mathbf{p} \rightarrow \mathbf{o}) \approx \frac{1}{N} \sum_{n=1}^N \frac{f(\mathbf{p}, \mathbf{o}, \omega_n)}{p(\omega_n)} \quad (1)$$

In Equation (1), $p(\omega_n)$ is a Probability Density Function (PDF), the closer it matches the numerator of the Equation the lower the number of samples N will be required to con-

[†] e-mail: heqi.lu@inria.fr

[‡] e-mail: romain.pacanowski@institutoptique.fr

[§] e-mail: xavier.granier@institutoptique.fr

verge to the solution. Although the best matching is reached when p is proportional to $f(\mathbf{p}, \mathbf{o}, \boldsymbol{\omega})$, it is hard to obtain an exact integrable analytical form in most cases.

A lot of work has been dedicated to derive an efficient sampling strategy for each component, or a subset, of the function $f(\mathbf{p}, \mathbf{o}, \boldsymbol{\omega})$ based on light [ODJ04, HSK*05, EKÖ12], or BRDF [AP07] or BRDF corrected by the cosine factor [PCS*12]. For visibility, since no analytical form exists, precomputation [GH06] is generally required. In general, a product cannot be accurately approximated with only one of its components.

A better PDF can be obtained by an approximation of the whole integrand. Bidirectional Importance Sampling techniques [BGH05, TCE05] approximate this product on a preliminary sampling (either BRDF-based or light-based): a weight is computed for each sample according to the shading functions. Then, a subset of these samples is selected by using these weights. A potential problem is due to the over-sampling needed to perform the second-step (resampling): some generated samples will not be used for the final estimator. Wang et al. [WA09] improved the techniques by combining lightcuts [WFA*05] with a BRDF-based sampling to reduce the initial bias caused by the initial distribution. Rouselle et al. [RCL*08] have extended these techniques by taking into account visibility.

More generally, precomputing conservative approximations of lighting environment and BRDF helps computing dynamically PDFs based on the product of the lighting and BRDF. For this purpose, Cline et al. [CETC06] split the environment map according to BRDF peaks computed from BRDF-based importance sampling. To be more conservative, Clarberg et al. [CAM08b] use a wavelet representation of the environment map, a wavelet approximation of the BRDF, and a fast computation of their product. These techniques require dedicated representations to support a per-pixel dynamic computation of the importance sampling structure. Although photon maps can also be used (e.g., [PBPP11]), in this paper we focus on a technique that does not require any additional storage.

Compared to these approaches, Multiple Importance Sampling (MIS) [Vea98] is still a simple and efficient way to combine different sampling techniques by first balancing the number of samples between the different strategies and second combining them into a weighted estimator (balance, power, maximum heuristics [Vea98] or α -max heuristic [GKPS12]). Despite the average improvement over the use of a unique strategy, the optimal balancing may differ from pixel to pixel in an unpredictable way [OZ00]. To reduce this problem, Pajot et al. [PBPP11] have introduced the empirical notion of *representativity* that estimates the concentration information for BRDF or photon maps. Intuitively, narrow lobe-shaped PDFs will request more samples. General guidelines to develop new representativity functions are introduced, but problems may occur when multiple concentrations exist such as in environment maps. Therefore,

our work differs on the following points. First, we do not need to introduce a dedicated representativity function when introducing a new sampling strategy. Second, we do not make any assumption on the shape of the PDFs, which is difficult to achieve on generic environment maps. Third, our balancing is completely done per-pixel and takes into account the local shading configuration.

Contributions In this paper, we introduce a second-order approximation for variance reduction when using MIS approaches. We demonstrate that our approach leads to a new and automatic way to distribute samples between different strategies. We show that for our test scenes and for direct lighting, our approximation also improves the average variance, validating its accuracy. We also demonstrate that we can use our approach for offline rendering with a better efficiency compared to previous balancing techniques. Its accuracy also permits to visualize what should be the dominant strategy for future investigations.

2 Theoretical Variance Reduction

To estimate the integral of a function $L = \int f(\boldsymbol{\omega}) d\boldsymbol{\omega}$ using random samples $\{\boldsymbol{\omega}_{i=B|L, n=1..N_i}\}$ from two PDFs p_B (for BRDF-based strategy to generate N_B samples) and p_L (for light-based strategy to generate N_L samples), Veach [Vea98] has introduced the following MIS primary estimator

$$L_{N_L, N_B} = \sum_{i=B|L} \frac{1}{N_i} \sum_{n=1}^{N_i} w_i(\boldsymbol{\omega}_{i,n}) \frac{f(\boldsymbol{\omega}_{i,n})}{p_i(\boldsymbol{\omega}_{i,n})} \quad (2)$$

where $w_i(\boldsymbol{\omega}_{i,n})$ is a weighting function. For balance heuristic,

$$w_i(\boldsymbol{\omega}_{i,n}) = \frac{N_i p_i(\boldsymbol{\omega}_{i,n})}{N_B p_B(\boldsymbol{\omega}_{i,n}) + N_L p_L(\boldsymbol{\omega}_{i,n})}.$$

By denoting $N = N_B + N_L$ (N is the total number of samples) with $N_B = \alpha N$, $N_L = (1 - \alpha)N$, $\alpha \in [0, 1]$ and as pointed by Georgiev et al. [GKPS12], this estimator is equivalent to the secondary estimator

$$L_{N, \alpha} = \frac{1}{N} \sum_{n=1}^N \frac{f(\boldsymbol{\omega}_n)}{\alpha p_B(\boldsymbol{\omega}_n) + (1 - \alpha) p_L(\boldsymbol{\omega}_n)}. \quad (3)$$

This is an estimator for the Defensive Importance Sampling strategy [Hes95] where a set of N samples $\boldsymbol{\omega}_n$ is generated according to $\alpha p_B + (1 - \alpha) p_L$. In this paper, we focus on minimizing the variance of this estimator to improve the sampling quality.

2.1 Second-Order Approximation of Variance

Finding the best α that minimizes the variance of $L_{N, \alpha}$ (i.e., $V[L_{N, \alpha}]$) is equivalent to minimize the sampling variance of the strategy (i.e., $V[L_{1, \alpha}]$) since

$$V[L_{N, \alpha}] = \frac{1}{N} V[L_{1, \alpha}].$$

If $\alpha p_B + (1 - \alpha) p_L$ leads to an unbiased strategy for all α (this is the case when both p_B and p_L generate unbiased samples), the expected value of $L_{1, \alpha}$ is independent of α since, by definition of unbiased techniques, $E[L_{1, \alpha}] = L$. Minimizing the

sampling variance is therefore similar to minimize $E[L_{1,\alpha}^2]$ since $V[L_{1,\alpha}] = E[L_{1,\alpha}^2] - E[L_{1,\alpha}]^2 = E[L_{1,\alpha}^2] - L^2$.

Nonetheless, the analytical form of the minimum for such a function is too complex to obtain. Therefore, we approximate the problem using a second order Taylor expansion around $\alpha = 1/2$:

$$E[L_{1,\alpha}^2] \approx \int \frac{f^2}{\bar{p}} - (2\alpha - 1) \int \frac{f^2 \Delta p}{\bar{p}^2} + (2\alpha - 1)^2 \int \frac{f^2 \Delta p^2}{\bar{p}^3}$$

where $\bar{p} = (p_B + p_L)/2$ and $\Delta p = (p_B - p_L)/2$. We demonstrate (cf. the derivation in the supplemental material) that this quadratic objective function is minimized when

$$\alpha = \frac{1}{4} \left(2 + \int \frac{f^2 \Delta p}{\bar{p}^2} / \int \frac{f^2 \Delta p^2}{\bar{p}^3} \right) \quad (4)$$

We do a Taylor expansion around $\alpha = 1/2$ since it ensures that the best approximation is centered in the definition domain of α , which is $[0, 1]$. As discussed in Section 2.3, when the best sampling strategy is to use a unique strategy (i.e., $\alpha = 0$ or $\alpha = 1$), our approximation becomes less valid. However, as demonstrated in the Section 3, our estimation of α still reduces the variance compared to the default balance heuristic.

2.2 Estimation of α

Despite the fact that we have an analytical formula (cf. Equation 4) to compute what would be the best α according to a second-order approximation of the sampling variance, two numerical integrations have to be computed. For this purpose, we use a MIS approach by classically selecting the same number of samples from each strategy, and combining them with balance heuristic. Then, the two integrals are estimated using M samples:

$$\begin{aligned} \int \frac{f^2 \Delta p}{\bar{p}^2} &\approx \frac{1}{M} \sum_{m=1}^M \left(\frac{f(\omega_m)}{\bar{p}(\omega_m)} \right)^2 \Delta p(\omega_m) \frac{1}{\bar{p}(\omega_m)} \\ \int \frac{f^2 \Delta p^2}{\bar{p}^3} &\approx \frac{1}{M} \sum_{m=1}^M \left(\frac{f(\omega_m)}{\bar{p}(\omega_m)} \right)^2 \frac{\Delta p^2(\omega_m)}{\bar{p}(\omega_m)} \frac{1}{\bar{p}(\omega_m)} \end{aligned} \quad (5)$$

Therefore, like Bidirectional Importance Sampling techniques and oppositely to [PBPP11], we compute and make use of the evaluated shading function on the generated samples: $f(\omega_m)$. However, instead of rejecting the samples, we use them to estimate the radiance as well.

2.3 Theoretical Accuracy and Improvement

Since our approach is based on an approximation, it has limited accuracy. We estimate its accuracy for two cases: perfect mirror BRDF where α has to be 1 and Dirac light source where α has to be 0. For highly specular cases, BRDFs are zero in almost every direction except in a small direction domain. In this domain, $p_B(\omega)$ is very large, therefore, Δp and \bar{p} are well approximated by $p_B(\omega)$. One can demonstrate (cf. the supplemental material) that $\alpha = 3/4$ instead of $\alpha = 1$. Similarly, with extremely directional light sources, α becomes equal to $1/4$ for our approximation

Method	Scene					
	1	2	8	9	10	11
BRDF	67.3	4597	2.070	0.0198	139.75	0.378
Light	0.47	2.56	0.036	118.20	0.776	0.127
Balanced heur.	0.90	4.76	0.043	0.250	0.402	0.055
Power heur.	0.99	4.87	0.046	0.255	0.440	0.059
Maximum heur.	1.24	5.11	0.061	0.256	0.618	0.075
Preprocessed α	0.42	2.95	0.033	0.095	0.301	0.045

Table 1: Comparisons, for each scene, of different balancing methods in terms of average sampling variance. The average variances are computed over all the pixels using the same number of samples $N = 256$ for each light source.

whereas $\alpha = 0$ is expected. Note that $1/4$ and $3/4$ are not bounds of our estimation, but theoretical values for two selected cases. This demonstrates that, in these two cases, we cannot find the optimal value for α . However, our estimated α still clearly indicates the dominant strategy.

We demonstrate (cf. supplemental material) that our variance reduction is proportional to $|2\alpha - 1|$:

$$\left| V[L_{1,\alpha}] - V[L_{1,1/2}] \right| \leq |2\alpha - 1| \left(V[L_{1,1/2}] + L^2 \right). \quad (6)$$

The closer the estimated α is to $1/2$, the less efficient becomes our approach.

3 Numerical Validation

We have implemented our α estimation into PBRT [PH04] and tested our technique on many different scenes (cf. Figures 1 and 2 and to the supplemental material for the full set). All the results are computed using 4 camera rays per pixel. For each camera ray, we use the default PBRT approach and sample uniformly all light sources and accumulate the contribution of each light source to estimate the final radiance. In the following, we will indicate the number of samples per light source.

3.1 Variance Reduction

To demonstrate the limited accuracy of our approach, we compare (cf. Table 1) the use of a reliable estimation of α , which is computed using $M = 4096$ samples for each light source, against different sampling and combination strategies: only light-based importance sampling, only BRDF-based one, and the half-half strategies that use the same number of samples between the two sampling techniques, combined with balance, power, or maximum heuristics. We do not compare with the α -max heuristic [GKPS12] since it requires additional knowledge about which strategy achieves more often a better sampling than the others.

As shown in Table 1, in most cases our technique reduces the variance compared to previous methods. As discussed in Section 2.3, since our second-order approximation is done around $\alpha = 1/2$, the approximation is getting worse for $\alpha = 0$ or $\alpha = 1$. Therefore, we cannot guarantee to detect correctly if using only light-based or BRDF-based strategy is the best solution. For example in scene 2 only light-based sampling performs better in average, and in scene 9 only BRDF-based sampling is better. However, our approach is




Figure 2: Subset of test scenes, ranging from diffuse (10) to mirror (8) materials including complex visibility (9-10) with different types of illumination. These images are computed using the estimator of Eq. (7) with $N = 256$ and $M = 128$. Since only one sample is required for point and directional light sources, their contribution are added separately and are removed from the variance computation.




Figure 3: Plot of the variance $V[L_{N=256,\alpha}]$ for increasing number of samples M used to estimate α . The plot is normalized by dividing the variance value by the one obtained when estimating α with $M = 4096$. This illustrates the convergence of the α estimation.

always better than default balance/power/maximum heuristics: this demonstrates that our second-order approximation is sufficient to improve (sometimes slightly) these strategies. The improvements may be small in average since, for pixels where the best α is $1/2$, our approach does not improve anything (cf. Equation 6). Moreover, selecting either light-based or BRDF-based requires an a priori knowledge, which is unavailable in most practical cases, but our α estimation provide an indication of what may be the dominant strategy.

3.2 Convergence Speed and Clamping Interval for α

For practical reasons, the number of samples to estimate α should be as low as possible. In Figure 3 we investigate how many samples are required to bring a satisfying variance reduction. Depending on the scene, the number of samples per light source can be low (upper curve) but conservatively, at least 128 samples per light source are required for a good estimation of α that will reduce the variance of the $L_{N,\alpha}$ estimator. We will use this number for the following tests.

We already know from Section 2.3 that our estimation is only an approximation of the best α . In order to investigate its limited accuracy in a more general context, we have tested (cf. supplemental material for details) the effect of clamping the estimated α . The main result of this experiment is that, for some scenes, the variance increases if the value of α is

outside the interval $[0.025, 0.975]$ (i.e., with a size of 0.95). Furthermore, this experience confirms two facts. First, the second-order approximation is not sufficient to capture some variance oscillations in the interval $[0, 1]$ since clamping improves the results. Second, it also confirms that our estimation still provides a good hint of what is the dominant strategy. Finally, it justifies the clamping of α to $[0.025, 0.975]$ for all the following results.

4 Practical Implementations and Results

4.1 Integrated α Estimation

In Section 3.1, we have demonstrated the accuracy of our α estimation with M samples. Moreover, instead of losing these M samples we would like to reuse them to compute the radiance $L_{N,\alpha}$. One solution would be to store and use them with the estimated α to compute the shading. However, this would require an extra-storage of size M per pixel. Therefore, we prefer to use a less optimal but simpler solution that reduces the memory footprint.

We first estimate α using a low number of M samples as in Equation 5. At the same time, we estimate a first radiance value $L_{M/2,M/2}$ as in Equation 2 with the default balance heuristic. Once α is known, we estimate the second radiance value $L_{\alpha(N-M),(1-\alpha)(N-M)}$ and combine these two estimators:

$$L_N = \frac{M}{N} L_{M/2,M/2} + \frac{N-M}{N} L_{\alpha(N-M),(1-\alpha)(N-M)} \quad (7)$$

With this formula, when $\alpha = 1/2$, our estimator is exactly the same as the default balance heuristic one.

Variance Reduction The maximum expected variance reduction is the one obtained without any combination as in Section 3.1 or when N is very large compared to M (cf. the supplemental material). In most cases, the expected reduction will be lower than this maximum, but better than the default balance heuristic. Even when using $M = N/2$ our approach still preserves the variance reduction (cf. Table 2-top). To visualize the variance reduction for each pixel, we compute the difference of variance between the default balance heuristic and our approach (cf. Figure 4). As shown

Method	Scene						
	1	2	8	9	10	11	
for Light and BRDF, cf. Table 1							
Var.	Est. 7	0.66	3.88	0.037	0.1700	0.359	0.051
Time (s)	BRDF	158.6	138.1	36.5	49.2	354	91.0
	Light	299.7	102.8	34.2	27.1	178	53.0
	Est. 7	250.8	128.3	34.5	40.9	281	72.0
Efficiency	BRDF	0.0001	0.0000	0.0132	1.0265	0.0000	0.0291
	Light	0.0071	0.0038	0.8122	0.0003	0.0072	0.1486
	Est. 7	0.0060	0.0020	0.7834	0.1438	0.0099	0.2723

Table 2: Average variance, rendering time and efficiency for different sampling strategies (for the full table, read the supplemental material). We use $N = 256$ samples per light source for each method. Est. 7 refers to the estimator introduced in Equation 7 where we use half of the total samples to evaluate α ($M = 128$) and the remaining half to evaluate the radiance with the estimated α . Efficiency [Vea98] is the inverse of the product between variance and rendering time.

by the preponderance of the red color, our method outperforms the default balance heuristic for most pixels. When our technique is worse (green pixels), the value of α oscillates around $1/2$ (cf. the inset image in Figure 1-left): this illustrates where slow convergence for α -estimation takes place.

Efficiency and Computation Time We have also investigated the overhead of the required α estimation on an Intel i7-3930K@3.2 GHz PC. This overhead takes place only during the first step and is consistently about 1% of the rendering time. This comes from the fact that most of the rendering cost is due to shading computation and PDF evaluation: the estimation of α , based on Equation 5, simply reuses these values. That is also why the fastest strategy is either light-based or BRDF-based sampling (cf. Table 2).

We have computed the efficiency [Vea98] of the sampling strategies: as shown in supplemental material, for all tested scenes our strategy outperforms the previous ones (balance, power and max heuristics). Moreover as shown in Table 2, for some scenes (e.g., 10 and 11), our method is even better than the two extreme strategies: BRDF-based or light-based sampling. Figure 6 shows that for the same rendering time our approach is slightly better for Lab error (1.51 vs 1.57) and estimator variance $V[L_{N,\alpha}]$ ($1.99 \cdot 10^{-4}$ vs $2.22 \cdot 10^{-4}$).

4.2 GPU Implementation

As shown in the companion video, the low overhead of our approach is suitable for a GPU implementation. It could help exploring new balancing strategies for dynamic scenes. For this purpose, we have implemented the two-step approach, described in the previous section, into a full dedicated GPU solution for dynamic environment maps [LPG13] lighting. The solution is implemented on a NVIDIA 580 GTX with 1.5 GB of memory. For one environment map and without any visibility computation, we achieve a frame rate of 67 fps with 256 samples per pixels at a 1024×768 resolution. Finally, Figure 5 and the companion video illustrate that our approach is perfectly suited for dynamic scenes where it is difficult to know a priori the scene characteristics.




Figure 4: Per-pixel comparison for variance reduction. This figure shows the difference of variance between the default balance heuristic and our method. Red corresponds to pixels where our method is better whereas Green corresponds to those where the default balance heuristic is better. Black means that variances are equal.




Figure 5: Dynamic sample distribution for dynamic environment maps (same color legend as figure 1): Red channel represents α (samples from light sources), whereas Green channel represents $1 - \alpha$ (samples from BRDF).

5 Final Discussion and Future Work

Accuracy of α Estimation As pointed out in Sections 2.3 and 3.2, the main limitation of our approach is due to the second-order approximation of the variance around $\alpha = 1/2$. We have investigated third-order approximation but there is generally no real solution to the objective equation. In the case of forth-order approximation, the theoretical upper bound is only slightly improved to $\alpha \approx 0.8$ at the price of a larger computational cost. Finally, we have shown that the use of a non-converged estimation of α may lead to some variance improvements.

Despite these limitations and since our technique exhibits some consistent improvements for variance reduction, we believe that it might be a good framework for future studies on balancing criteria. We think that our second-order approximation provides a good trade-off between accuracy and computational cost. It may also be used to estimate α progressively using a non-linear optimization method.

More Sampling Strategies We have experimented our approach for balancing only between two strategies that do not require a precomputation step. The same approach might be used to combine them with other strategies such as the ones based on Photon Mapping as in [PBPP11] or visibility-based ones [GH06]. For more than two strategies, the minimization of the second-order approximation will result in solving a linear system to find the strategy weights.

Since our approach reduces the sampling variance, it can also benefit to Bidirectional Importance Sampling strategies.




Figure 6: Equal time (72ms) comparisons for scene 11. Our solution is slightly smoother and also closer (1.51 vs 1.57 mean Lab error) to the reference solution than the default balance heuristic (cf. supplemental material for more comparisons).

For example, Burke et al. [BGH05] use either BRDF-based or light-based strategies as initialization step. Improved sampling would certainly improve it. Similarly, our approach can serve as a support for Control Variate techniques (e.g., [CAM08a]) since they can be used on top of state-of-the-art schemes for importance sampling.

6 Conclusion

In this paper, we have introduced a second-order approximation of variance for Multiple Importance Sampling. This approximation leads to an automatic distribution of samples between different sampling strategies. We have demonstrated that, for all our test scenes, our balancing technique reduces (sometimes slightly) the variance compared to previous MIS approaches and outperforms them as well in terms of efficiency. Finally, we have also shown that it fits into existing MIS approaches and that it can be implemented on GPU. We believe that our approach will help further investigations on how to develop improved balancing strategies.

Acknowledgments

We would like to thank Laurent Belcour for his suggestions as well as Petrik Clarberg for sharing scene 2. Heqi Lu's PhD scholarship is funded by the Région Aquitaine. This research has been supported by the ALTA project (ANR-11-BS02-006).

References

- [AP07] ASHIKHMIN M., PREMOZE S.: *Distribution-based BRDFs*. Tech. rep., Univ. of Utah, 2007. <http://www.cs.utah.edu/~premoze/dbrdf>. 2
- [BGH05] BURKE D., GHOSH A., HEIDRICH W.: Bidirectional Importance Sampling for Direct Illumination. In *Eurographics Symposium on Rendering* (2005), pp. 147–156. 2, 6
- [CAM08a] CLARBERG P., AKENINE-MÖLLER T.: Exploiting Visibility Correlation in Direct Illumination. *Computer Graphics Forum* 27, 4 (2008), 1125–1136. 6
- [CAM08b] CLARBERG P., AKENINE-MÖLLER T.: Practical Product Importance Sampling for Direct Illumination. *Computer Graphics Forum* 27, 2 (2008). 2
- [CETC06] CLINE D., EGBERT P. K., TALBOT J., CARDON D. L.: Two Stage Importance Sampling for Direct Lighting. In *Proc. Eurographics Symposium on Rendering* (2006), pp. 103–113. 2
- [EKÖ12] ERGUN S., KURT M., ÖZTÜRK A.: Real-time Kd-tree Based Importance Sampling of Environment Maps. In *Proc. Spring Conference on Computer Graphics* (2012), Comenius University Press, pp. 84–91. 2
- [GH06] GHOSH A., HEIDRICH W.: Correlated visibility sampling for direct illumination. *Vis. Comput.* 22, 9 (2006), 693–701. 2, 5
- [GKPS12] GEORGIEV I., KŘIVÁNEK J., POPOV S., SLUSALLEK P.: Importance Caching for Complex Illumination. *Computer Graphics Forum* 31, 2pt3 (2012), 701–710. 2, 3
- [Hes95] HESTERBERG T.: Weighted Average Importance Sampling and Defensive Mixture Distributions. *Technometrics* 37, 2 (1995), 185–194. 2
- [HSK*05] HAVRAN V., SMYK M., KRAWCZYK G., MYSZKOWSKI K., SEIDEL H.-P.: Interactive System for Dynamic Scene Lighting using Captured Video Environment Maps. In *Eurographics Symposium on Rendering* (2005), pp. 31–42. 2
- [Kaj86] KAJIYA J. T.: The rendering equation. In *Proc. SIGGRAPH '86* (1986), ACM, pp. 143–150. 1
- [LPG13] LU H., PACANOWSKI R., GRANIER X.: Real-Time Importance Sampling of Dynamic Environment Maps. In *Eurographics Annual Conference (Short Papers)* (2013), Eurographics. 5
- [ODJ04] OSTROMOUKHOV V., DONOHUE C., JODOIN P.-M.: Fast hierarchical importance sampling with blue noise properties. In *Proc. SIGGRAPH '04* (2004), ACM, pp. 488–495. 2
- [OZ00] OWEN A., ZHOU Y.: Safe and Effective Importance Sampling. *J. American Statistical Association* 95, 449 (2000), 135–143. 2
- [PBPP11] PAJOT A., BARTHE L., PAULIN M., POULIN P.: Representativity for robust and adaptive multiple importance sampling. *IEEE Trans. Visualization and Computer Graphics* 17, 8 (2011), 1108–1121. 2, 3, 5
- [PCS*12] PACANOWSKI R., CELIS O. S., SCHLICK C., GRANIER X., POULIN P., CUYT A.: Rational BRDF. *IEEE Trans. Visualization and Computer Graphics* 18, 11 (2012), 1824–1835. 2
- [PH04] PHARR M., HUMPHREYS G.: *Physically Based Rendering: From Theory to Implementation*. Morgan Kaufmann, 2004. 3
- [RCL*08] ROUSSELLE F., CLARBERG P., LEBLANC L., OSTROMOUKHOV V., POULIN P.: Efficient Product Sampling using Hierarchical Thresholding. *The Visual Computer* 24, 7-9 (2008), 465–474. 2
- [TCE05] TALBOT J. F., CLINE D., EGBERT P.: Importance resampling for global illumination. In *Proc. Eurographics Symposium on Rendering* (2005), pp. 139–146. 2
- [Vea98] VEACH E.: *Robust monte carlo methods for light transport simulation*. PhD thesis, 1998. Stanford University. 1, 2, 5
- [WA09] WANG R., AKERLUND O.: Bidirectional Importance Sampling for Unstructured Direct Illumination. *Computer Graphics Forum* 28, 2 (2009), 269–278. 2
- [WFA*05] WALTER B., FERNANDEZ S., ARBREE A., BALA K., DONIKIAN M., GREENBERG D. P.: Lightcuts: a scalable approach to illumination. *ACM Trans. Graph.* 24, 3 (2005), 1098–1107. 2