

Adaptive BRDF-Oriented Multiple Importance Sampling of Many Lights

Yifan Liu¹, Kun Xu^{1*}, and Ling-Qi Yan²

¹ BNRist, Department of Computer Science and Technology, Tsinghua University, Beijing, China

² Department of Computer Science, University of California, Santa Barbara, USA

* Corresponding author, xukun@tsinghua.edu.cn




Figure 1: We present an adaptive BRDF-oriented multiple importance sampling method for many-light rendering. In equal time (55s), our method results in significantly better quality as compared to state of the art method combining light and BRDF sampling [EK18]. The 3D models are courtesy of Sketchfab.com users tkachsb and Comrade1280.

Abstract

Many-light rendering is becoming more common and important as rendering goes into the next level of complexity. However, to calculate the illumination under many lights, state of the art algorithms are still far from efficient, due to the separate consideration of light sampling and BRDF sampling. To deal with the inefficiency of many-light rendering, we present a novel light sampling method named BRDF-oriented light sampling, which selects lights based on importance values estimated using the BRDF's contributions. Our BRDF-oriented light sampling method works naturally with MIS, and allows us to dynamically determine the number of samples allocated for different sampling techniques. With our method, we can achieve a significantly faster convergence to the ground truth results, both perceptually and numerically, as compared to previous many-light rendering algorithms.

CCS Concepts

• Computing methodologies → Ray tracing;

1. Introduction

As computer graphics approaches the next level of details and realism, it is more and more common that we need to accurately model real-world luminaires and use them to render complex scenes. This indicates that we often have millions of lights in one scene, es-

pecially for production scenes that involve small particles or large scale views. The large number of light sources introduces the many-light rendering problem. That is, how to efficiently calculate the direct illumination — the sum of the contributions of all individual lights.

Various methods have been proposed to deal with the many-light rendering problem [VKK18, EK18]. However, existing methods still introduce a considerable amount of noise, especially for glossy scenes with a large amount of lights. The reason is that, these methods share a common problem — the separate consideration of light sampling and BRDF sampling. On one hand, to importance sample a light, it is difficult to find lights with large contributions to shading points without utilizing BRDF information. On the other hand, pure BRDF importance sampling results in rays that constantly miss scattered light sources. Moreover, inefficiency within these two techniques may even happen simultaneously, indicating that a combination of them using multiple importance sampling (MIS) still cannot solve the problem.

Our method is aimed to improve the efficiency of many-light rendering. Specifically, a new light sampling method named *BRDF-oriented light sampling* is proposed. Our key idea is to associate each light or light cluster with an importance value, quickly estimated using the BRDF at the shading point, so that we can importance sample the lights both individually and in a hierarchical fashion (Sec. 4). Besides, we combine our BRDF-oriented light sampling using MIS with existing sampling approaches, such as traditional irradiance-oriented light sampling and classic BRDF importance sampling. Furthermore, we propose an adaptive heuristic that dynamically determines the number of samples allocated for different sampling techniques (Sec. 5). With our method, we can achieve a significantly faster convergence to the ground truth results, both perceptually and numerically, as compared to previous many-light rendering algorithms (Sec. 6).

2. Related Work

Illumination with many lights. Historically, many-light methods were proposed to approximate indirect illumination, by turning the problem into calculating direct illumination with a large number of virtual point lights (VPLs) [Kel97]. With the development of modern rendering, the use of many lights was later extended to directly illuminate complex scenes. In either type of applications, the amount of lights can be large, thus scalable methods [WFA*05, WABG06, BMB15] were proposed to reduce the computational complexity to be sublinear in the number of lights. For a more detailed introduction to these many-light methods, we refer readers to [DKH*14] for an overview.

Different to the goal of approximating illumination in those many-light methods, we aim at selecting lights based on our proposed importance metric. At the same time, similar to these methods, we also focus on direct illumination, since indirect illumination is automatically handled by the recursion of light transport.

Importance sampling of light sources is a long-standing problem that dates back to Monte Carlo path tracing. In [War94, SWZ96], the idea of importance sampling light sources according to their contributions was proposed. Since often applied in many-light rendering, importance sampling light sources was focused more on selecting light clusters rather than individual lights. Wang et al. [WÅ09] sample light clusters by their solid angle coverage on directions generated according to BRDF. Wu et al. [WC13] estimate visibility by shooting shadow rays and take visibility

into account in importance sampling of light clusters. Vévoda et al. [VKK18] further introduce an online Bayesian regression method that adaptively updates the visibility estimation. Instead of sampling light clusters, Estevez et al. [EK18] compute the overall direction illumination by summing up contributions from all light clusters, while the contribution of each light cluster is estimated through importance sampling.

Range evaluation of BRDFs. Our importance metric of light sampling (Sec. 4) is closely related to calculating the BRDF integration within a certain solid angle subtended by different light sources. To do so, researchers have proposed to represent BRDFs using wavelets [CJAMJ05], quad-trees [CAM08], or spherical harmonics [JCJ09] to approximate BRDF integration. However, all these methods can only compute BRDF integrals over the entire sphere or hemisphere, and cannot be applied to integrals within a certain solid angle.

To solve the problem, Xu et al. [XCM*14] proposed a method to integrate a BRDF over a spherical triangle. Linearly Transformed Cosines (LTCs) [HDHN16] were introduced to analytically integrate a BRDF within a polygonal area. Spherical Pivot Transformed Distributions (SPTDs) [DHB17] were proposed to analytically integrate a BRDF within solid angles subtended by spheres. We refer to SPTDs to perform our range evaluation, not only for BRDF integration, but also for full contribution of the rendering equation.

Multiple Importance Sampling (MIS). Veach [Vea97] proposed MIS to combine multiple sampling strategies for estimating the same integration. MIS not only considers a properly balanced weight to each sampling strategy, but also supports different numbers of samples for different strategies. Thus, a considerable amount of work has been devoted to adaptively determine the sample allocation, either based on analysis of the variance of MIS estimators [LPG13, HS14, SHSK16, SH17, SHSKE18], or by dynamically increasing the sampling rate of different strategies with various heuristics [PBPP11].

3. Background and Motivation

In this section, we first briefly introduce the theory of Monte Carlo direct illumination estimation in the context of many lights. Then we analyze the two most commonly used types of approaches, light sampling and BRDF sampling, and their own disadvantages when applied individually and combined together. Based on the analysis, we propose the idea of our BRDF-oriented light sampling method.

3.1. Direct illumination estimation

According to the rendering equation, the direct illumination at a shading point \mathbf{x} could be computed by:

$$L_o(\mathbf{x}, \omega_o) = \int_{\Omega} L(\mathbf{x}, \omega_i) B(\mathbf{x}, \omega_i, \omega_o) \cos \theta_{\mathbf{x}} d\omega_i, \quad (1)$$

where ω_i, ω_o are incident and outgoing directions, respectively. L_o is the outgoing radiance of direct illumination. L refers to the incoming radiance, but only considers energy directly emitted from light sources. B is the BRDF at \mathbf{x} . $\theta_{\mathbf{x}}$ is the angle between the incident direction and surface normal. Ω is the upper hemisphere. Note that self emission is not included.

In the context of many-light rendering, it is common practice that we reinterpret Eqn. 1 as an integration over the area A of all surfaces on all lights:

$$L_o(\mathbf{x}, \omega_o) = \int_A F(\mathbf{y} \rightarrow \mathbf{x} \rightarrow \omega_o) d\mathbf{y}, \quad (2)$$

where \mathbf{y} denotes a light point (i.e. a point on a light source), and the integrand F is computed by:

$$F(\cdot) = L(\mathbf{y} \rightarrow \mathbf{x})B(\mathbf{y} \rightarrow \mathbf{x} \rightarrow \omega_o)V(\mathbf{y} \leftrightarrow \mathbf{x})G(\mathbf{y} \leftrightarrow \mathbf{x}), \quad (3)$$

where L and B are defined earlier. $V(\mathbf{y} \leftrightarrow \mathbf{x})$ is the binary visibility function between \mathbf{y} and \mathbf{x} . The geometry factor $G(\mathbf{y} \leftrightarrow \mathbf{x})$ is computed as $\frac{\cos \theta_x \cos \theta_y}{d^2(\mathbf{y} \leftrightarrow \mathbf{x})}$, where θ_y is the angle between incident direction $\mathbf{y} \rightarrow \mathbf{x}$ and the surface normal at \mathbf{y} , and $d(\mathbf{y} \leftrightarrow \mathbf{x})$ is the Euclidean distance between \mathbf{y} and \mathbf{x} .

The direct illumination in Eqn. 2 could be estimated using Monte Carlo integration, given by:

$$\langle L_o(\mathbf{x}, \omega_o) \rangle = \frac{F(\mathbf{y} \rightarrow \mathbf{x} \rightarrow \omega_o)}{p(\mathbf{y}|\mathbf{x}, \omega_o)}, \quad (4)$$

where $p(\mathbf{y}|\mathbf{x}, \omega_o)$ is the probability density function (pdf) for sampling the light point \mathbf{y} .

In order to reduce the variance of this estimation, the pdf $p(\cdot)$ is expected to be as close to the integrand F as possible. An ideal pdf $p(\cdot)$ will be proportional to $LBVG$, according to the definition of F in Eqn. 3.

3.2. An analysis of light sampling

Light sampling is essentially using a pdf p proportional to L , or more precisely, LG . Recently, Vévoda et al. [VKK18] improve light cluster sampling by taking visibility into account. They partition the scene into disjoint spatial regions and adaptively update the visibility estimation using online Bayesian regression. However, they do not consider visibility in light sampling within a cluster. Thus, it is partially considering LVG . Estevez et al. [EK18] adopt an approach based on stratified sampling. They first perform dynamic light clustering for each shading point, so that the lights in each cluster exhibit a certain extent of similarity. To sample a light within a cluster, they consider LG only.

From the analysis, we find that none of these methods consider BRDF information during light sampling. Thus, when the BRDF is glossy or complicated, the lights selected by these methods may not contribute much when they are not inside the BRDF lobe.

3.3. An analysis of BRDF sampling

Another way of sampling the direct illumination estimation in Eqn. 2 is that, we may primarily sample the BRDF, while consider the lights' irradiance at the same time.

These methods are mostly designed for environment lighting or path guiding. In both cases, illumination on the hemisphere will be represented as a spherical function. For environment maps, the spherical function is usually fixed. For path guiding, the function is used to represent the indirect illumination, usually learned, dynamically updated, and stored in different kinds of basis functions, such




Figure 2: BRDF sampling often selects invalid directions when a scene contains many small lights, thus is not efficient in many-light rendering.

as spherical Gaussians [VKŠ*14, HEV*16] and piecewise constant grids [MGN17]. The illumination is used together with the BRDF to select the direction of the next bounce of a light path.

In the applications of these methods, the sampled directions will always be valid. And in most cases, they are indeed the directions with most energy arriving at the shading point. However, these methods would often fail in many-light rendering. The discrete many lights are essentially very high-frequency contents on the hemisphere, but the sampled directions are from a low-frequency estimation of the illumination. So, a lot of samples will not hit the lights, and will be discarded immediately, resulting in high variance.

3.4. An analysis of multiple importance sampling

To alleviate the above problems in many-light rendering, existing methods [VKK18, EK18] use MIS that combines light sampling and BRDF sampling, to compensate the ignorance of BRDF information during light sampling. These approaches sometimes would work, but often fails in many-light rendering. This is because MIS only works when at least one of its component would work. Unfortunately, in addition to the inefficient light sampling, the BRDF sampling itself also fails frequently, since it cannot easily find a direction that hits the discrete light sources.

As Fig. 2 shows, when the light sources are small or far from the shading point, the directions selected using BRDF sampling can hardly hit these light sources. So the BRDF sampling technique will fail in this case. In the light sampling procedure, since no BRDF information is considered, it may often result in lights that are away from the BRDF lobe, thus contribute little. Hence, combining these two sampling techniques using MIS will not help much, since neither of them work well alone.

Motivated by the difficulties of light sampling and BRDF sampling as well as their MIS combination, we come up with our BRDF-oriented light sampling method that selects lights according to their importance that estimates their contributions to the shading point.

4. BRDF-Oriented Multiple Importance Sampling

As analyzed in Sec. 3, inefficiency can happen simultaneously in both light sampling and BRDF sampling. This problem indicates that, we need a sampling method that is reliable in most cases. In this section, we introduce our BRDF-oriented light sampling, and describe how to apply our new sampling technique together with traditional methods using multiple importance sampling.

In a high level, our full MIS framework contains the following techniques:

- BRDF-oriented light sampling: our new sampling technique that considers BRDF, light intensity and geometry term together as the sample importance.
- Traditional light sampling: a defensive sampling technique similar to that in light tree splitting [EK18], considers light intensity and geometry term, but does not consider BRDF.
- Traditional BRDF importance sampling: a classic sampling technique that considers only BRDF.

4.1. BRDF-oriented light sampling

Hierarchical light tree construction. We construct a hierarchical light tree for organizing all lights. The light tree construction happens only once for the entire scene. Specifically, we start from a tree node that contains all the lights, and construct the light tree from top to bottom, similar to building a Bounding Volume Hierarchy (BVH). We split the nodes using the SAOH metric similar to [EK18]. The SAOH metric is based on the classic partition metric Surface Area Heuristic (SAH) and includes two additional weights in relation to bounding cone and emit energy. It tends to keep lights in a cluster when they have similar spatial positions, orientations and lighting energy. When the light tree is constructed, every node stores a spatial bounding sphere centered at C_C , an orientation cone O_C which is an angular bounding cone of all the contained lights' surface normals, and the total energy E_C of all the lights within. Each light's energy is defined as the maximum emitted radiance integrated over its total area.

Light cut construction. After the light tree is constructed, during runtime, we traverse the light tree from top to bottom for each shading point to form a light cut [WFA*05]. We name this process as light cut construction.

During light cut construction, we propose two criteria to determine whether a node in the light tree (a light node) should be further split or not: First, when the shading point is inside the bounding sphere of the light node, we always split since the light node covers a full-sphere solid angle that is too large as seen from the shading point. The other criterion is that, we also split a light node when the integral of the cosine weighted BRDF within the solid angle subtended by its bounding sphere Ω_C is larger than a predefined threshold δ . Intuitively, this is because we would like to form smaller light clusters where the BRDF is large. This split criterion is formally computed as:

$$S_C(\mathbf{x}) > \delta, \text{ where } S_C(\mathbf{x}) = \int_{\Omega_C} B(\mathbf{x}, \omega_i, \omega_o) \cos \theta_{\mathbf{x}} d\omega_i. \quad (5)$$

In practice, we find that setting $\delta = 0.5$ produces a reasonable




Figure 3: Evaluate the sample importance of a light cluster C

amount of clusters throughout all our test cases. We use Spherical Pivot Transformed Distributions (SPTDs) [DHB17] to quickly compute an approximation of the integral. Other solutions, such as the Linear Transformed Cosines (LTCs) [HDHN16] could give more accurate results, but with much more overhead.

Light cluster sampling. For each shading point, once we have determined the light cut, i.e. a set of light clusters, we have two choices to estimate their direct illumination. The first choice is to sample a light for each light cluster, then adds up the contributions of all light clusters [EK18]. The second choice is to sample a light cluster based on discrete importance values of light clusters, then sample a light from the selected cluster [VKK18].

These choices do not make big differences in previous works. However, for scenes with glossy BRDFs, since large BRDF values are usually observed within a small solid angle, lights in other regions may easily result in little contributions. This observation indicates that it is superior to use the second choice in our case, i.e., sample a light cluster instead of summing up contributions of all clusters.

Then the question is how to determine the importance of a specific light cluster. As introduced earlier in Sec. 3.1, ideally, we would like to relate the importance of a cluster C to its contribution to the shading point \mathbf{x} . Thus, we define the importance of each cluster as

$$I_C(\mathbf{x}) = E_C \cdot G_C \cdot B_C \\ = \frac{E_C \cdot \overline{\cos \theta_C}}{d(\mathbf{x}, C_C)^2} \cdot \frac{1}{\Omega_C} \int_{\Omega_C} B(\mathbf{x}, \omega_i, \omega_o) \cos \theta_{\mathbf{x}} d\omega_i, \quad (6)$$

where E_C is the total energy of the cluster as pre-computed during light tree construction, G_C is estimation of the geometry term, and B_C is the averaged BRDF over the solid angle Ω_C of the cluster's bounding sphere (as shown in Fig. 3). The product of these three terms gives us a good estimation of the contribution of a cluster of lights in Eqn. 2.

Also note that we split the geometry term, so that the integration of the BRDF match the form in Eqn. 5 thus saving computation time. Then the remaining terms $d(\cdot)$ is the distance between the shading point \mathbf{x} and the center of the cluster C_C , and $\overline{\cos \theta_C}$ is the upper bound of all the dot products between any direction in the orientation cone O_C and any direction from the surface of




Figure 4: Our BRDF-oriented light sampling technique may sometimes be affected too much by a large BRDF term that magnifies the inaccuracy in estimating a light cluster’s importance. In (a), cluster C_1 may actually contribute less than C_2 , but was given more samples as predicted in (b), green curve.




Figure 5: Combining our BRDF-oriented light sampling with traditional light sampling increases robustness and sample quality. (a) BRDF-oriented light sampling only; (b) BRDF-oriented light sampling + traditional light sampling. Images are rendered in equal time (24s).

the bounding sphere to the shading point x . As Fig. 3 illustrates, $\frac{\cos\theta_C}{\cos\theta_C}$ is essentially $\cos(\max(\theta - \theta_o - \theta_u, 0))$, as estimated similarly in [EK18].

With the importance metric defined for each light cluster, we can now importance sample a light cluster with a probability proportional to its importance. We perform the light cluster sampling recursively similar to [EK18]. That is, once a cluster has been selected, we continue to select its sub-clusters in the light tree. Finally, we stop at a leaf node that contains only one light.

Light sampling. When a light is finally selected in the light cluster sampling step, the final procedure of our BRDF-oriented light sampling technique is to randomly sample a point on the light (a light point). Similar to [SWZ96, PJH16], we uniformly sample on the light’s surface area. Now, we have successfully sampled a light point by taking the BRDF into consideration.

4.2. Combining sampling techniques

With our proposed BRDF-oriented light sampling technique, in most cases we already have better sampling quality as compared




Figure 6: Combining with BRDF sampling brings even less noise, especially when the scene contains large area lights. (a) BRDF-oriented light sampling + traditional light sampling; (b) BRDF-oriented light sampling + traditional light sampling + traditional BRDF sampling. Images are rendered in equal time (8s).

to traditional light sampling and BRDF sampling. Fig. 14 shows a comparison of rendering using these methods individually.

However, the sampling quality can be further improved by combining our BRDF-oriented light sampling with other defensive techniques [GKPS12]. Since we introduce approximations, our estimation of each cluster’s importance could be inaccurate. Fig. 4 shows one possible case where our method tends to “trust” the BRDF more than other terms, such as energy and geometry. When inaccuracy happens, it is not robust that we simply rely on our BRDF-oriented light sampling.

To increase the robustness of our BRDF-oriented light sampling, we combine it with traditional light sampling and BRDF sampling. The combination with traditional light sampling improves our method’s robustness, since it alleviates our dependence on BRDF. Specifically, we refer to the light sampling method proposed by [EK18]. As Fig. 5 indicates, when combined with traditional light sampling, we achieve a much better sampling quality. And the combination with BRDF sampling is especially useful when there are large area lights thus BRDF sampling becomes very efficient, as shown in Fig. 6.

Finally, in order to balance these sampling techniques, we need to calculate the pdf of our BRDF-oriented light sampling. This is simply the product of a series of conditional probabilities when we select the light clusters, multiplied by the uniform pdf that we use to sample a light point.

With our full MIS framework that combines different sampling techniques, we now have a robust solution of rendering with many lights. Note that, for BRDFs with only a constant diffuse term, our BRDF-oriented light sampling becomes traditional light sampling. For efficiency, in our implementation we use only two sampling techniques (traditional light sampling and BRDF sampling) for such degenerated cases.

5. Adaptive Sample Allocation

In Sec. 4, we have proposed our BRDF-oriented light sampling method, and we combine it with traditional light sampling and BRDF sampling with MIS. Although it is simple and convenient to allocate the same number of samples to each sampling technique, a better sample allocation scheme will always be useful — suppose when a specific sampling method does not work at a certain shading point, we should be able to decrease its sample count for lower variance, and vice versa.

There are generally two kinds of sample allocation methods. One is based on variance [LPG13, HS14, SHSK16, SH17, SHSKE18]. These methods often require a large (up to hundreds) amount of pre-allocated samples, then analyze each sampling technique's variance theoretically. However, these methods are not convenient in many-light rendering, since usually only a small amount of direct illumination samples are drawn per shading point for more path samples per pixel. For this reason, we refer to the second kind of sample allocation based on heuristics that adaptively allocates samples to different techniques.

Our first insight is that, since both the BRDF sampling and our BRDF-oriented light sampling takes into account the BRDF, it is helpful that we explicitly find scenarios where one of these two methods fails but the other still works. Then we can allocate more samples to the more efficient technique. As analyzed in Sec. 3, we have shown that when a direction sample is drawn according to the BRDF, it is difficult that direction hits a light source, especially when the lights are small. This indicates that we should depend more on our BRDF-oriented light sampling in this case.

We adaptively determine which sampling technique to use (i.e., traditional BRDF sampling or BRDF-oriented light sampling) for each sample. We keep two variables n_{hit} , n_{nohit} , recording the accumulated times that the sampled rays from traditional BRDF sampling succeed or fail to hit a light source, respectively. The two variables are both initialized as one, and they are constantly updated during the sample allocation process. Intuitively, we should give more samples to BRDF-oriented light sampling if traditional BRDF sampling always fail to hit. Hence, we select BRDF-oriented light sampling with a probability p proportional to the rate of failed hits:

$$p = n_{nohit} / (n_{nohit} + \lambda \cdot n_{hit}), \quad (7)$$

Note that traditional BRDF sampling will still miss lights frequently even in cases when it is better than BRDF-oriented light sampling. Hence, we empirically set $\lambda = 10$ to compensate it. The pseudocode of our sample allocation process is given in Algorithm 1. Fig. 7 (a) visualizes the sample ratio of traditional BRDF sampling.

Then, the remaining question is how to balance the samples allocated between BRDF-aware methods (traditional BRDF sampling and our BRDF-oriented light sampling) and the non-BRDF-aware method (traditional light sampling). For convenience, we fix the number of samples allocated to traditional light sampling as 1, since it is actually summing up samples from each light cluster and is thus essentially using quite a lot of samples already. Then we just need to find a proper number of samples for the BRDF-aware methods in total.

Algorithm 1 Sample allocation between BRDF-oriented light sampling and BRDF sampling

```

1: procedure BRDF_AWARE_SAMPLING(TotalSampleNum)
2:    $n_{nohit} \leftarrow 1$ 
3:    $n_{hit} \leftarrow 1$ 
4:    $u \leftarrow \text{SAMPLE1D}()$ 
5:   for  $i = 1$  to TotalSampleNum do
6:      $p \leftarrow n_{nohit} / (n_{nohit} + \lambda \cdot n_{hit})$ 
7:     if  $u < p$  then
8:       BRDF_ORIENTED_LIGHT_SAMPLING()
9:        $u \leftarrow u/p$ 
10:    else
11:       $isHit \leftarrow \text{BRDF_SAMPLING}()$ 
12:      if  $isHit$  then
13:         $n_{hit} \leftarrow n_{hit} + 1$ 
14:      else
15:         $n_{nohit} \leftarrow n_{nohit} + 1$ 
16:      end if
17:       $u \leftarrow (1 - u) / (1 - p)$ 
18:    end if
19:  end for
20: end procedure

```




Figure 7: The Car Exhibition scene showing (a) the rate of the sample count of BRDF sampling compared to total BRDF-aware sampling methods, and (b) the importance of BRDF at each shading point used to determine the sample counts of BRDF-aware methods.

So, our second insight is that, we should allocate more samples to the BRDF-aware methods only when the BRDF is at least important to one light. Instead of traversing every light, we find it very effective to define the importance using the upper bound of the BRDF integration within any of the clusters:

$$I_B(\mathbf{x}) = \max\{S_C\}, \quad (8)$$

where C is any cluster within the selected light cut, S_C is the previously defined split criterion of a cluster in Eqn. 5, which is exactly the BRDF integration that we need.

When the integral upper bound I_B is small, we immediately know that the BRDF has a limited contribution to the shading point \mathbf{x} . From Fig. 7 (b), we can see that those regions with a high value of I_B correspond to glossy highlights very well. So, we convert the integral upper bound into an actual number of samples in a linear fashion:

$$N = N_{\max} \cdot I_B, \quad (9)$$

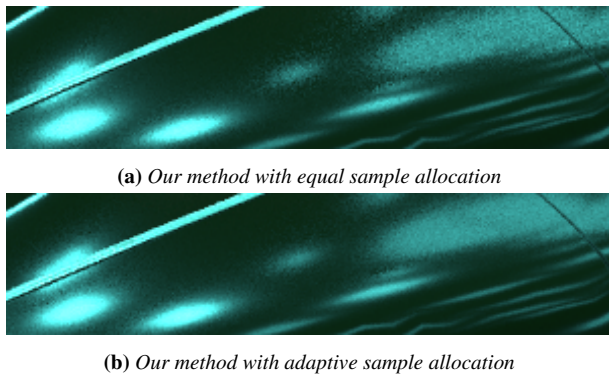


Figure 8: The same crop from the Car Exhibition scene that compares (a) equal and (b) adaptive sample allocation methods. Thank to the adaptive determination and allocation of samples for different sampling techniques in MIS at each shading point, our adaptive strategy achieves even better results on both small and large high-lights. Images are rendered in equal time (110s).

where N_{\max} is a pre-defined maximum number of samples. In practice, we set it as 50 for all our test scenes, which gives reasonable average sample counts to BRDF-aware methods (see Table 1).

With the given heuristics above, now we are able to determine the total number of samples for BRDF-aware methods, then distribute these samples between our BRDF-oriented light sampling and traditional BRDF sampling. In Fig. 8 and Fig. 11, we compare our MIS solution with and without adaptive sample allocation, to demonstrate that our adaptive sample allocation further improves performance. We show more results and comparisons in Sec. 6.

6. Results and Evaluation

In this section, we show rendering results using our model, and compare with previous methods. Our method is implemented using the PBRT renderer [PJH16]. All the results are generated on a PC with an Intel Core i9-9900K CPU with 16 cores, 16 threads and 32 GB of RAM.

6.1. Results

We show our results (direct illumination only) on 5 different test scenes. Parameters of the scenes, including resolution, complexity of meshes and lights are listed in details in Table 1, along with statistics such as averaged light cut size, average sample counts for BRDF-aware methods, rendering time, and the percentage of timing for each sampling technique. We describe each scene and its related comparisons below.

Buddha. The Buddha scene consists of a buddha model and a (triangle-tessellated) spiral curve around it as the light source. This scene contains complex geometry and light visibility. As we can see in Fig. 12, although [VKK18] is specifically designed to handle the lights' visibility, our method still outperforms it, thanks to our BRDF-oriented light sampling method. In Fig. 9, we also compare




Figure 9: Logarithm plot of RMSE as a function of render time for different methods used in the Buddha scene. Our adaptive BRDF-oriented multiple importance sampling converges much faster than state of the art methods.

the convergence curve of our method against state of the art methods [VKK18] and [EK18], and show that we achieve much faster convergence.

Car Exhibition. This scene contains a glossy car in the center, illuminated by many disk lights around and above the car, together with a large area light on the ceiling. We can see from Fig. 1 that our method works reasonably well for both large and small lights, while [EK18] cannot efficiently deal with small lights. The difference verifies our analysis in Sec. 3 that when BRDF sampling does not work well, traditional MIS will not help much.

Cornell Box. We modify the classic Cornell Box scene by replacing the boxes inside it with an armadillo and a bunny, both are assigned with plastic materials that exhibit diffuse and glossy appearances simultaneously. We also place many tiny spherical lights on the left and right walls. As Fig. 14 shows, traditional BRDF sampling barely works in this case. In contrast, our proposed BRDF-oriented light sampling itself outweighs not only the traditional light sampling and BRDF sampling methods individually, but also their combination using MIS [EK18]. Moreover, our BRDF-oriented multiple importance sampling is already able to converge to the reference very fast, even without enabling our adaptive sample allocation.

Lamp. The Lamp scene consists of tens of thin curves (also tessellated) as light sources growing out from a box. The box creates a region on the ground with significant occlusion of lights, which we use to test the robustness of our method. As Fig. 11 shows, using only our BRDF-oriented light sampling is not enough (as analyzed in Sec. 3), although it already produces better reflection than light tree splitting [EK18] in non-occluded regions. However, when we combine it with defensive light sampling and BRDF sampling with MIS, the robustness of our method significantly improves. With our adaptive sample allocation, the result is even better.

Staircase. The Staircase scene contains two large area lights and a few small lights. We use this scene to demonstrate the consistency of our method's performance and robustness even for simple scenes (Fig. 13). Our adaptive sample allocation scheme prefers traditional BRDF sampling for computing reflections from large area lights (see left side), and prefers BRDF-oriented light sampling for

Table 1: Statistics of our test scenes. For each scene, from left to right, we provide: image resolution, the number of lights, the number of triangles, average light cut size in traditional light sampling, average light cut size in BRDF-oriented light sampling (taken over glossy shading points only), average sample count of BRDF-aware sampling techniques (taken over glossy shading points only), rendering time, break-down timings for BRDF-oriented light sampling, traditional light sampling, traditional BRDF sampling and others.

	Res.	#Lights	#Triangles	Cut size in TL	Cut size in BL	#Samples	Time(s)	Percentage of Timings(%)
Buddha	1024×1024	6912	1094.6k	8.43	22.88	14.57	25	24.7 / 60.2 / 4.6 / 10.5
Car Exhibition	1280×720	43486	1825.6k	11.29	5.07	15.06	55	22.1 / 68.6 / 3.9 / 5.4
Cornell Box	1024×1024	1494	356.6k	10.35	8.58	6.68	10	14.4 / 60.17 / 6.0 / 19.4
Lamp	1024×1024	53760	54.2k	10.73	20.30	4.79	36	26.4 / 63.2 / 1.7 / 8.7
Staircase	1024×1024	21	30.9k	8.04	7.51	23.60	7.5	33.7 / 30.4 / 18.9 / 17.0

Table 2: Statistics with different δ for rendering the Car Exhibition scene. From left to right, we provide: value of δ , average light cut size in BRDF-oriented light sampling, average sample count of BRDF-aware sampling methods, RMSE and rendering time.

δ	Avg. cut size	#Avg.sample	RMSE	Time(s)
0.005	50.87	3.18	0.106008	29.1
0.05	13.49	4.41	0.106646	26.8
0.1	9.48	5.85	0.105999	28.5
0.3	3.68	11.04	0.102140	41.0
0.5	3.40	15.09	0.099717	45.2
0.7	3.35	18.32	0.097589	53.4

computing reflections from small lights (see the highlight on the right side).

6.2. Evaluation

Parameter δ in Eqn. 5. In our BRDF-oriented light sampling, the number of light clusters (i.e., cut size) decreases as δ becomes larger, and vice versa. On one hand, a larger number of light clusters will increase the time of light cluster sampling and slightly improve rendering quality. On the other hand, a smaller number of light clusters will result in larger BRDF integral upper bound (Eqn. 9), hence leading to more samples and more time for BRDF-aware sampling methods. In Table 2, we list statistics of rendering time, rendering quality (in terms of RMSE) using different δ for the Car Exhibition scene. Overall, we find setting $\delta = 0.5$ provides a good trade-off between rendering quality and efficiency.

Scalability. We also evaluated the scalability of our method with respect to the number of lights. To do so, we render the modified Cornell Box scene (as shown in Fig. 14) using different number of tiny spherical lights under the same parameter settings. For each number, we randomly place the given number of lights in specific regions on the left and right walls and record the rendering time. Fig. 10 shows how it changes with the number of lights. The rendering time increases from 11s to 20s when the number of lights increases from 1,000 to 100,000. It demonstrates that our method scales well with the number of lights.




Figure 10: Rendering time with respect to the number of lights for the modified Cornell Box scene. Rendering time increases only 2 times (i.e., from 11s to 20s) when the number of lights increases 100 times (i.e., from 1,000 to 100,000).

7. Conclusion and Future Work

We present BRDF-oriented light sampling that selects lights with regard to BRDF's contribution for many-light rendering, together with a full Multiple Importance Sampling (MIS) framework. We further propose an adaptive sample allocation heuristic that dynamically determines the number of samples for different sampling techniques within MIS. Our sampling method produces significantly lower variance than other methods, and we achieve much faster convergence to the reference as compared to previous many-light rendering algorithms.

Our method still has limitations which are worth investigating in the future. First, the evaluation of the importance of BRDFs in our algorithm relies on existing analytic area lighting methods (LTCs or SPTDs), whose efficiency and accuracy may be further improved. Secondly, after a light is selected, our BRDF-oriented light sampling technique samples a point on the light through uni-

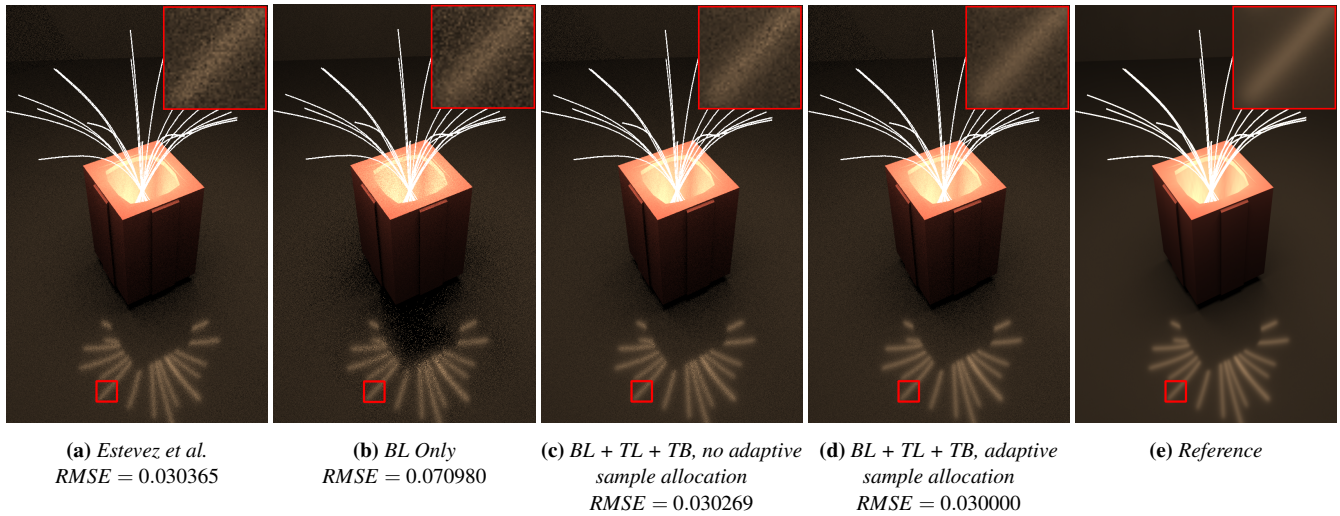


Figure 11: Equal-time comparison (36s) of our adaptive sampling method and Estevez et al. [EK18]. (a) Estevez et al.; (b) Our method with BRDF-oriented light sampling only; (c) Our method with MIS of all three sampling techniques, but without adaptive sample allocation; (d) Our method with MIS of all three sampling techniques and adaptive sample allocation. (e) Reference. The 3D Models are courtesy of BlendSwap.com user shmuel245.




Figure 12: Equal-time comparison (25s) of (a) Vévoda et al [VKK18] and (b) our adaptive sampling method with multiple importance sampling. Although [VKK18] is specifically designed to handle the lights' visibility, our method still outperforms it, thanks to our BRDF-oriented light sampling method.

formly random sampling on the light's surface area. A more sophisticated approach may consider BRDF into the point sampling process. Besides, while our adaptive sample allocation scheme provides a nice heuristic in choosing between the two BRDF-aware sampling methods, it is far from optimal.

Additionally, it is straightforward to extend our BRDF-oriented light sampling to general BSDFs, taking refractive materials into account. It is also worth thinking about visibility, either by combining our method with visibility-aware methods [VKK18] that estimates the visibility of light clusters from a shading point, or by

efficiently incorporating visibility tests into our sampling method. Exploring better MIS could also be promising, such as more optimized combination heuristics and more effective sampling allocation techniques. In summary, we believe that our method is an important step towards practical many-light rendering, and is also a good start that contributes to the next generation of rendering with complex lighting condition and materials.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their valuable comments.




Figure 13: Equal-time comparison (7.5s) of (a) Estevez et al. [EK18] and (b) our adaptive sampling method with multiple importance sampling. The scene contains very few lights (see Table 1) but our method is still able to perform slightly better, thus is consistently robust. This scene is courtesy of [Bit16].

This work is supported by the National Natural Science Foundation of China (Project Numbers: 61822204, 61521002), and the National Key Research and Development Program of China (Project Number: 2017YFB1002604).

References

- [Bit16] BITTERLI B.: Rendering resources, 2016. <https://benedikt-bitterli.me/resources/>. 10
- [BMB15] BUS N., MUSTAFA N. H., BIRI V.: Illuminationcut. In *Computer Graphics Forum* (2015), vol. 34, Wiley Online Library, pp. 561–573. 2
- [CAM08] CLARBERG P., AKENINE-MÖLLERY T.: Practical product importance sampling for direct illumination. In *Computer Graphics Forum* (2008), vol. 27, Wiley Online Library, pp. 681–690. 2
- [CJAMJ05] CLARBERG P., JAROSZ W., AKENINE-MÖLLER T., JENSEN H. W.: Wavelet importance sampling: efficiently evaluating products of complex functions. In *ACM Transactions on Graphics (TOG)* (2005), vol. 24, ACM, pp. 1166–1175. 2
- [DHB17] DUPUY J., HEITZ E., BELCOUR L.: A spherical cap preserving parameterization for spherical distributions. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 139. 2, 4
- [DKH*14] DACHSBACHER C., KRÍVÁNEK J., HAŠAN M., ARBREE A., WALTER B., NOVÁK J.: Scalable realistic rendering with many-light methods. In *Computer Graphics Forum* (2014), vol. 33, Wiley Online Library, pp. 88–104. 2
- [EK18] ESTEVEZ A. C., KULLA C.: Importance sampling of many lights with adaptive tree splitting. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 1, 2 (2018), 25. 1, 2, 3, 4, 5, 7, 9, 10, 11
- [GKPS12] GEORGIEV I., KRÍVÁNEK J., POPOV S., SLUSALLEK P.: Importance caching for complex illumination. In *Computer Graphics Forum* (2012), vol. 31, Wiley Online Library, pp. 701–710. 5
- [HDHN16] HEITZ E., DUPUY J., HILL S., NEUBELT D.: Real-time polygonal-light shading with linearly transformed cosines. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 41. 2, 4
- [HEV*16] HERHOLZ S., ELEK O., VORBA J., LENSCH H., KRÍVÁNEK J.: Product importance sampling for light transport path guiding. In *Computer Graphics Forum* (2016), vol. 35, Wiley Online Library, pp. 67–77. 3
- [HS14] HAVRAN V., SBERT M.: Optimal combination of techniques in multiple importance sampling. In *Proceedings of the 13th ACM SIGGRAPH International Conference on Virtual-Reality Continuum and its Applications in Industry* (2014), ACM, pp. 141–150. 2, 6
- [JJC09] JAROSZ W., CARR N. A., JENSEN H. W.: Importance sampling spherical harmonics. In *Computer Graphics Forum* (2009), vol. 28, Wiley Online Library, pp. 577–586. 2
- [Kel97] KELLER A.: Instant radiosity. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1997), SIGGRAPH '97, ACM Press/Addison-Wesley Publishing Co., pp. 49–56. 2
- [LPG13] LU H., PACANOWSKI R., GRANIER X.: Second-order approximation for variance reduction in multiple importance sampling. In *Computer Graphics Forum* (2013), vol. 32, Wiley Online Library, pp. 131–136. 2, 6
- [MGN17] MÜLLER T., GROSS M., NOVÁK J.: Practical path guiding for efficient light-transport simulation. In *Computer Graphics Forum* (2017), vol. 36, Wiley Online Library, pp. 91–100. 3
- [PBPP11] PAJOT A., BARTHE L., PAULIN M., POULIN P.: Representativity for robust and adaptive multiple importance sampling. *IEEE transactions on visualization and computer graphics* 17, 8 (2011), 1108–1121. 2
- [PJH16] PHARR M., JAKOB W., HUMPHREYS G.: *Physically based rendering: From theory to implementation*. Morgan Kaufmann, 2016. 5, 7
- [SH17] SBERT M., HAVRAN V.: Adaptive multiple importance sampling for general functions. *The Visual Computer* 33, 6-8 (2017), 845–855. 2, 6
- [SHSK16] SBERT M., HAVRAN V., SZIRMAY-KALOS L.: Variance analysis of multi-sample and one-sample multiple importance sampling. In *Computer Graphics Forum* (2016), vol. 35, Wiley Online Library, pp. 451–460. 2, 6
- [SHSKE18] SBERT M., HAVRAN V., SZIRMAY-KALOS L., ELVIRA V.: Multiple importance sampling characterization by weighted mean invariance. *The Visual Computer* 34, 6-8 (2018), 843–852. 2, 6




Figure 14: Equal-time comparison (10s) of our adaptive sampling method, Estevez et al. [EK18], and individual sampling techniques. (a) traditional light sampling only; (b) traditional BRDF sampling only; (c) Our method with BRDF-oriented light sampling only; (d) Estevez et al. [EK18]; (e) Our method with MIS of all three sampling techniques. Since both traditional light sampling and BRDF sampling are very inefficient in this case, when combined with MIS in (e), the sampling overhead outweighs the quality gain, and a slightly higher RMSE than (c) in equal time is expected.

[SWZ96] SHIRLEY P., WANG C., ZIMMERMAN K.: Monte carlo techniques for direct lighting calculations. *ACM Transactions on Graphics (TOG)* 15, 1 (1996), 1–36. 2, 5

[Vea97] VEACH E.: *Robust Monte Carlo methods for light transport simulation*, vol. 1610. Stanford University PhD thesis, 1997. 2

[VKK18] VÉVODA P., KONDAPANENI I., KRIVÁNEK J.: Bayesian on-line regression for adaptive direct illumination sampling. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 125. 2, 3, 4, 7, 9

[VKŠ*14] VORBA J., KARLÍK O., ŠIK M., RITSCHER T., KRIVÁNEK J.: On-line learning of parametric mixture models for light transport simulation. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 101. 3

[WÅ09] WANG R., ÅKERLUND O.: Bidirectional importance sampling for unstructured direct illumination. In *Computer Graphics Forum* (2009), vol. 28, Wiley Online Library, pp. 269–278. 2

[WABG06] WALTER B., ARBREE A., BALA K., GREENBERG D. P.: Multidimensional lightcuts. In *ACM Transactions on graphics (TOG)* (2006), vol. 25, ACM, pp. 1081–1088. 2

[War94] WARD G. J.: The radiance lighting simulation and rendering system. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques* (1994), ACM, pp. 459–472. 2

[WC13] WU Y.-T., CHUANG Y.-Y.: Visibilitycluster: Average directional visibility for many-light rendering. *IEEE transactions on visualization and computer graphics* 19, 9 (2013), 1566–1578. 2

[WFA*05] WALTER B., FERNANDEZ S., ARBREE A., BALA K., DONIKIAN M., GREENBERG D. P.: Lightcuts: a scalable approach to illumination. In *ACM Transactions on graphics (TOG)* (2005), vol. 24, ACM, pp. 1098–1107. 2, 4

[XCM*14] XU K., CAO Y.-P., MA L.-Q., DONG Z., WANG R., HU S.-M.: A practical algorithm for rendering interreflections with all-frequency brdfs. *ACM Transactions on Graphics* 33, 1 (2014), 10:1–10:16. 2