SUNG-EUI YOON, KAIST

RENDERING

FREELY AVAILABLE ON THE INTERNET

FREELY AVAILABLE ON THE INTERNET

http://sglab.kaist.ac.kr/~sungeui/render

First printing, July 2018

16
Importance Sampling

In the last chapter, we discussed Monte Carlo (MC) ray tracing, especially, path tracing that generates a light path from the camera to the light source. While it is an unbiased estimator, it has significant variance, i.e., noise, when we have a low ray samples per pixel. To reduce the noise of MC generated images, we studied quasi-Monte Carlo technique in Sec. 15.3.

In this chapter, as an effective way of reducing the variance, we discuss importance sampling. We first discuss an importance sampling method considering light sources, called direct illumination method. We then discuss other importance sampling methods considering various factors of the rendering equation.

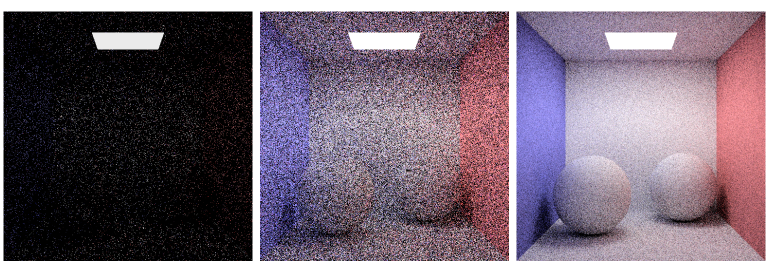

(a) Results w/o direct illumination. From the left, 1 spp, 4 spp, and 16 spp are used.

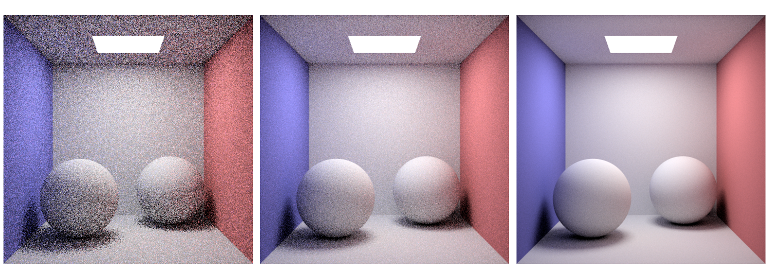

(b) Results w/ direct illumination.

Figure 16.1: These images are generated by path tracer w/ and w/o direct illumination. They are created by using a path tracer created by Ritchie et al. `http://web.stanford.edu/~dritchie/path/index.html`.

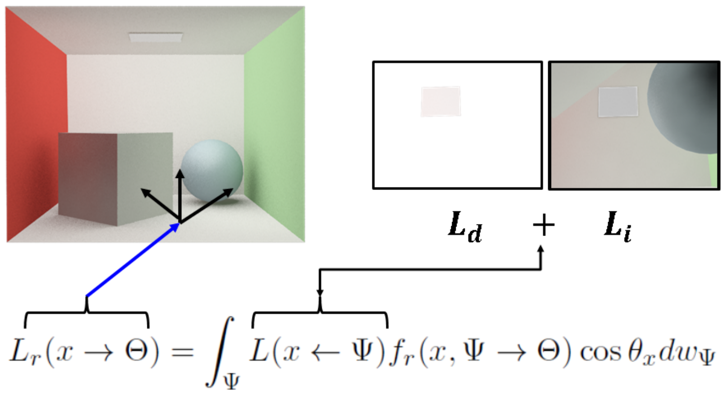

Figure 16.2: This figure illustrates the factorization of the reflected radiance into direct and indirect illumination terms.

16.1 *Direct Illumination*

Fig. 16.1 show rendering results w/ and w/o direct illumination. The first row shows rendering results w/o direct illumination under 1, 4, and 16 spp. In this scene, we adopt path tracing and observe severe noise even when we use 16 spp. This noise is mainly from the variance of the MC estimator. Note that we use random sampling on the hemisphere to generate a reflected ray direction, and it can keep bounce unless arriving at the light source located at the ceiling of the scene. Furthermore, since we are using the Russian roulette, some rays can be terminated without carrying any radiance, resulting in dark colors.

A better, yet intuitive approach is to generate a ray directly toward the light source, since we know that the light source is emitting energy and brightens the scene. The question is how we can accommodate this idea within the MC estimation framework! If we just generate a ray toward the light source, it will introduce a bias and we may not get a correct result, even when we generate an infinite number of samples.

Let's consider the rendering equation that computes the radiance $L(x \to \Theta)$, from a location x in the direction of Θ [1]. The radiance is composed of the self-emitted energy and reflected energy (Fig. 13.1):

$$L(x \to \Theta) = L_e(x \to \Theta) + L_r(x \to \Theta). \tag{16.1}$$

For the reflected term $L_r(\cdot)$, we decompose it into two terms: direct illumination term, $L_d(\cdot)$, and indirect illumination term, $L_i(\cdot)$:

$$L_r(x \to \Theta) = L_d(x \to \Theta) + L_i(x \to \Theta). \tag{16.2}$$

Fig. 16.2 illustrates an example of this decomposition.

Once we decomposed the radiance term into the direct and indirect illumination terms, we apply two separate MC estimators for

[1] This notation is introduced in Sec. 13.1

those two terms. For the direct illumination term, we cannot use the hemispherical integration described in Sec. 13.1, since we need to generate rays to the light source. For generating rays only to the light source, we use the area formulation, Eq. 13.5 explained in Sec. 13.2.

Rays corresponding to the direct illumination should be not duplicated considered for indirect illumination.

For estimating the indirect illumination, we use the hemispherical integration. The main difference to the regular hemispherical integration is that a ray generated from the hemispherical integration should not accumulate energy directly from the light source. In other words, when the ray intersects with the light source, we do not transfer the energy emitted from the light source, since the ray in this case is considered in the direct illumination term, and thus its energy should not be considered for the indirect illumination to avoid duplicate computation.

Many light problems. We discussed a simple importance sampling with the direct illumination sampling to reduce the variance of MC estimators. What if we have so many lights? In this case, generating rays to many lights can require a huge amount of time. In practice, simulating realistic scenes with complex light setting may require tens or hundreds of thousands of point light sources. This problem has been known as the many light problem. Some of simple approaches are to generate rays to those lights with probabilities that are proportional to their light intensity.

16.2 Multiple Importance Sampling

In the last section, we looked into direct illumination sampling as an importance sampling method. While it is useful, it cannot be a perfect solution, as hinted in our theoretical discussion (Sec. 14.3)

There are many other different terms in the rendering equation. Some of them are incoming radiance, BRDF, visibility, cosine terms, etc. The direct illumination sampling is a simple heuristic to consider the incoming radiance, while there could be many other strong indirect illuminations such as strong light reflection from a mirror. BRDF of an intersected object and cosine terms are available, and thus we can design importance sampling methods considering those factors. Nonetheless, these different importance sampling methods are designed separately and may work well in one case, but not in other cases.

Multiple importance sampling (MIS) is introduced to design a combined sampling method out of separately designed estimators. Suppose that there are n different sampling methods, and we allocate n_i samples for each sampling method. Given the total number of samples N, $n_i = c_i N$ with independent $X_{i,j}$ samples. The whole

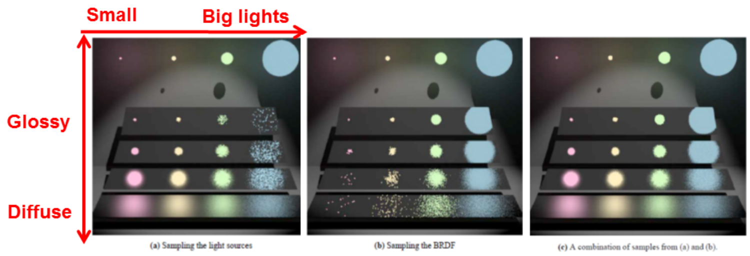

Figure 16.3: These figures show rendering results with different sampling methods. From the left, we use sampling light sources, BRDF, and both of them w/ multiple importance sampling.

distribution, $\bar{p}(x)$, combined with those n different methods, is defined as the following:

$$\bar{p}(x) = \sum_{i}^{n} c_i p_i(x), \tag{16.3}$$

where $p_i(x)$ is a i-th sampling distribution. $\bar{p}(x)$ is also called combined sample distribution [2], whose each sample $X_{i,j}$ has $1/N$ sampling probability.

[2]

By applying the standard MC estimator with the combined sampling distribution, we get the following estimator:

$$I = \frac{1}{N} \sum_{i} \sum_{n_i} \frac{f(X_{i,j})}{\bar{p}(X_{i,j})}. \tag{16.4}$$

This estimator is also derived by assigning the relative importance, i.e., probability, of a sampling method among others. In this perspective, this is also known as to be derived under balance heuristic. Surprisingly, this simple approach has been demonstrated to work quite well as shown in Fig. 16.3; these figures are excepted from the paper of Veach et al. [3]. A theoretical upper bound of the variance error of this approach is available in the original paper.

[3]