

On Learning the Best Local Balancing Strategy

D. Murray¹ and S. Benzait¹ and R. Pacanowski² and X. Granier^{2,3}

1 - INRIA 2 - LP2N (IOGS, CNRS, U. Bordeaux) 3 - ARCHEOVISION (CNRS, U. Bordeaux, U. Bordeaux Montaigne)

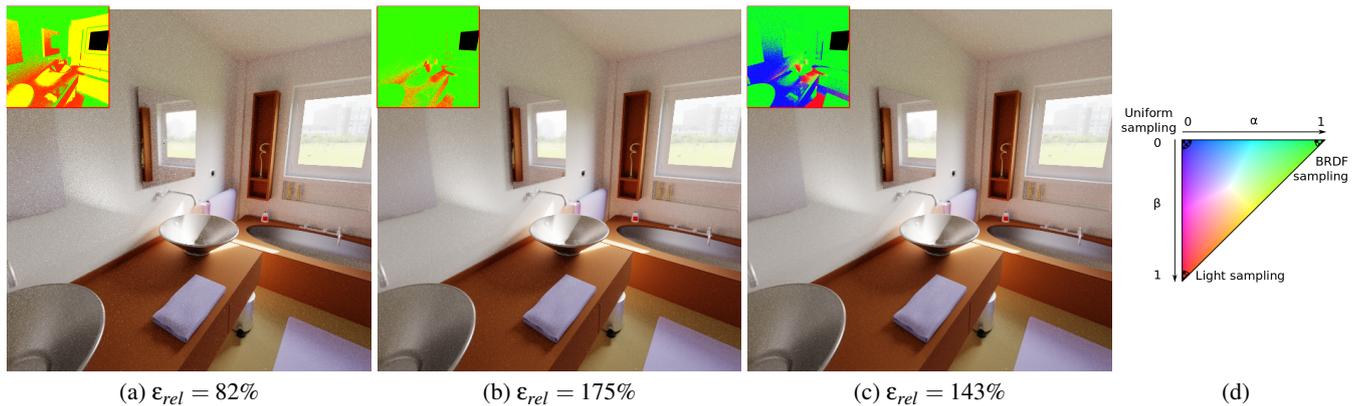


Figure 1: Optimizing sample allocation comparisons for the first bounce between (a) two sampling strategies (light and BRDF) using [SHSK19], (b) two sampling strategies using *our solution*, (c) three strategies (by adding a uniform sampling strategy). (d) Color map used to display the sample allocation (top left corner of each image) for each strategy. The balance maps are computed using 512 learning samples and 8 iterations. The efficiency is given relative to the reference computed with the balanced heuristic.

Abstract

Fast computation of light propagation using Monte Carlo techniques requires finding the best samples from the space of light paths. For the last 30 years, numerous strategies have been developed to address this problem but choosing the best one is really scene-dependent. Multiple Importance Sampling (MIS) emerges as a potential generic solution by combining different weighted strategies, to take advantage of the best ones. Most recent work have focused on defining the best weighting scheme. Among them, two paper have shown that it is possible, in the context of direct illumination, to estimate the best way to balance the number of samples between two strategies, on a per-pixel basis.

In this paper, we extend this previous approach to Global Illumination and to three strategies.

1. Motivation

To generate a synthetic image, the light propagation is computed by solving the rendering equation [Kaj86]. This equation computes the radiance for a position \mathbf{x} , as viewed from direction \mathbf{o} , such as:

$$L(\mathbf{x} \rightarrow \mathbf{o}) = \int_{\Omega_n} \rho(\mathbf{x}, \mathbf{o}, \boldsymbol{\omega}) (\boldsymbol{\omega} \cdot \mathbf{n}) L(\mathbf{x} \leftarrow \boldsymbol{\omega}) d\boldsymbol{\omega}, \quad (1)$$

where \mathbf{n} is the normal of the surface at \mathbf{x} , Ω_n is the unit hemisphere, $\rho(\mathbf{x}, \mathbf{o}, \boldsymbol{\omega})$ is the reflectance function (BRDF), and $L(\boldsymbol{\omega} \rightarrow \mathbf{x})$ is the incoming radiance at position \mathbf{x} from direction $\boldsymbol{\omega}$. To compute this integral efficiently, most techniques rely on Monte Carlo (MC) methods where $L(\boldsymbol{\omega} \rightarrow \mathbf{x})$ is approximated by a MC-estimator F_N

defined by

$$F_N = \frac{1}{N} \sum_i^N \frac{\rho(\mathbf{x}, \mathbf{o}, \boldsymbol{\omega}_i) (\boldsymbol{\omega}_i \cdot \mathbf{n}) L(\mathbf{x} \leftarrow \boldsymbol{\omega}_i)}{p(\boldsymbol{\omega}_i)}, \quad (2)$$

where $p(\boldsymbol{\omega})$ represents the probability density function (PDF) and N the number of samples. Without prior knowledge of the light transport in the scene, finding the optimal PDF is not possible. Therefore, common sampling techniques are mostly based on either BRDF (or cosine times BRDF) or light sampling strategies or even a combination of multiple strategies by using Multiple Importance Sampling (MIS [Vea97]). Given S sampling strategies, an MIS estimator is defined by

$$F_N = \sum_{s=1}^S \frac{1}{N_s} \sum_{i=0}^{N_s} w_s(\boldsymbol{\omega}_{i,s}) \frac{f(\mathbf{x}, \mathbf{o}, \boldsymbol{\omega}_{i,s})}{p_s(\boldsymbol{\omega}_{i,s})}, \quad (3)$$

where $f(\mathbf{x}, \boldsymbol{o}, \boldsymbol{\omega}) = \rho(\mathbf{x}, \boldsymbol{o}, \boldsymbol{\omega})(\boldsymbol{\omega} \cdot \mathbf{n})L(\mathbf{x} \leftarrow \boldsymbol{\omega})$, p_s represents the PDF associated with the s -th sampling strategy, $N = \sum_s N_s$, and w_s are weights computed using different heuristic (e.g., max, power, balanced). By exploiting different sampling strategies, the MIS estimator relies on the hope that at least one of them performs well, for a limited overhead as it should not perform worse than any of them.

Many sampling strategies (e.g., [KVG⁺19,GGSK19,HEV⁺16]) have been proposed, along with solutions to determine which one is good. However, most of these solutions focus on finding efficient strategies using equal sample distribution and only a few [LPG13,SHSK19] focus on optimizing sample allocation. Lu et al. [LPG13] compute a balancing factor α per pixel between two sampling strategies. They use a second-order approximation of the variance to obtain the theoretical optimal α . However, their solution is limited to direct lighting and only provide guidelines for the best strategy as it cannot identify the cases where a single strategy is optimal. Sbert et al. [SHSK19] demonstrate that using iterative optimization of α converges toward the optimal strategy, still for direct lighting only.

In this paper, we focus on solutions that seek the optimal *per-pixel* balancing allocation among different importance sampling strategies. **First**, we analyse, in terms of efficiency, the recent work of Sbert et al. [SHSK19] when used in the context of Global Illumination for scenes either lit by direct or indirect lighting. **Second**, we propose a more efficient way to address these cases by modifying the approach. **Third**, we introduce an extension to account for more than two strategies and show the benefit of adding them. All presented results are generated using the Malia Rendering Framework [DMP⁺], a GPU Path-Tracing engine, with an Nvidia[®] Titan V graphics cards. All images are computed with 32768 samples per pixel (including the learning samples).

2. Principles of the Approach and the Study

As pointed by Georgiev et al. [GKPS12], the MIS estimator in Eq. 3 can be rewritten in terms of Defensive importance sampling. For the case of two sampling strategies, the MIS estimator becomes

$$F_N = \frac{1}{N} \sum_{i=1}^N \frac{f(\mathbf{x}, \boldsymbol{o}, \boldsymbol{\omega}_i)}{\alpha p_{\text{BRDF}}(\boldsymbol{\omega}_i) + (1 - \alpha) p_{\text{light}}(\boldsymbol{\omega}_i)}, \quad (4)$$

where $\alpha \in [0, 1]$ is the balancing factor between using BRDF sampling and light sampling (with respective PDFs p_{BRDF} and p_{light}). The PDF $p_\alpha = \alpha p_{\text{BRDF}} + (1 - \alpha) p_{\text{light}}$ of this estimator is indeed a weighted combination of the PDF of each sampling strategy.

The best α can be iteratively estimated [SHSK19] using a second-order approximation [LPG13] by

$$\alpha_{n+1} = \alpha_n - \frac{V'_\alpha[F](\alpha_n)}{V''_{\alpha\alpha}[F](\alpha_n)}, \quad (5)$$

where $V[F]$ is the variance of the estimator F and where V'_α and $V''_{\alpha\alpha}$ are respectively the first and second derivatives according to parameter α . We refer to the supplemental material for the exact expression of $V'_\alpha[F](\alpha_n)$ and $V''_{\alpha\alpha}[F](\alpha_n)$.

With such an approach, only variance is optimized. In order to check the ability of the proposed method to provide an accurate solution quickly, we need to check its efficiency according to Veach's

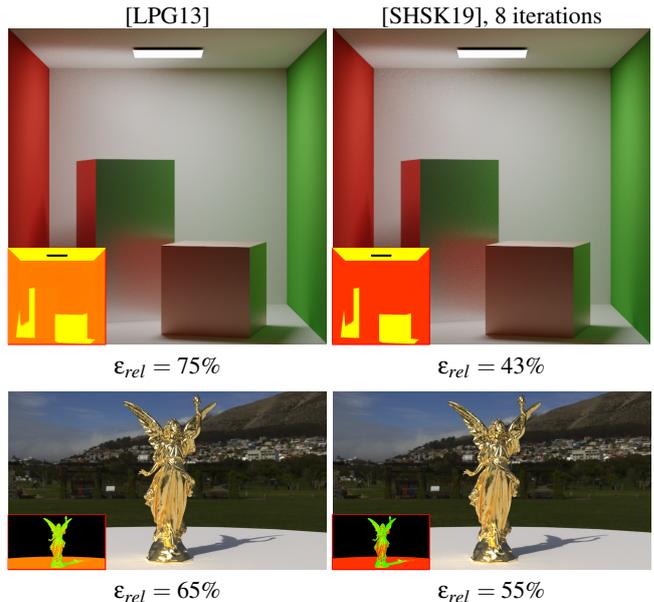


Figure 2: Global Illumination using precomputed α -maps, obtained with [LPG13] and [SHSK19], using direct lighting only and 256 learning samples per iteration. **Top Part:** Closed scene with direct and indirect illumination. The relative efficiency drops as the variance increases significantly due to indirect paths being under-sampled. **Bottom Part:** Open scene dominated by direct illumination. The efficiency drop is less important as most light paths are indeed direct ones, but still present because inter-reflection paths are poorly sampled.

definition ([Vea97], Equation 2.19). In this paper, all presented results are relative to the efficiency of the balanced heuristic with equal distribution.

3. Accounting for Indirect Lighting

Previous approaches have been used in direct lighting cases only (i.e., estimating the balancing between choosing a direction according to light sources or the BRDF). Most use-cases for Monte-Carlo methods are directed at performing Global Illumination, for which the aforementioned sample allocation has yet to be extended.

A first and simple extension to Global Illumination is to reuse the α value estimated from the direct illumination for the first ray-bounce. For the other bounces, a direction is sampled according to the BRDF and direct lighting is systematically estimated as in [Kaj86]. However, this is inefficient for most cases as illustrated in Figure 2. Indeed, this approach decreases the efficiency:

- slightly for open scenes, lit by an environment map, as most light paths consist in direct lighting and only a small amount consists in indirect paths.
- significantly for closed scene, the efficiency drops due to indirect light paths being widely under-sampled as light sampling as been identified as the best strategy.

With this approach, the method of [LPG13] to compute alpha

performs better than [SHSK19] α values are more restricted ($[0.25, 0.75]$ vs. $[0.1, 0.9]$) for "perfect" cases, thus preventing to use only a single strategy that may not be appropriate in the context of Global Illumination.

Another approach is to simply account for the radiance from *indirect* lighting contribution $L(p \leftarrow \omega)$ when computing α instead of the radiance issued from direct lighting only. This means that we still use α to guide the samples allocation only for the first bounce of the light path, but this value will now include information regarding the full light path. Note that we cannot simply compute α per bounce per pixel as we would be mixing spatially incompatible information. By doing so, we can extend the previous methods to increase their efficiency in the context of Global Illumination (cf. Figure 3). As in the previous solution, this converges toward a similar optimum independently of the initial choice of balancing (as illustrated in the supplemental video).

4. Optimizing between more strategies

We now extend this method to include more than just light or BRDF sampling strategies. In this paper, we study only the case with three strategies at the same time, the general formulation for S strategies can be found in the supplemental material. As a third strategy, we choose uniform sampling on the hemisphere: it ensures that the third strategy is not correlated to the other two while ensuring the exploration of the full space of light paths.

With three strategies (light, BRDF and uniform sampling), the PDF relative to defensive sampling becomes

$$P_{\alpha,\beta} = \alpha P_{\text{BRDF}} + \beta P_{\text{light}} + (1 - \alpha - \beta) P_{\text{uniform}}, \quad (6)$$

with $\alpha, \beta \in [0, 1]; \alpha + \beta \leq 1$. In this case, minimizing the variance means that we are looking for α, β such that

$$\nabla V[F](\alpha, \beta) = \begin{bmatrix} V'_{\alpha}[F](\alpha, \beta) \\ V'_{\beta}[F](\alpha, \beta) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \quad (7)$$

$$\text{determinant}(\mathcal{H}(V[F](\alpha, \beta))) > 0 \quad (8)$$

$$\text{trace}(\mathcal{H}(V[F](\alpha, \beta))) > 0, \quad (9)$$

where $\mathcal{H}(V[F](\alpha, \beta))$ is the 2×2 Hessian matrix of the variance. The full formulation of $\mathcal{H}(V[F](\alpha, \beta))$ is provided in the supplemental material in which we demonstrate that assumptions of Equations 8 and 9 are valid in our case where the strategies are not correlated. Thus, we only need to solve Equation 7, using the Newton-Raphson method, as follows:

$$\begin{bmatrix} \alpha_{n+1} \\ \beta_{n+1} \end{bmatrix} = \begin{bmatrix} \alpha_n \\ \beta_n \end{bmatrix} - \mathcal{H}^{-1}(V[F](\alpha_n, \beta_n)) \nabla V[F](\alpha_n, \beta_n). \quad (10)$$

As in [LPG13, SHSK19], all these derivatives are estimated by Monte Carlo integration. The learning samples are also used for rendering the final image. Moreover, as we did for two strategies, we must clamp $\alpha + \beta$ in the range $[0.1; 0.9]$ to ensure that we keep exploring all light paths.

Our third strategy (uniform sampling) is not expected to perform well and we expect that in most cases, $\alpha + \beta \geq 0.9$, yielding to an allocation close to the one we obtained with only two strategies. However, our main goal is to highlight that our technique is not

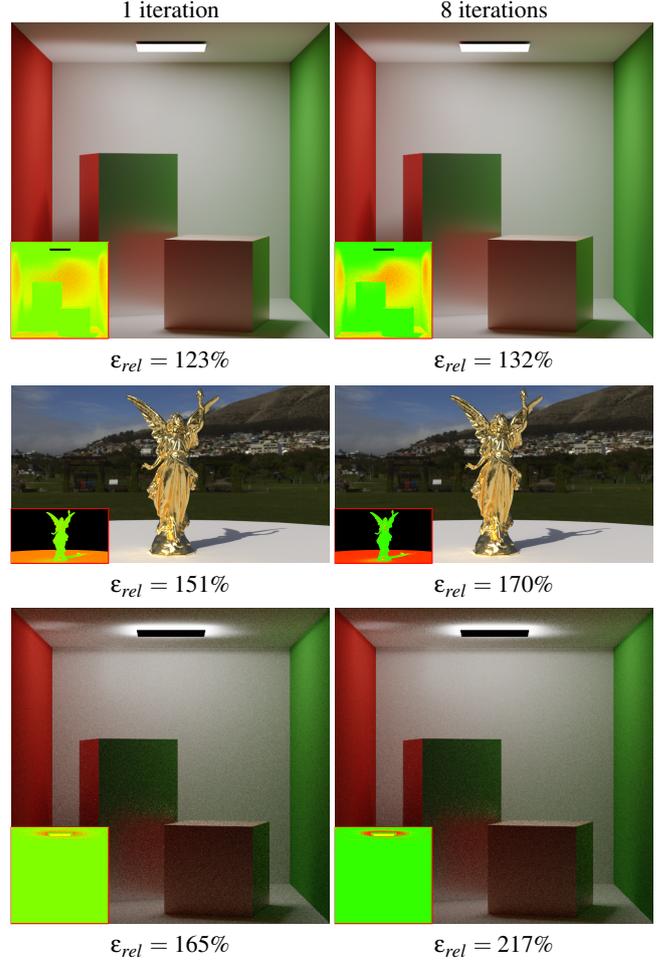


Figure 3: Comparison between one iteration and 8, using Global Illumination to compute α and 256 learning samples per iteration. **Top Part:** In areas where indirect contributions are greater than direct ones, BRDF sampling is identified as optimal whereas light sampling is optimal where direct contributions are greater. This results in a decreased variance at nearly equal time. **Middle Part:** The behavior is close to the one in Figure 2 except on areas where there were many inter-reflections, where BRDF sampling is optimal, leading to a reduced variance and increased efficiency. **Bottom Part:** BRDF sampling is almost a global optimal strategy, yielding a significant gain of efficiency.

limited to two strategies and can identify the local optimal strategy between more than two, as illustrated on Figure 4. We can see that, even with a simple strategy, our approach still increases the efficiency since it gives higher weights to near-tangent contributing paths.

5. Limitations and Future Work

Like [SHSK19], we demonstrated that our approach yields results with a relatively small number of learning samples. The balancing percentages (e.g., α and β) must still be clamped to ensure that we have a sufficient number of samples to explore the other strategies. This limits the efficiency gain as some of these paths may not

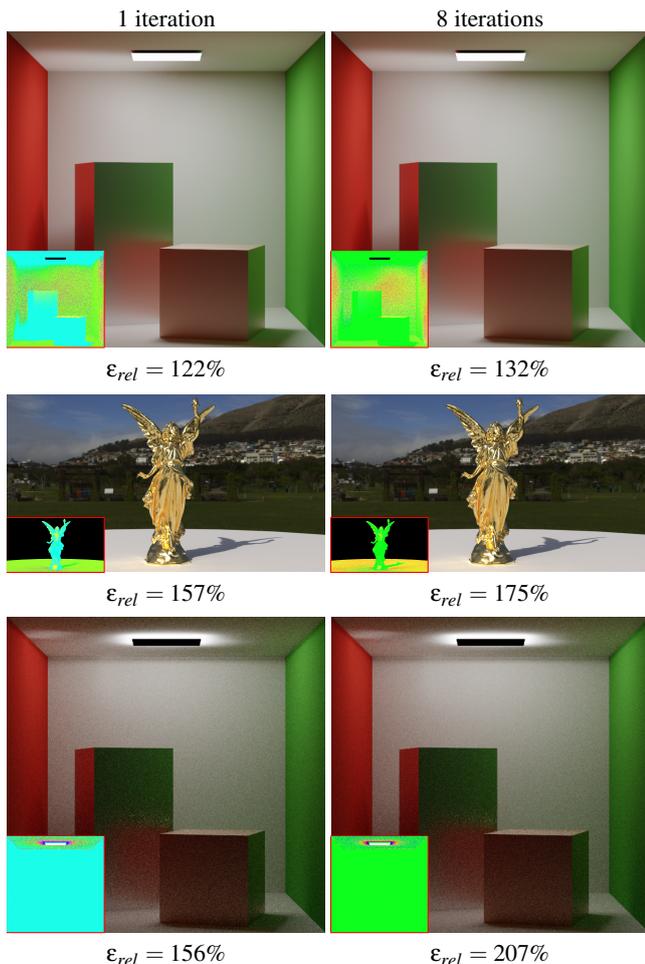


Figure 4: Comparisons between one and 8 iterations, using 256 learning samples per iteration. Relative efficiency is wrt. to the balance heuristic for two and three strategies. In all cases, there is a gain compared to the reference obtained with either two or three strategies.

contribute to the final image. During the learning step with a large number of samples, clamping is no longer necessary.

As a first experiment, we have limited the balancing to the first bounce. The resulting gains lead us to believe that it can be generalized to any bounce. For this purpose, a spatial structure has to be developed and trained (e.g., with a caching approach [GKPS12]). With such an extension, we may also learn how much we can limit the number of shadow rays for direct lighting at any bounce. Most approaches that aim at resolving complex light paths efficiently could benefit from our solution, in order to better identify these complex paths (e.g., [BJ19]).

Furthermore, we use a simple uniform sampling as a third strategy to test our extension. We have shown in our preliminary tests some gain in efficiency. This demonstrates that considering more tailored strategies may result in more significant gains. In particular, forcing the exploration the space excluded by the classical sampling strategies may be an interesting study-case, as well as using

cache-based PDF such as in [HEV⁺16] as a complementary strategy. Our generalized approach to more than two strategies, combined with the fact that the convergence is independent from the initial balancing strategy, opens a new field of studies to learn the best strategy in the context of MIS.

6. Conclusion

In this paper, we proposed a first solution to learn the best way to balance *per pixel* the number of samples in an MIS approach, extended for both Global Illumination and for more than two strategies. Based on these propositions, we performed a study on the convergence of the learning process and its impact on the efficiency. In particular, we showed that using this balancing for the first bounce of camera rays already yields a significant gain in term of efficiency, compared to the balance heuristic with equal sample. We also showed that introducing a simple third strategy that it is not correlated to the other ones also improves the efficiency compared to the same heuristic. As discussed in the previous section, we believe that our approach opens the path to: a better understanding of the impact of the sampling strategies, the definition of new complementary strategies, and a generic learning process for the best balancing.

7. Acknowledgements

This project is supported by the CaLiTrOp ANR Project (ANR-16-CE33-0026) and the VIDA ANR Project (ANR-17-CE23-0017). We thank The Stanford 3D Scanning Repository, M. McGuire’s Archive and B. Bitterli’s Resources for the datasets.

References

- [BJ19] B. Bitterli and W. Jarosz. Selectively Metropolis Monte Carlo light transport simulation. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)*, 38(6), November 2019. 4
- [DMP⁺] A. Dufay, D. Murray, R. Pacanowski, et al. Malia rendering framework. <https://pacanows.gitlabpages.inria.fr/MRF/main.md.html>. 2
- [GGSK19] P. Grittmann, I. Georgiev, P. Slusallek, and J. Křivánek. Variance-aware multiple importance sampling. *ACM Trans. Graph. (SIGGRAPH Asia 2019)*, 38(6), 2019. 2
- [GKPS12] I. Georgiev, J. Křivánek, S. Popov, and P. Slusallek. Importance caching for complex illumination. In *Computer Graphics Forum*, volume 31, pages 701–710. Wiley Online Library, 2012. 2, 4
- [HEV⁺16] S. Herholz, O. Elek, J. Vorba, H. Lensch, and J. Křivánek. Product importance sampling for light transport path guiding. In *Computer Graphics Forum*, volume 35, pages 67–77. Wiley Online Library, 2016. 2, 4
- [Kaj86] James T. Kajiya. The rendering equation. *SIGGRAPH Comput. Graph.*, 20(4):143–150, August 1986. 1, 2
- [KVG⁺19] I. Kondapaneni, P. Vévoda, P. Grittmann, T. Skřivan, P. Slusallek, and J. Křivánek. Optimal multiple importance sampling. *ACM Trans. Graph.*, 38(4), 2019. 2
- [LPG13] H. Lu, R. Pacanowski, and X. Granier. Second-order approximation for variance reduction in multiple importance sampling. *Computer Graphics Forum*, 32(7):131–136, October 2013. 2, 3
- [SHSK19] M. Sbert, V. Havran, and L. Szirmay-Kalos. Optimal deterministic mixture sampling. In *Eurographics (Short Papers)*, pages 73–76, 2019. 1, 2, 3
- [Vea97] Eric Veach. *Robust Monte Carlo methods for light transport simulation*, volume 1610. Stanford University PhD thesis, 1997. 1, 2