Computer graphics III – Multiple Importance Sampling

Jaroslav Křivánek, MFF UK

Jaroslav.Krivanek@mff.cuni.cz

Multiple Importance Sampling in a few slides

Motivation

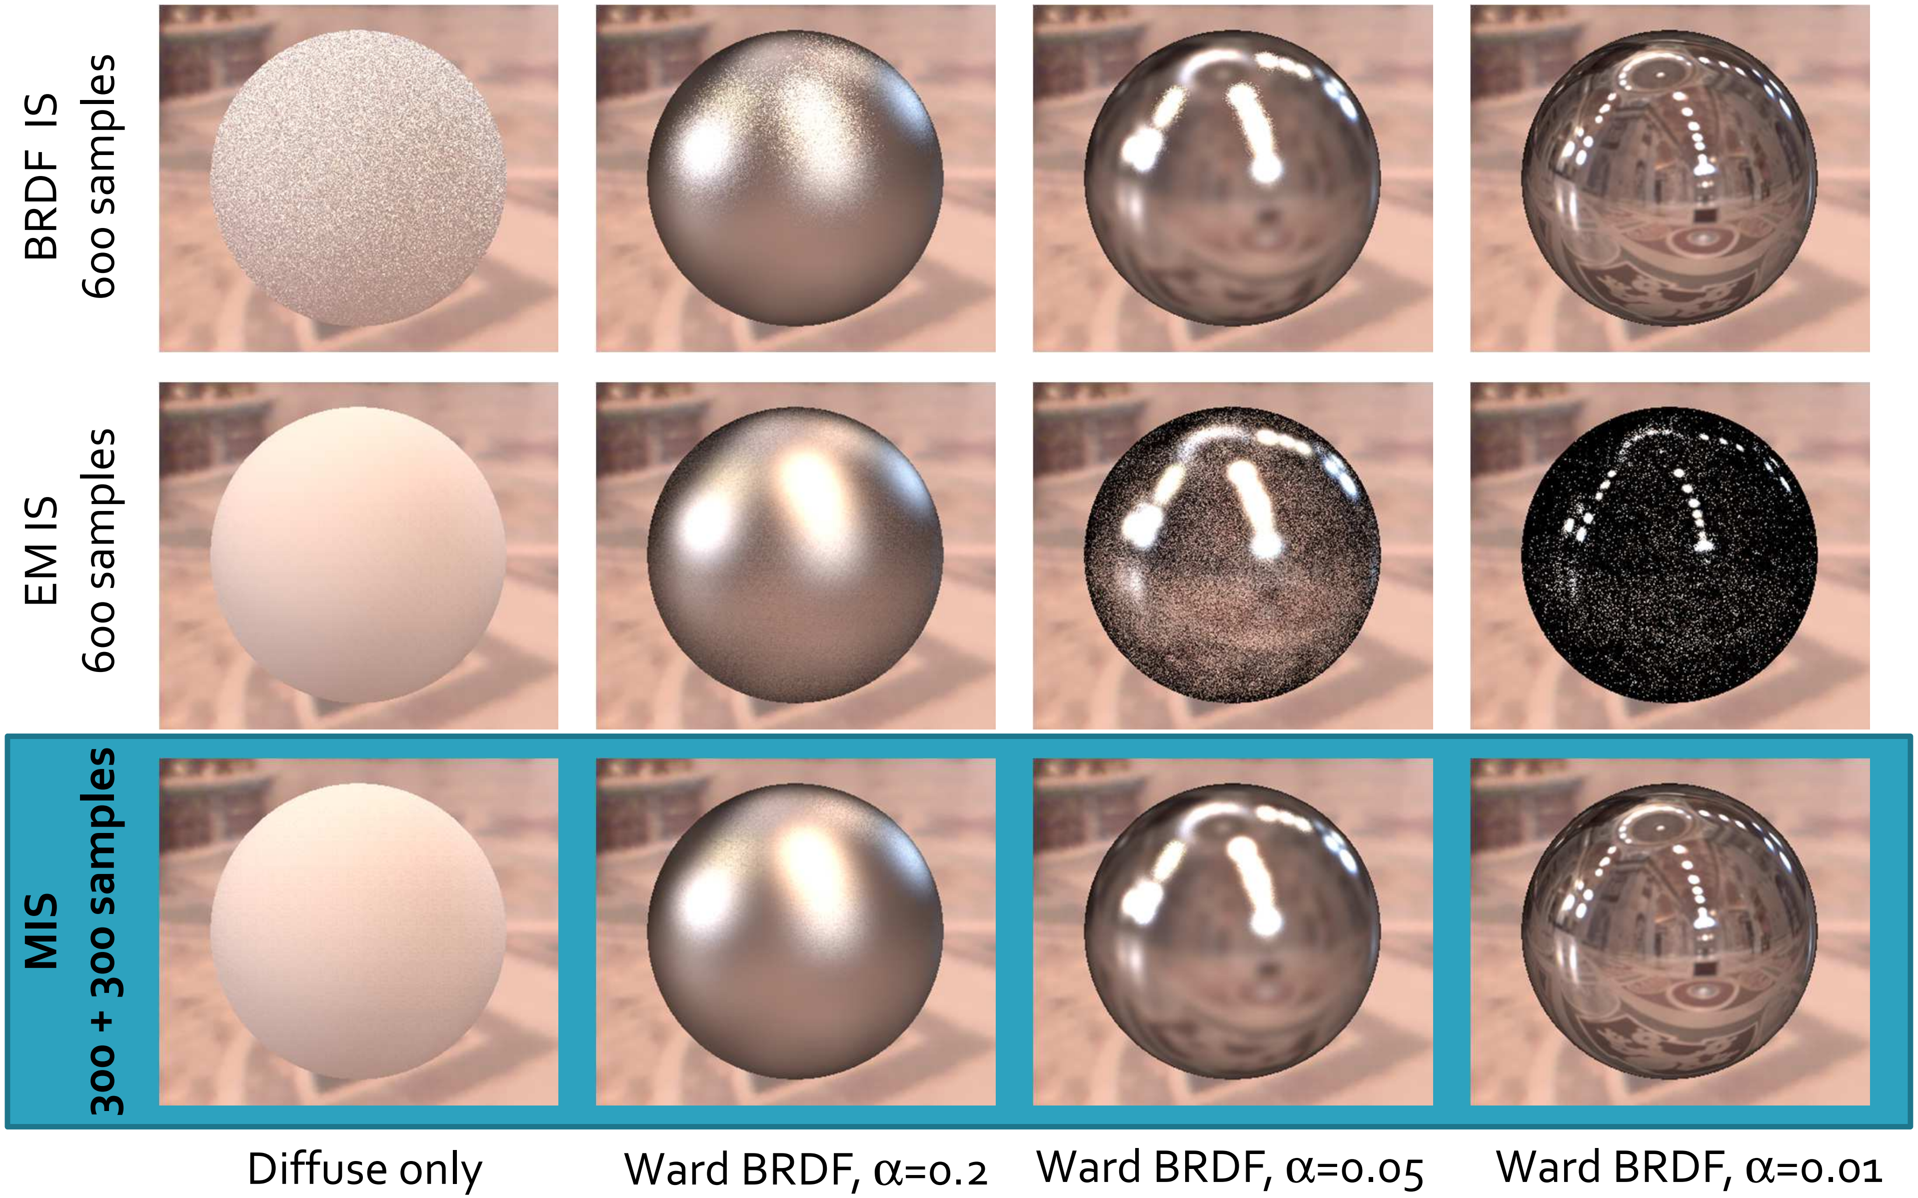
BRDF IS
600 samples
EM IS
600 samples
MIS
300 + 300 samples
Diffuse only
Ward BRDF, α=0.2
Ward BRDF, α=0.05
Ward BRDF, α=0.01

What is wrong with BRDF and light source sampling?

- **A:** None of the two is a good match for the entire integrand under all conditions

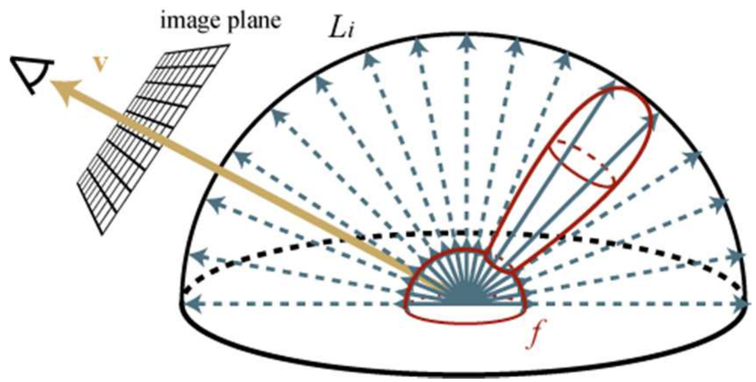

$$L_r(\omega_o) = \int_{H(\mathbf{x})} L_i(\omega_i) \cdot f_r(\omega_i \to \omega_o) \cdot \cos\theta_i \, \mathrm{d}\omega_i$$

Multiple Importance Sampling (MIS)

[Veach & Guibas, 95]

Combined estimator:

$$\langle I \rangle = \frac{f(x)}{[p_a(x) + p_b(x)]/2}$$

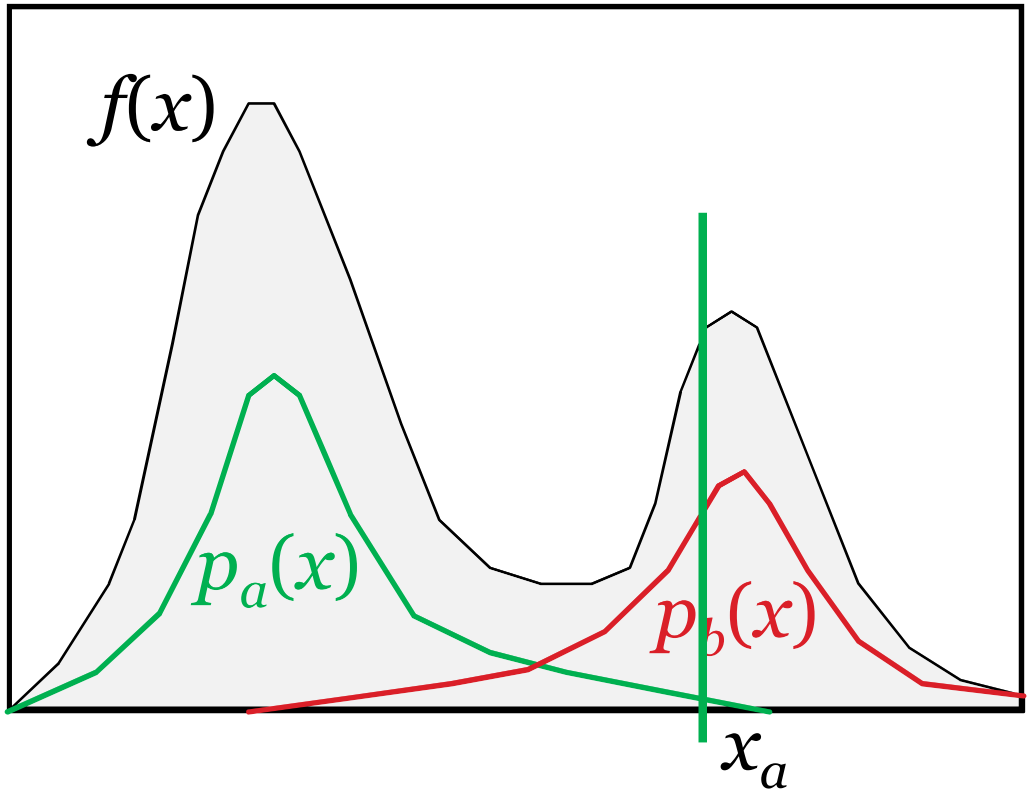

Notes on the previous slide

- We have a complex multimodal integrand $f(x)$ that we want to numerically integrate using a MC method with importance sampling.
- Unfortunately, we do not have a PDF that would mimic the integrand in the entire domain.
- Instead, we can draw the sample from two different PDFs, p_a and p_b each of which is a good match for the integrand under different conditions – i.e. in different part of the domain.
- However, the estimators corresponding to these two PDFs have extremely high variance – shown on the slide.
- We can use Multiple Importance Sampling (MIS) to combine the sampling techniques corresponding to the two PDFs into a single, robust, combined technique.
- The MIS procedure is extremely simple: it randomly picks one distribution to sample from (p_a or p_b , say with fifty-fifty chance) and then takes the sample from the selected distribution.
- This essentially corresponds to sampling from a weighted average of the two distributions, which is reflected in the form of the estimator, shown on the slide.

- This estimator is really powerful at suppressing outlier samples such as those that you would obtain by picking x_from the tail of p_a, where $f(x)$ might still be large.
- Without having p_b at our disposal, we would be dividing the large $f(x)$ by the small $p_a(x)$, producing an outlier.
- However, the combined technique has a much higher chance of producing this particular x (because it can sample it also from p_b), so the combined estimator divides $f(x)$ by $[p_a(x) + p_b(x)] / 2$, which yields a much more reasonable sample value.

- I want to note that what I'm showing here is called the "balance heuristic" and is a part of a wider theory on weighted combinations of estimators proposed by Veach and Guibas.

Application to direct illumination

- Two sampling strategies

1. **BRDF-proportional sampling - p_a**
2. **Environment map sampling - p_b**

… and now the (almost) full story

First for general estimators, so please forget the direct illumination problem for a short while.

Multiple Importance Sampling

(Veach & Guibas, 95)

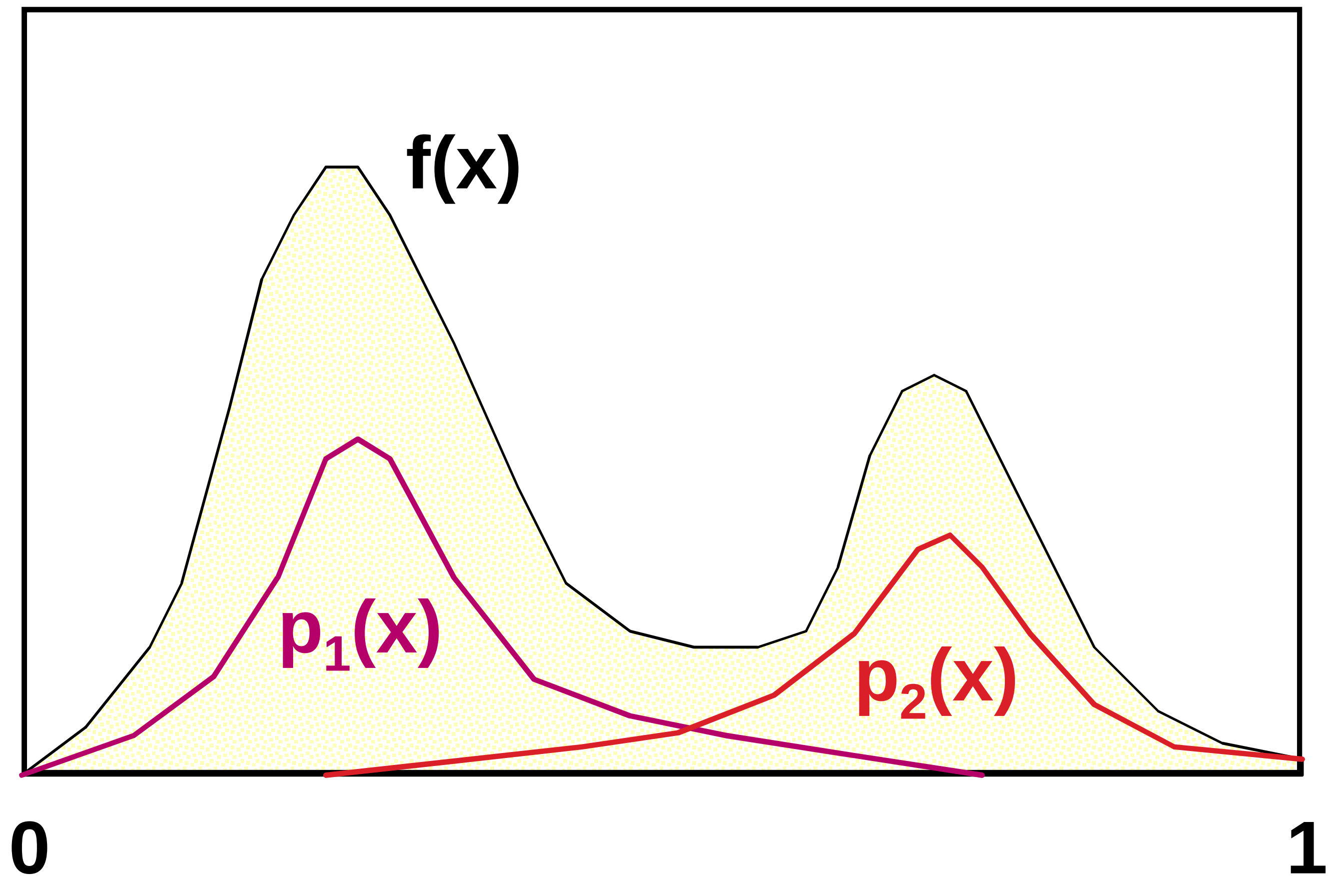

Multiple Importance Sampling

- Given n sampling techniques (i.e. pdfs) $p_1(x), .. , p_n(x)$
- We take n_i samples $X_{i,1}, .. , X_{i,n_i}$ from each technique
- **Combined estimator**

Combination weights
(different for each sample)

$$\boxed{F = \sum_{i=1}^{n} \frac{1}{n_i} \sum_{j=1}^{n_i} w_i(X_{i,j}) \frac{f(X_{i,j})}{p_i(X_{i,j})}}$$

samples from individual techniques

sampling techniques

Unbiasedness of the combined estimator

$$E[F] = \ldots = \int \left[\sum_{i=1}^{n} w_i(x) \right] f(x) \, \mathrm{d}x \equiv \int f(x)$$

- Condition on the weighting functions

$$\boxed{\forall x: \quad \sum_{i=1}^{n} w_i(x) = 1}$$

Choice of the weighting functions

- **Objective:** minimize the variance of the combined estimator

1. Arithmetic average (very bad combination)

$$w_i(x) = \frac{1}{n}$$

2. **Balance heuristic** (very good combination)
   - ....

Balance heuristic

- Combination weights

$$\hat{w}_i(\mathbf{x}) = \frac{n_i p_i(\mathbf{x})}{\sum_k n_k p_k(\mathbf{x})}$$

- Resulting estimator (after plugging in the weights)

$$\boxed{F = \sum_{i=1}^{n} \sum_{j=1}^{n_i} \frac{f(X_{i,j})}{\sum_k n_k p_k(X_{i,j})}}$$

  - i.e. the form of the contribution of a sample does not depend on the technique (pdf) from which it came

Balance heuristic

- The balance heuristic **is almost optimal**
  - No other weighting has variance much lower than the balance heuristic
- Further possible combination heuristics
  - **Power heuristic**
  - **Maximum heuristics**
  - See [Veach 1997]

One term of the combined estimator

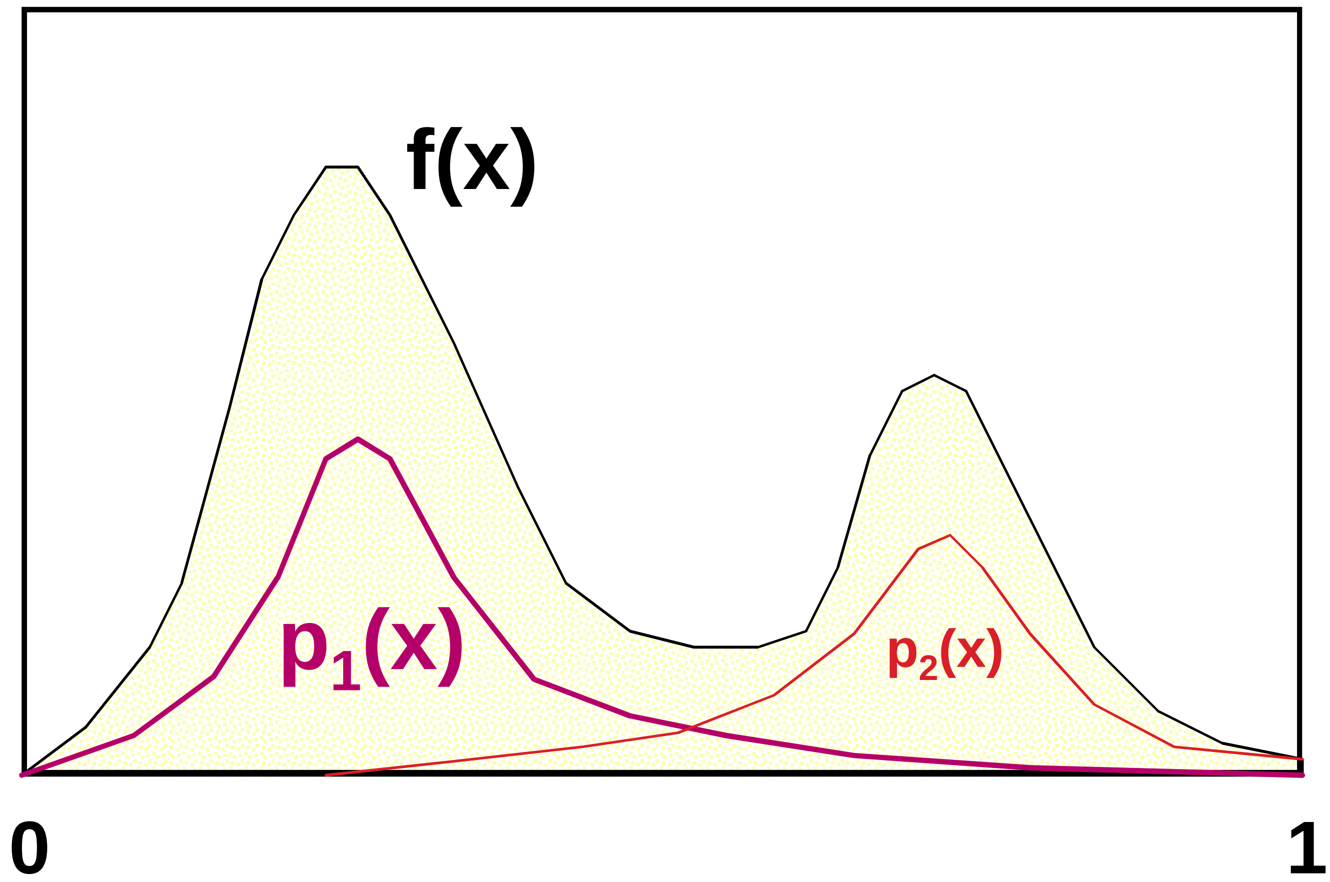

One term of the combined estimator: Arithmetic average

$$0.5\frac{f(x)}{p_1(x)} + 0.5\frac{f(x)}{p_2(x)}$$

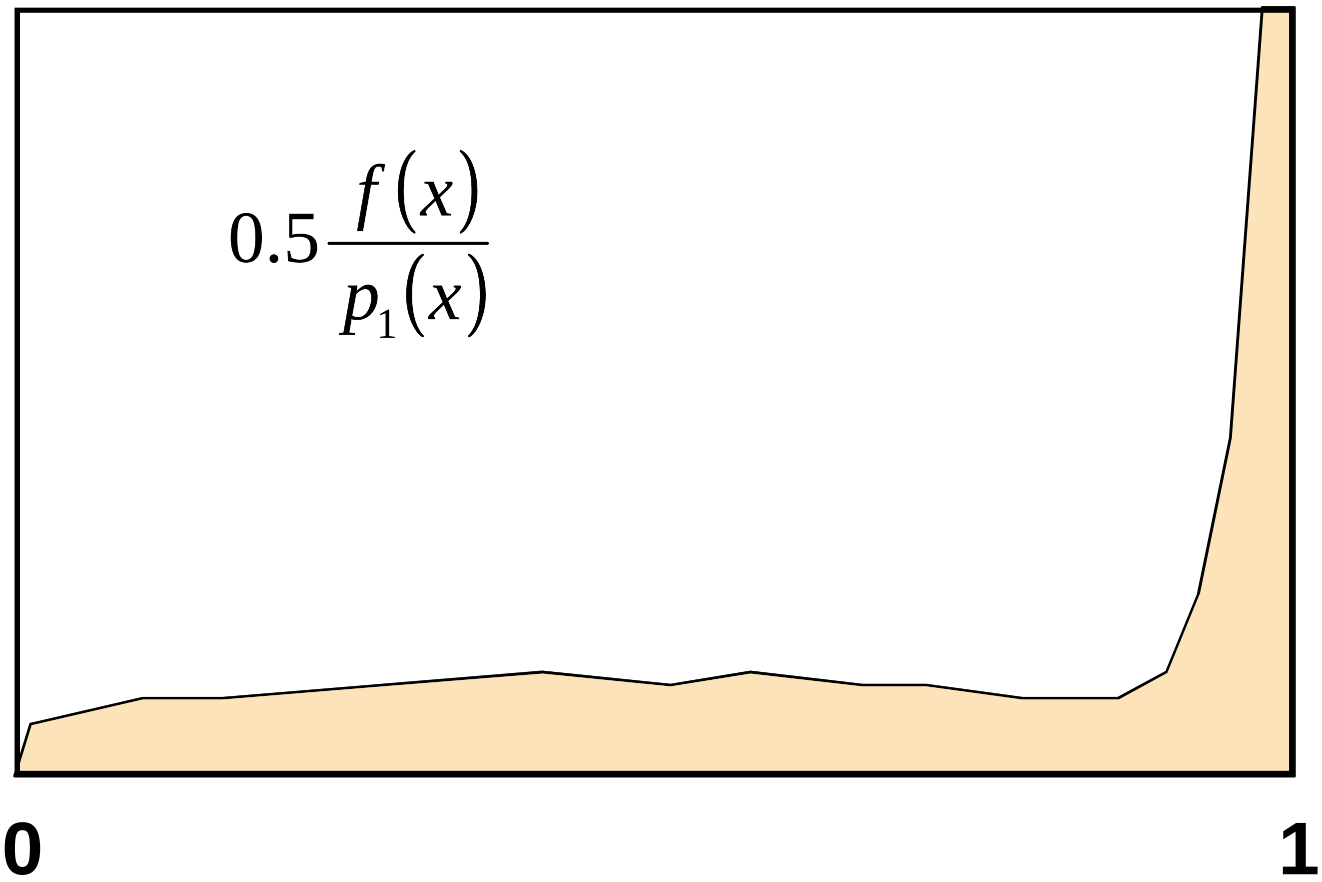

One term of the combined estimator: Balance heuristic

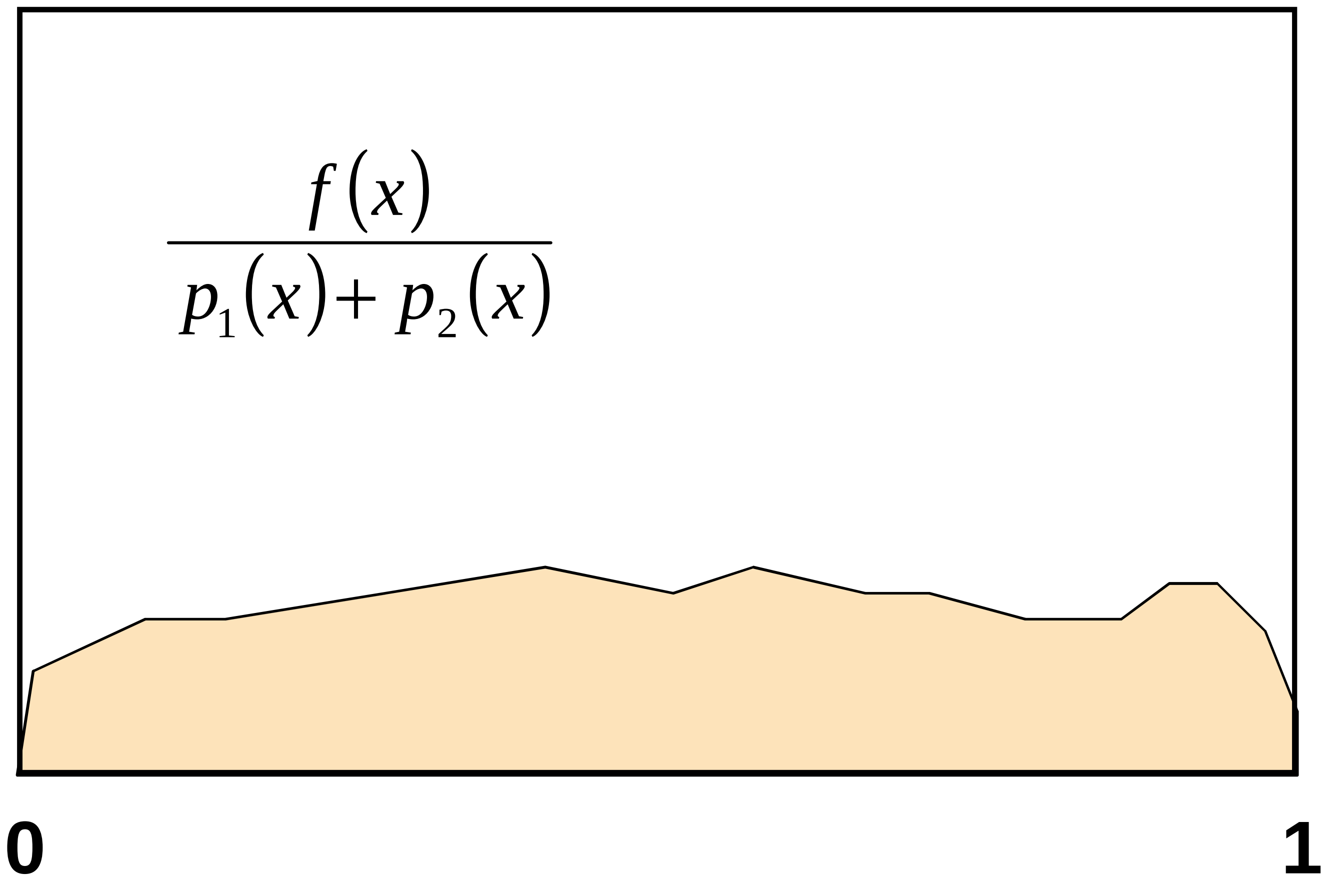

Direct illumination calculation using MIS

We now focus on area lights instead of the motivating example that used environment maps. But the idea is the same.

Problem: Is random BRDF sampling going to find the light source?

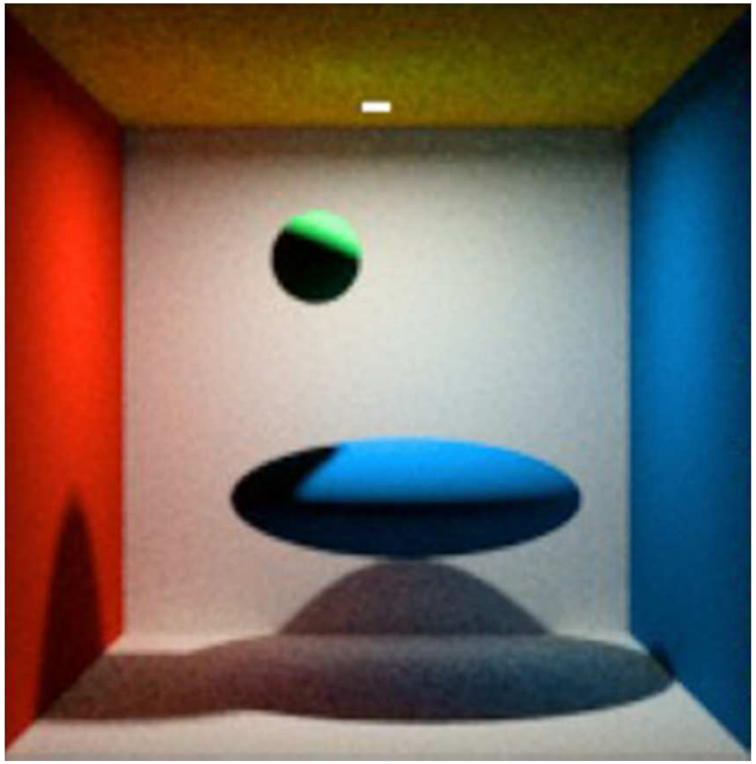

reference

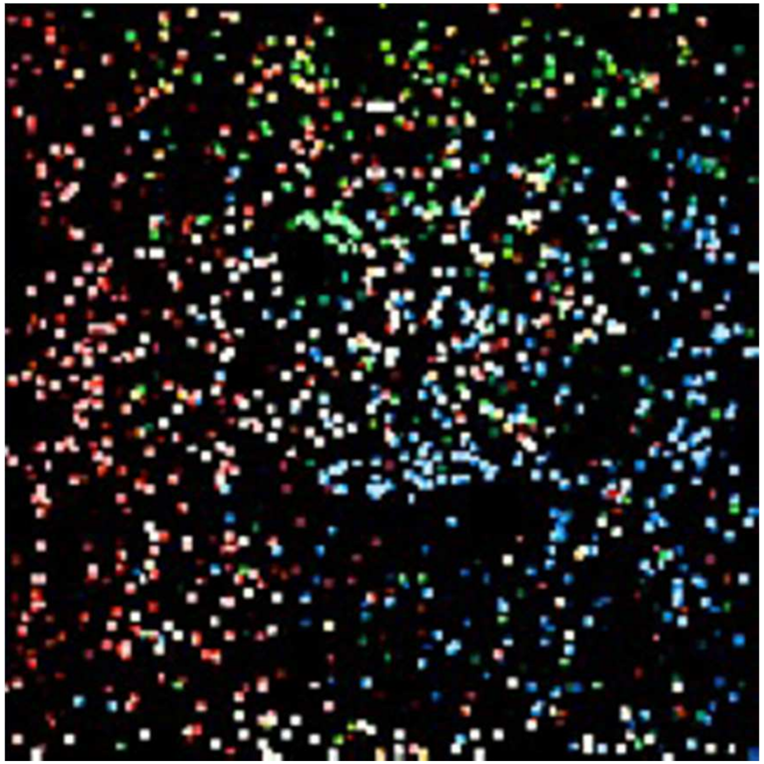

simple path tracer
(150 paths per pixel)

Images: Alexander Wilkie

Direct illumination: Two strategies

- We are calculating **direct illumination** due to a given light source.
  - i.e. radiance reflected from a point **x** on a surface exclusively due to the light coming directly from the considered source

- Two sampling strategies
  1. **BRDF-proportional sampling**
  2. **Light source area sampling**

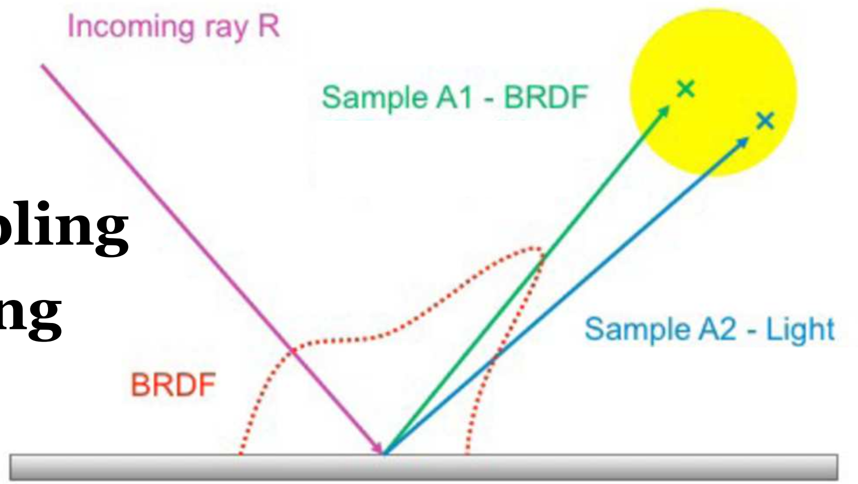

Image: Alexander Wilkie

Direct illumination: Two strategies

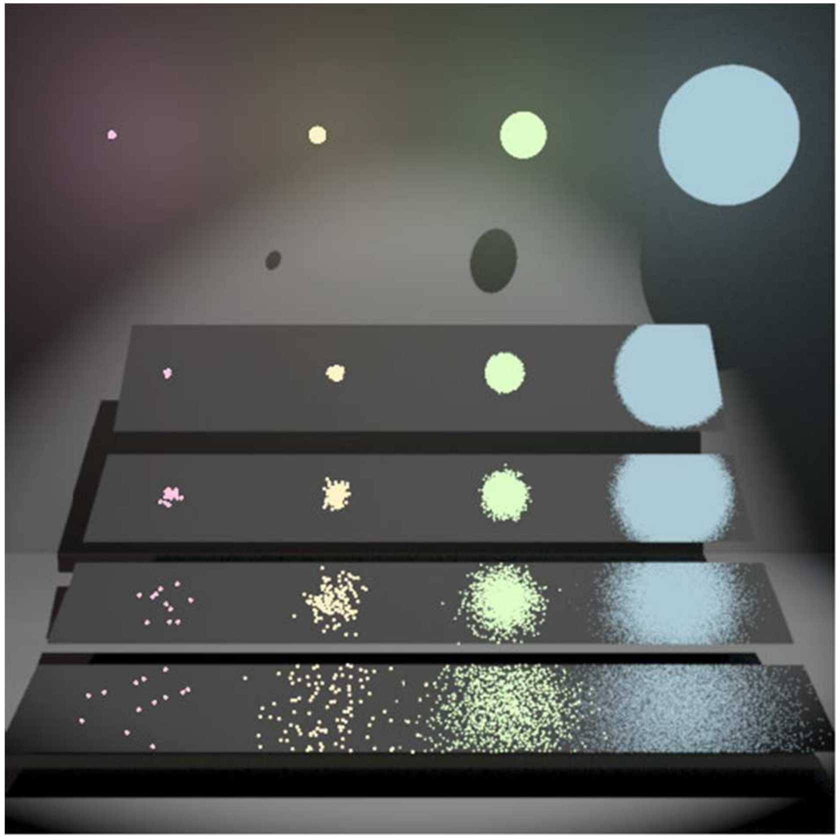

BRDF proportional sampling

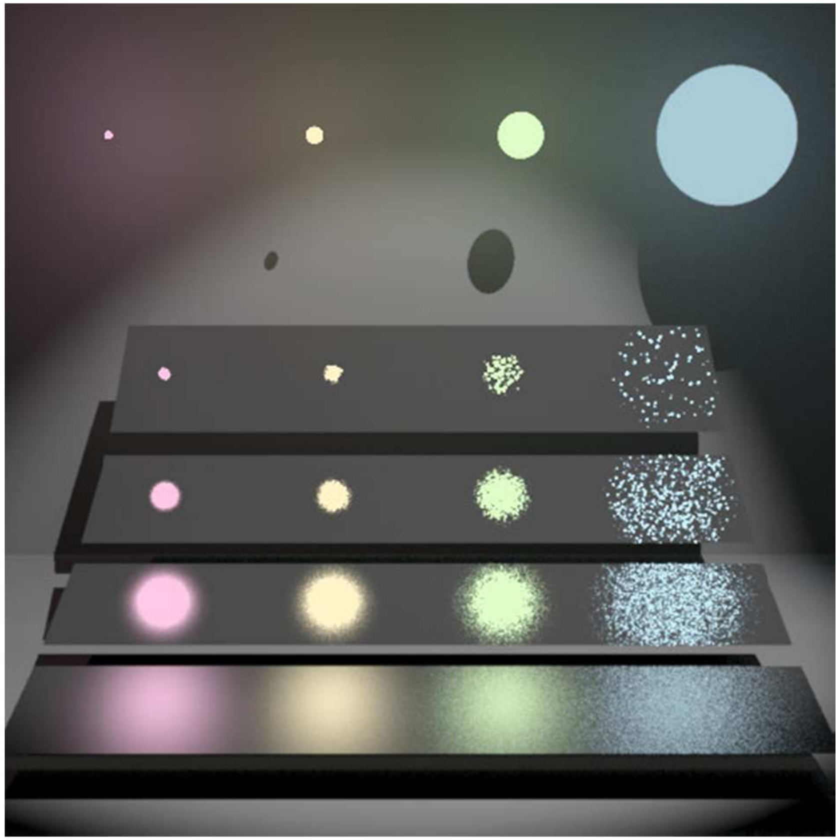

Light source area sampling

Images: Eric Veach

Direct illumination: BRDF sampling (rehash)

- **Integral** (integration over the hemisphere above **x**)

$$L_\mathrm{r}(\mathbf{x},\omega_\mathrm{o}) = \int_{H(\mathbf{x})} L_\mathrm{e}(\mathrm{r}(\mathbf{x},\omega_\mathrm{i}),-\omega_\mathrm{i})\cdot f_r(\mathbf{x},\omega_\mathrm{i}\to\omega_\mathrm{o})\cdot\cos\theta_\mathrm{i}\,\mathrm{d}\omega_\mathrm{i}$$

- **MC estimator**
  - Generate random direction $\omega_{\mathrm{i},k}$ from the pdf p
  - Cast a ray from the surface point **x** in the direction $\omega_{\mathrm{i},k}$
  - If it hits a light source, add $L_\mathrm{e}(.)\, f_r(.)$ cos/pdf

$$\hat{L}_\mathrm{r}(\mathbf{x},\omega_\mathrm{o}) = \frac{1}{N}\sum_{k=1}^{N}\frac{L_\mathrm{e}(\mathrm{r}(\mathbf{x},\omega_{\mathrm{i},k}),-\omega_{\mathrm{i},k})\cdot f_r(\mathbf{x},\omega_{\mathrm{i},k}\to\omega_\mathrm{o})\cdot\cos\theta_{\mathrm{i},k}}{p(\omega_{\mathrm{i},k})}$$

Direct illumination: Light source area sampling (rehash)

- **Integral** (integration over the light source area)

$$L_\mathrm{r}(\mathbf{x},\omega_\mathrm{o}) = \int_A L_\mathrm{e}(\mathbf{y}\to\mathbf{x})\cdot f_r(\mathbf{y}\to\mathbf{x}\to\omega_\mathrm{o})\cdot V(\mathbf{y}\leftrightarrow\mathbf{x})\cdot G(\mathbf{y}\leftrightarrow\mathbf{x})\,\mathrm{d}A_\mathbf{y}$$

- **MC estimator**
  - Generate a random position $\mathbf{y}_k$ on the source
  - Test the visibility V(x, y) between **x** and **y**
  - If V(x, y) = 1, add |A| $L_\mathrm{e}(\mathbf{y})\, f_r(.)$ cos/pdf

$$\hat{L}_\mathrm{r}(\mathbf{x},\omega_\mathrm{o}) = \frac{|A|}{N}\sum_{k=1}^{N} L_\mathrm{e}(\mathbf{y}_k\to\mathbf{x})\cdot f_r(\mathbf{y}_k\to\mathbf{x}\to\omega_\mathrm{o})\cdot V(\mathbf{y}_k\leftrightarrow\mathbf{x})\cdot G(\mathbf{y}_k\leftrightarrow\mathbf{x})$$

Direct illumination: Two strategies

- **BRDF proportional sampling**
  - Better for large light sources and/or highly glossy BRDFs
  - The probability of hitting a small light source is small -> high variance, noise

- **Light source area sampling**
  - Better for smaller light sources
  - It is the only possible strategy for point sources
  - For large sources, many samples are generated outside the BRDF lobe -> high variance, noise

Direct illumination: Two strategies

- Which strategy should we choose?
  - **Both!**

- Both strategies estimate the same quantity $L_r(\mathbf{x}, \omega_o)$
  - A mere sum would estimate 2 x $L_r(\mathbf{x}, \omega_o)$, which is wrong

- We need a weighted average of the techniques, but **how to choose the weights**? => MIS

How to choose the weights?

- **Multiple importance sampling** (Veach & Guibas, 95)

- Weights are functions of the pdf values

- Almost minimizes variance of the combined estimator

- Almost optimal solution

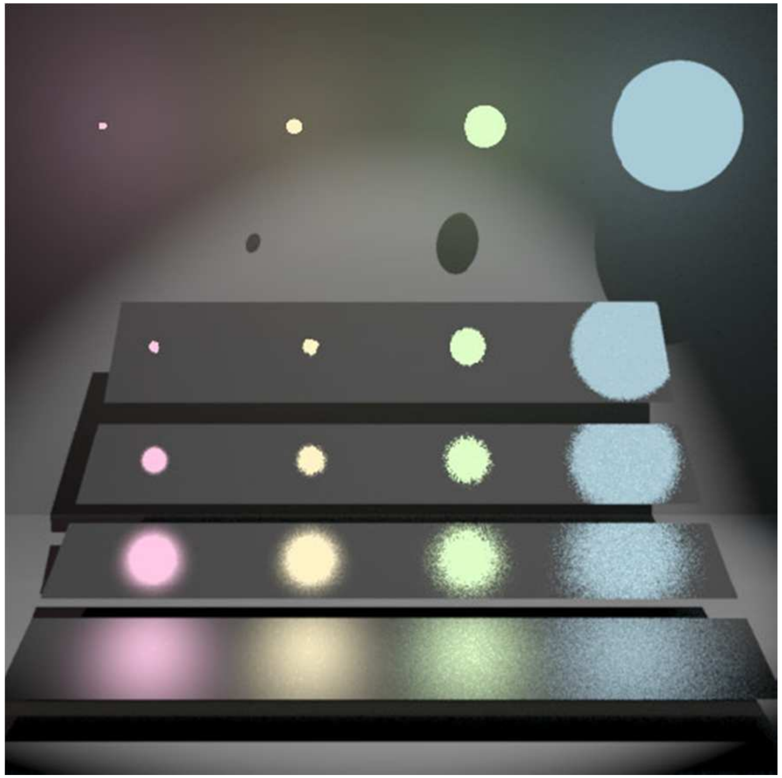

Image: Eric Veach

Direct illumination calculation using MIS

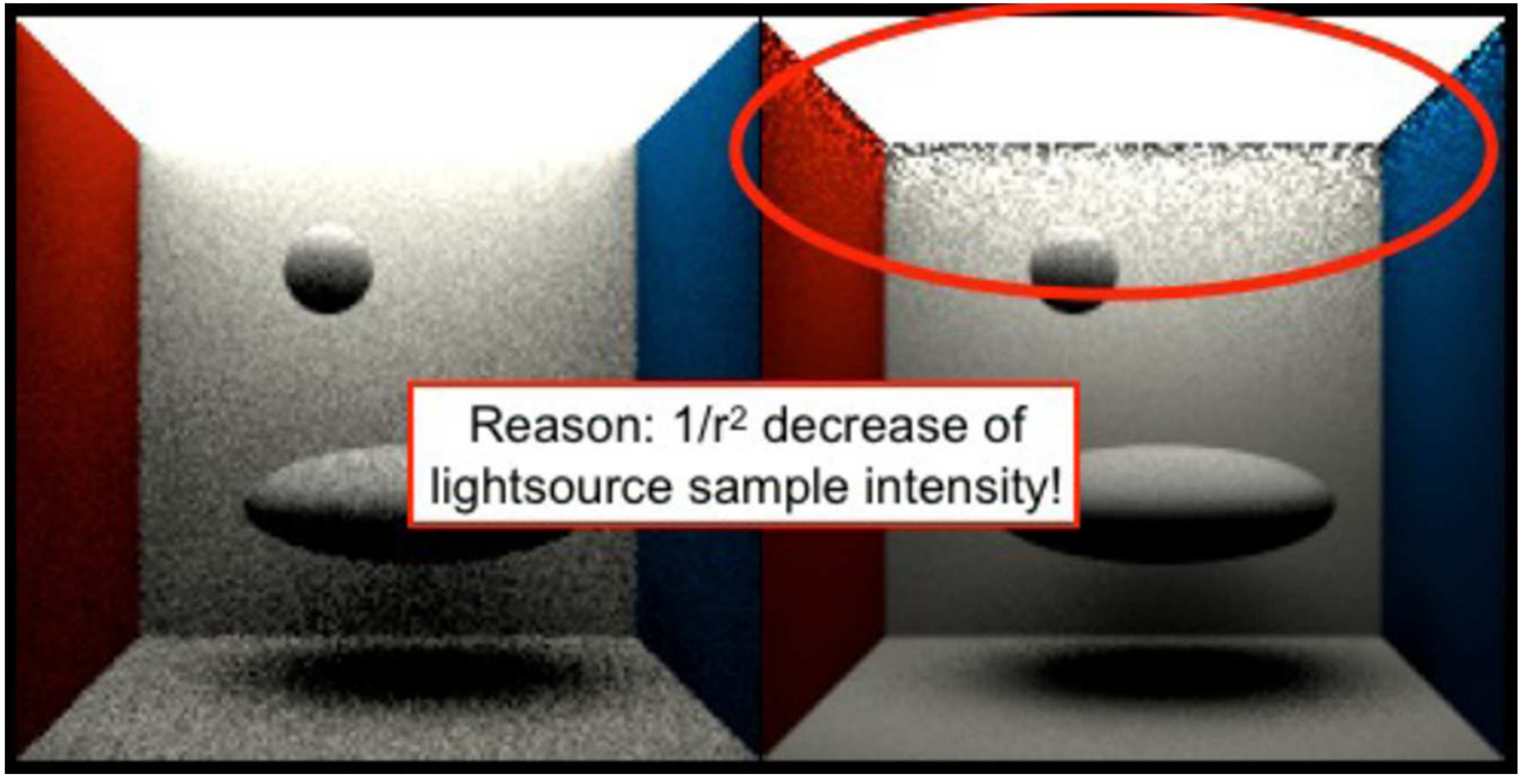

Sampling technique (pdf) p_1: BRDF sampling

Sampling technique (pdf) p_2: Light source area sampling

Image: Alexander Wilkie

Combination

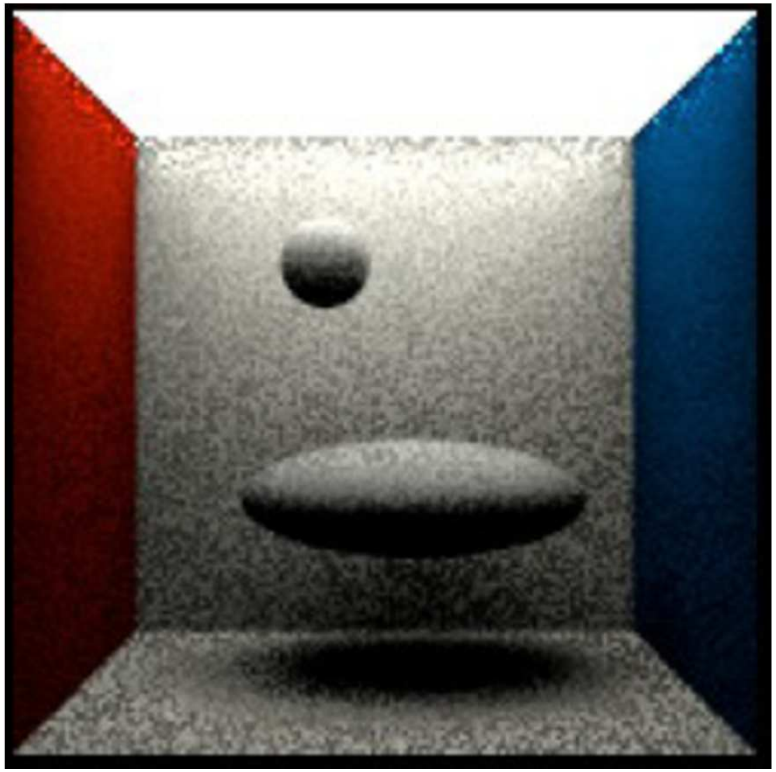

Arithmetic average
Preserves **bad** properties
of both techniques

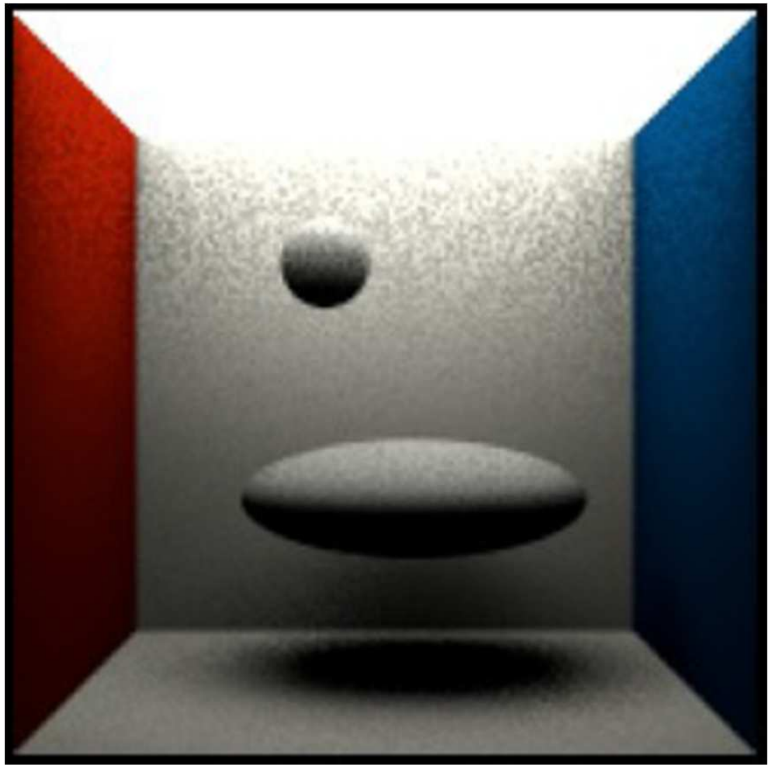

Balance heuristic
Bingo!!!

Image: Alexander Wilkie

MIS weight calculation

Sample weight for BRDF sampling

$$w_1(\omega_j) = \frac{p_1(\omega_j)}{p_1(\omega_j) + p_2(\omega_j)}$$

PDF for BRDF sampling

PDF with which the direction ω_j would have been generated, if we used light source area sampling

PDFs

- **BRDF sampling: $p_1(\omega)$**
  - Depends on the BRDF, e.g. for a Lambertian BRDF:

$$p_1(\omega) = \frac{\cos\theta_{\mathbf{x}}}{\pi}$$

- **Light source area sampling: $p_2(\omega)$**

$$p_2(\omega) = \frac{1}{|A|} \frac{\|\mathbf{x}-\mathbf{y}\|^2}{\underline{\cos\theta_{\mathbf{y}}}}$$

Conversion of the uniform pdf 1/|A| from the area measure (dA) to the solid angle measure (dω)

Contributions of the sampling techniques

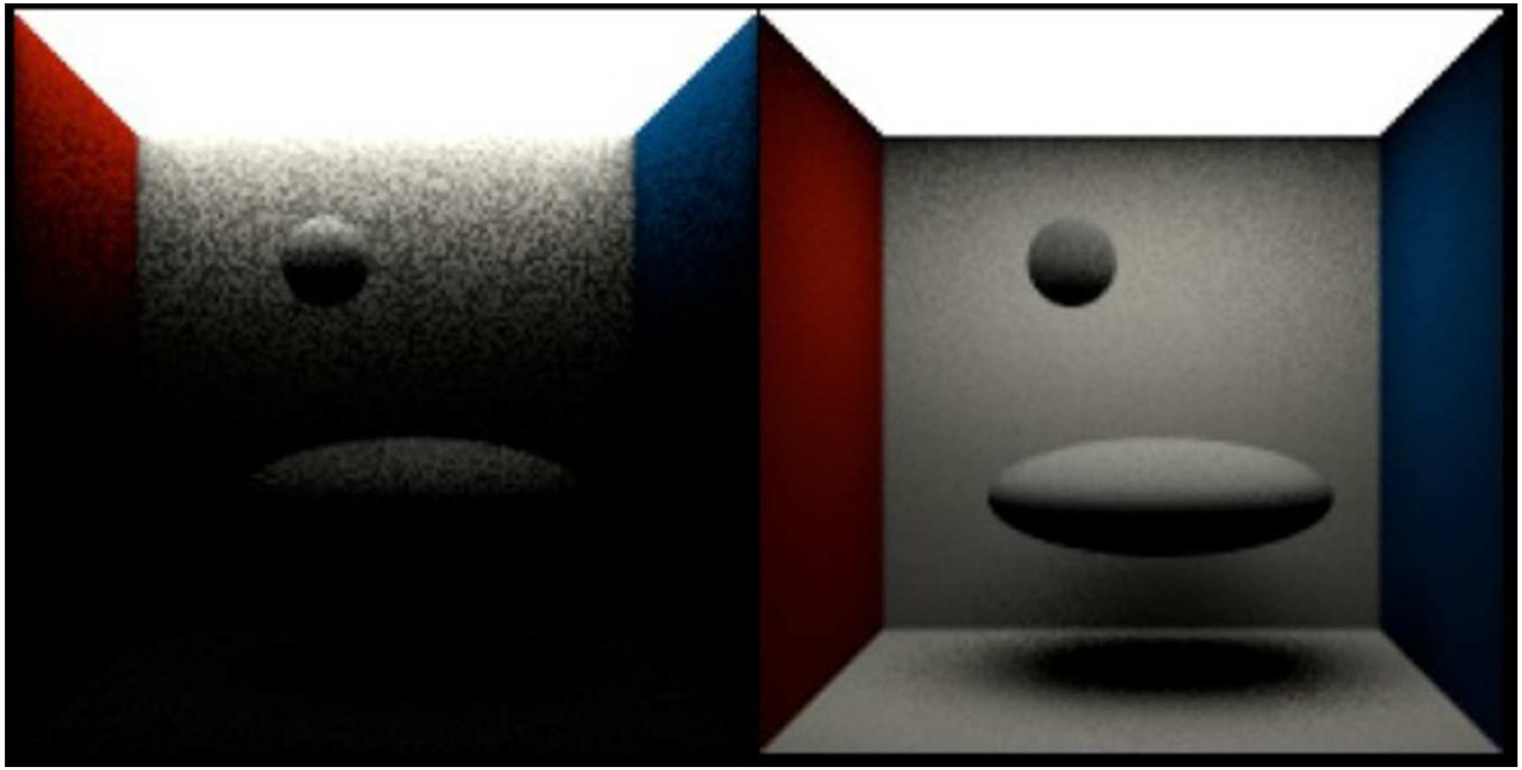

w1 * BRDF sampling **w2 * light source area sampling**

Image: Alexander Wilkie