Advanced 3D graphics for movies and games (NPGR010)

– Multiple Importance Sampling

Jiří Vorba, MFF UK/Weta Digital
jirka@cgg.mff.cuni.cz

Slides of prof. Jaroslav Křivánek, minor edits by Jiří Vorba

Sampling of environment lighting

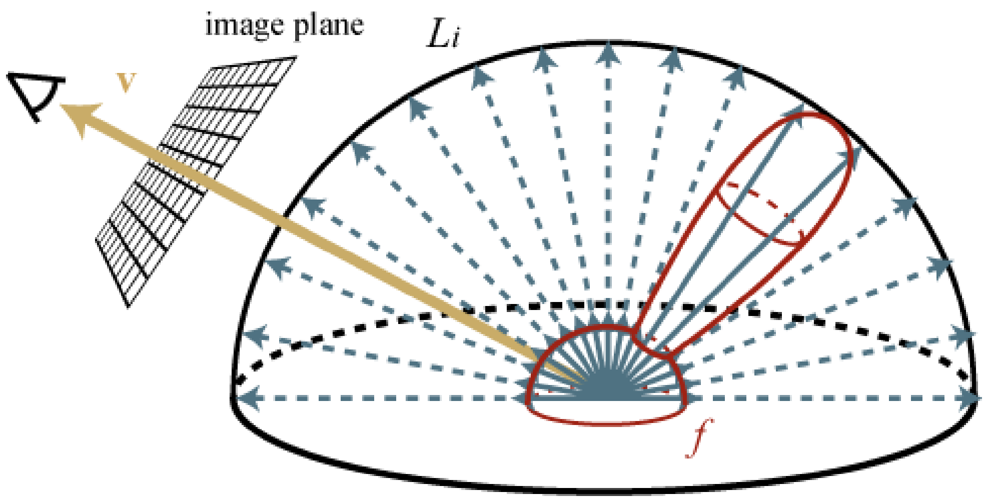

$$L_{\text{out}}(\omega_{\text{out}}) = \int_{H(\mathbf{x})} L_{\text{in}}(\omega_{\text{in}}) \cdot f_r(\omega_{\text{in}} \to \omega_{\text{out}}) \cdot \cos\theta_{\text{in}} \; \mathrm{d}\omega_{\text{in}}$$

Sampling of environment lighting

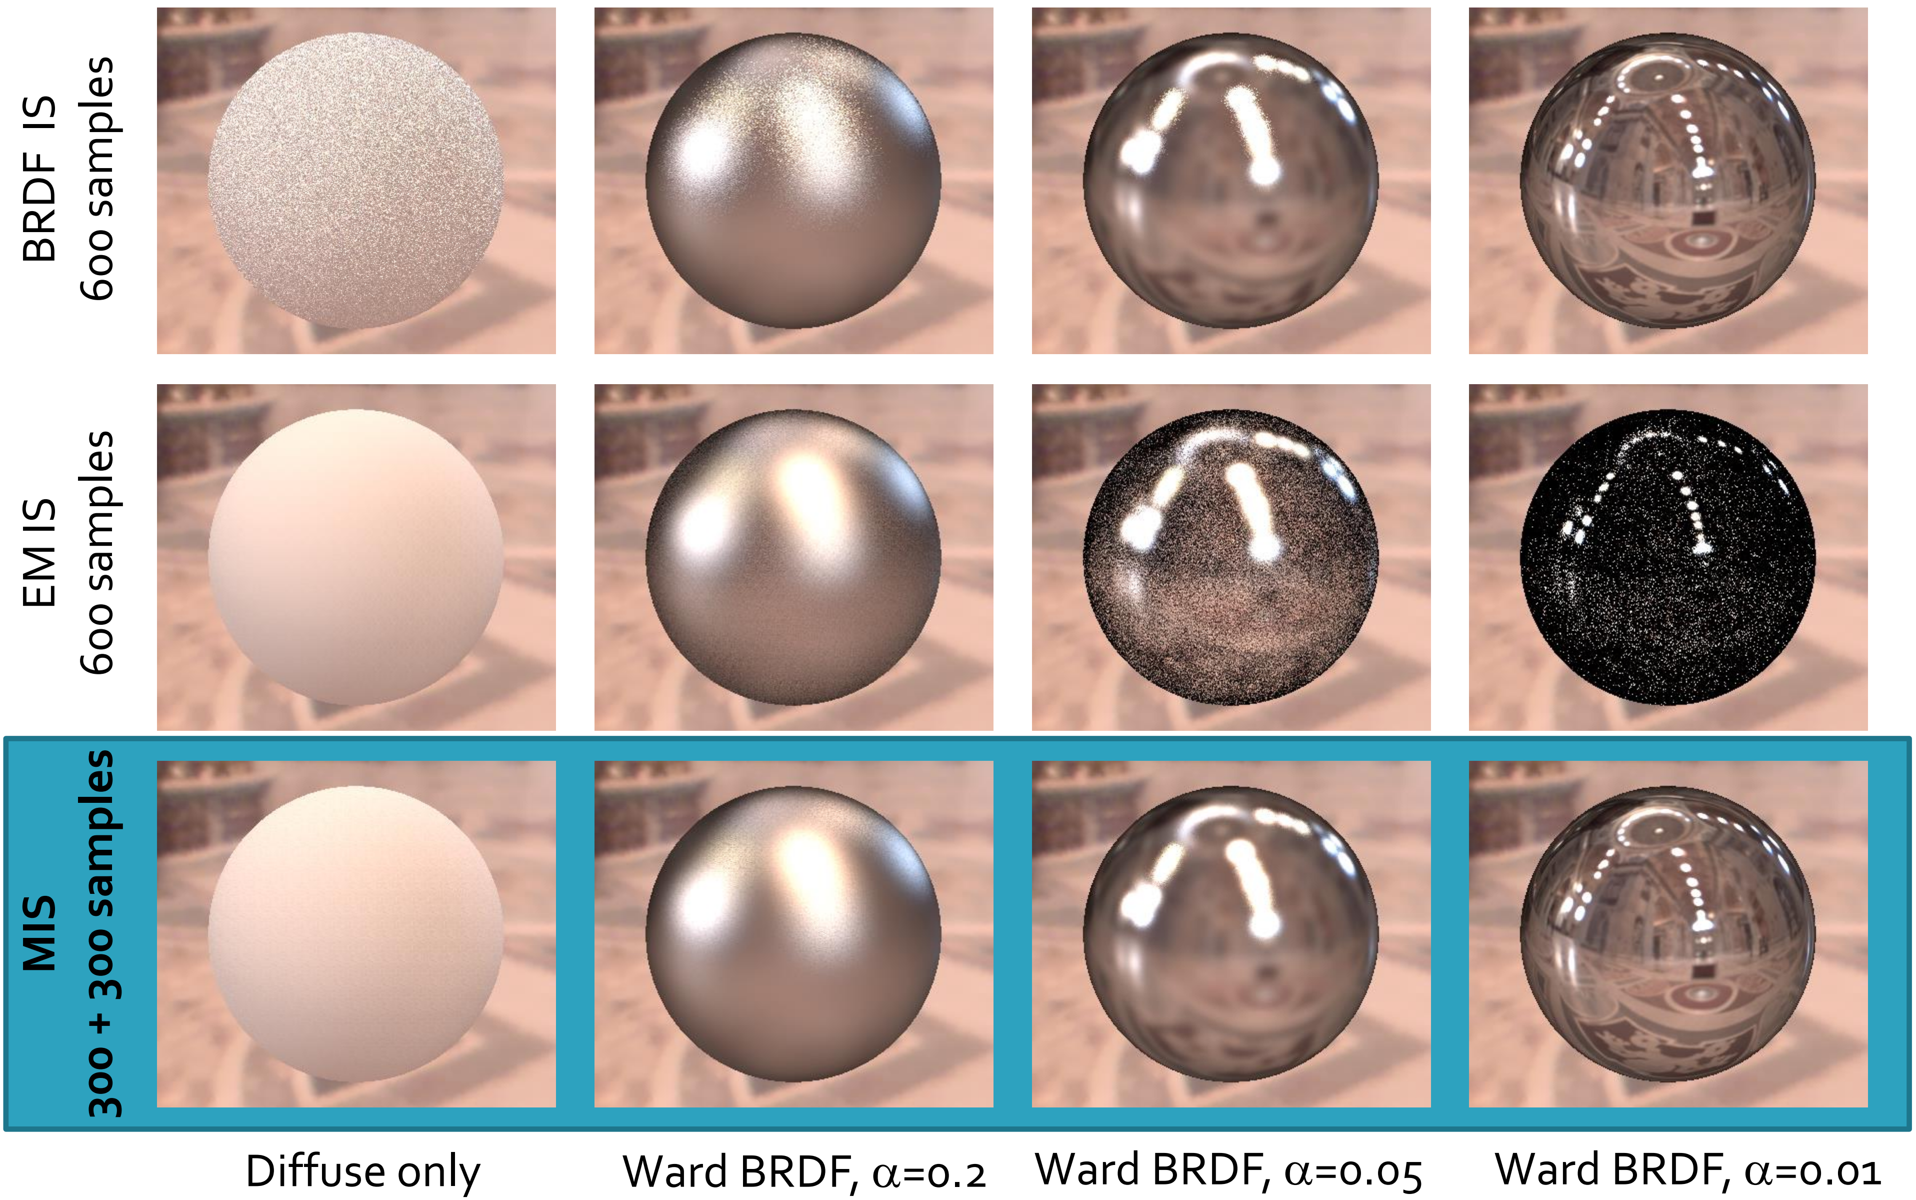

Sampling of environment lighting

- Two different sampling strategies for generating the incoming light direction ω_{in}

1. **BRDF-proportional sampling -** $p_a(\omega_{in})$

2. **Environment map-proportional sampling -** $p_b(\omega_{in})$

What is wrong with using either of the two strategies alone?

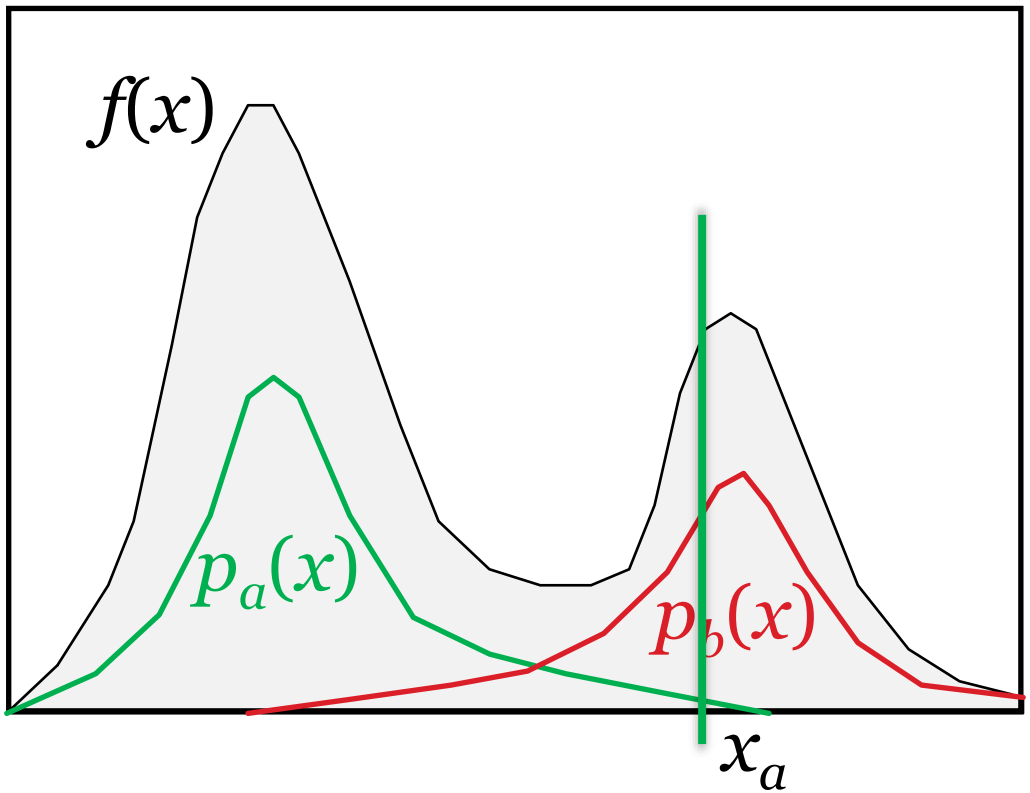

Better strategy

$$\frac{1}{2}p_a(x) + \frac{1}{2}p_b(x)$$

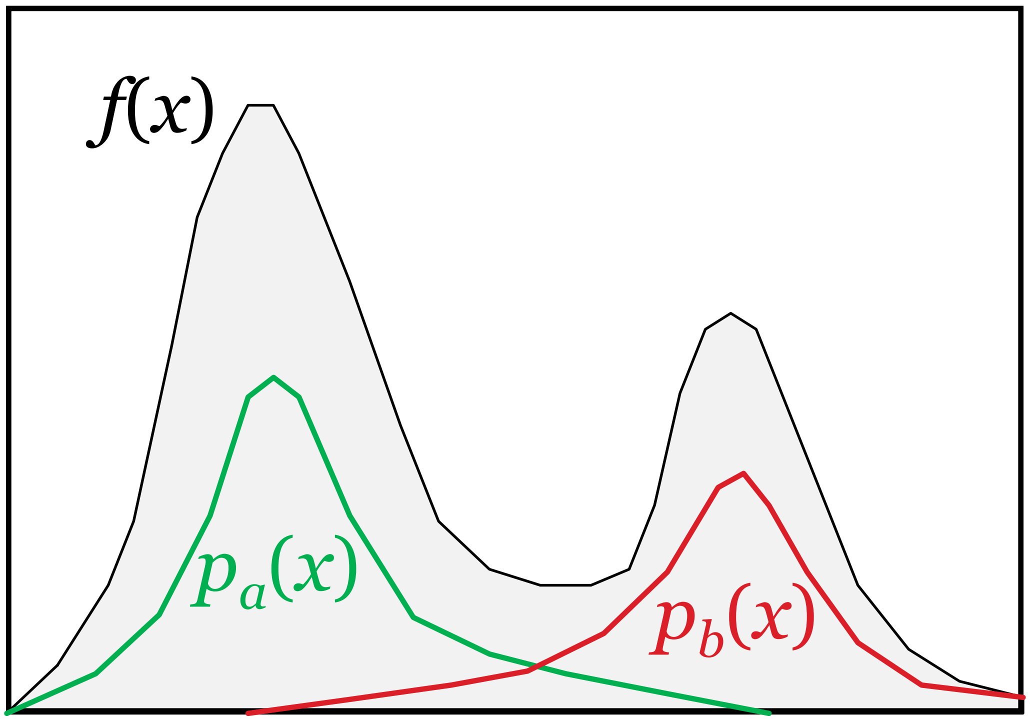

Example: Sum of two Gaussians

$\frac{e^{-(x+8)^2/(2\times 4^2)}}{\sqrt{2\pi}\times 4}$	$p_a(x)$
$\frac{e^{-(x-8)^2/(2\times 4^2)}}{\sqrt{2\pi}\times 4}$	$p_b(x)$
$\frac{1}{2}\times\frac{e^{-(x+8)^2/(2\times 4^2)}}{\sqrt{2\pi}\times 4}+\frac{1}{2}\times\frac{e^{-(x-8)^2/(2\times 4^2)}}{\sqrt{2\pi}\times 4}$	

$$\frac{1}{2}p_a(x)+\frac{1}{2}p_b(x)$$

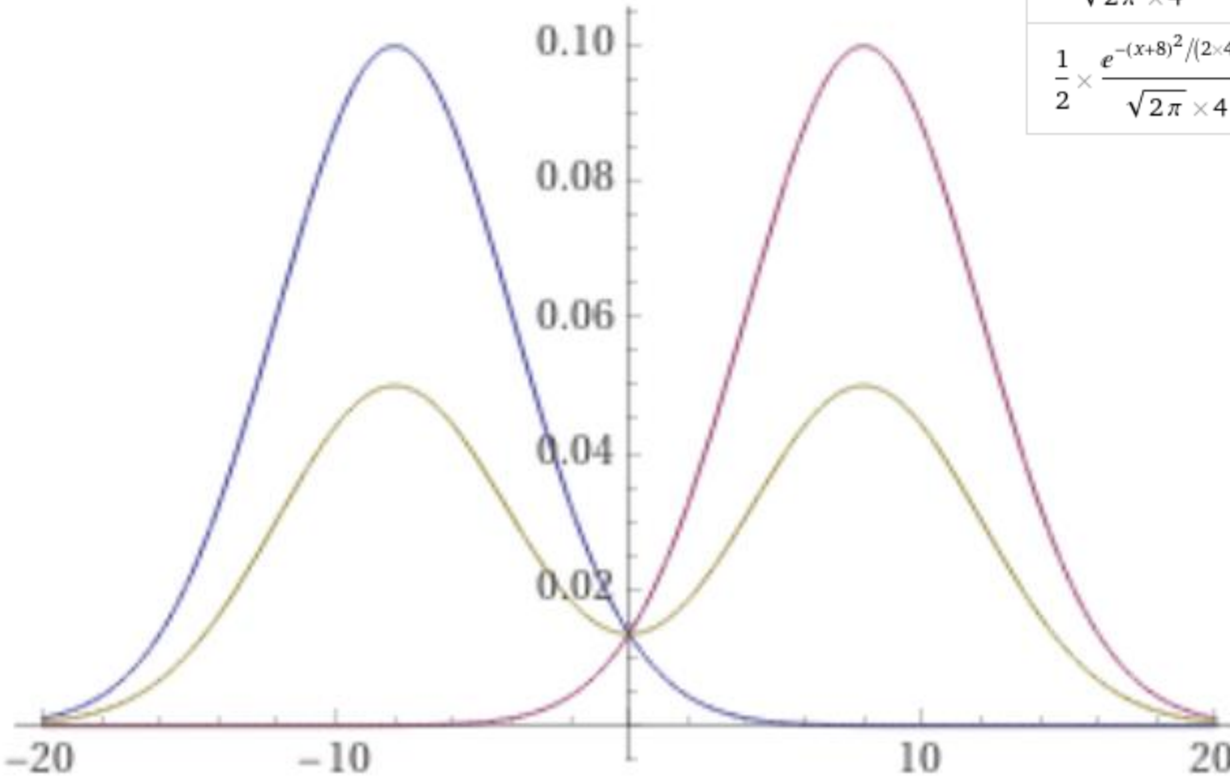

Notes on the previous slides

- We have a complex multimodal integrand g(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: sample from both techniques p_a and p_b , and then weight the samples appropriately.

- 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 g(x) might still be large. Without having p_b at our disposal, the MC estimator would be dividing the large g(x) by the small p_a (x), producing an outlier sample.
- 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 g(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.

Multiple Importance Sampling

Eric Veach [1997]

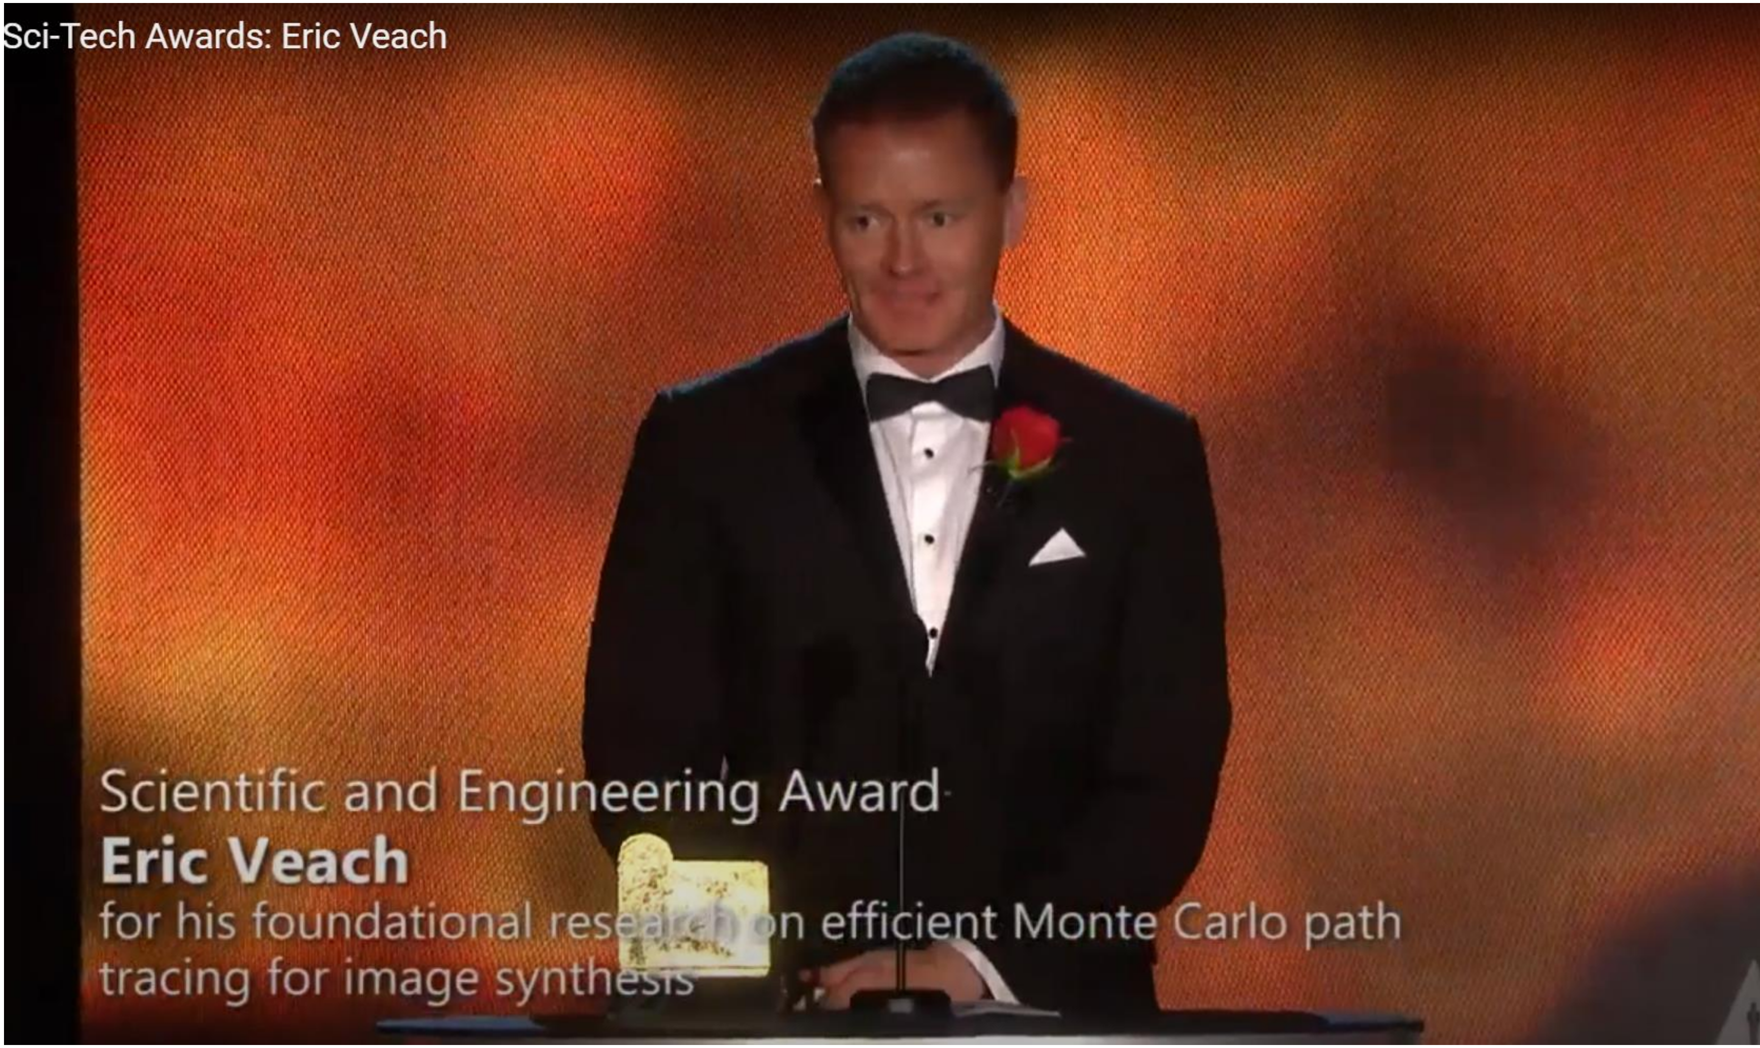

Multiple Importance Sampling

- 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)

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

sampling techniques

samples from individual techniques

Unbiasedness of the combined estimator

- The MIS estimator is unbiased…

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

- … provided the weighting functions sum up to 1

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

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

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

MIS estimator with the Balance heuristic

- Plugging Balance heuristic weights into the MIS formula

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

  - The contribution of a sample does not depend on which technique (pdf) it came from
  - Effectively, the sample is drawn from a weighted average of the individual pdfs – as can be seen from the form of the estimator

Balance heuristic

- The balance heuristic **is almost optimal** [Veach 97]
  - No other weighting has variance much lower than the balance heuristic

- Our work
  [Kondapaneni et al. 2018]
  revises MIS
  - If you allow negative weights,
    one can improve
    over the balance heuristic a lot

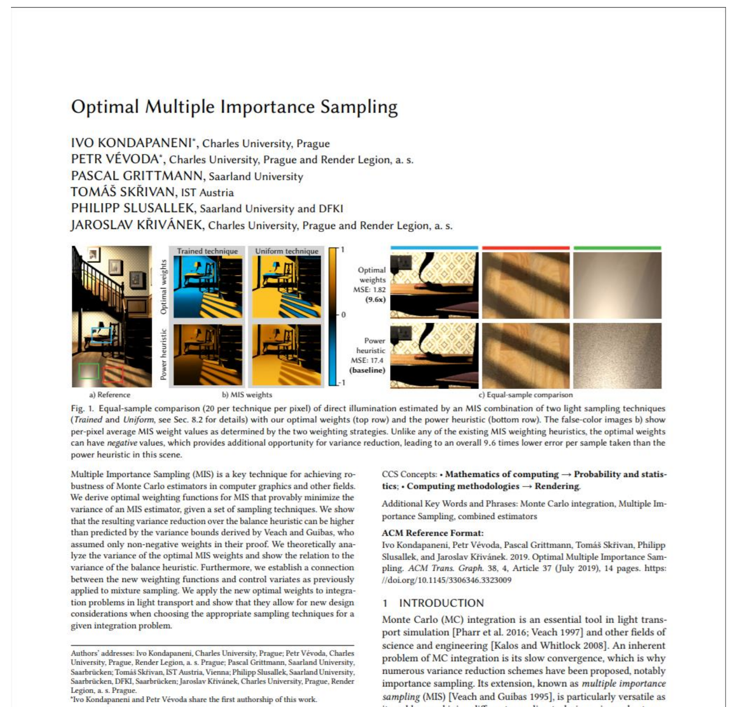

Optimal Multiple Importance Sampling

IVO KONDAPANENI*, Charles University, Prague
PETR VÉVODA*, Charles University, Prague and Render Legion, a. s.
PASCAL GRITTMANN, Saarland University
TOMÁŠ SKŘIVAN, IST Austria
PHILIPP SLUSALLEK, Saarland University and DFKI
JAROSLAV KŘIVÁNEK, Charles University, Prague and Render Legion, a. s.

Fig. 1. Equal-sample comparison (20 per technique per pixel) of direct illumination estimated by an MIS combination of two light sampling techniques (*Trained* and *Uniform*, see Sec. 8.2 for details) with our optimal weights (top row) and the power heuristic (bottom row). The false-color images b) show per-pixel average MIS weight values as determined by the two weighting strategies. Unlike any of the existing MIS weighting heuristics, the optimal weights can have *negative* values, which provides additional opportunity for variance reduction, leading to an overall 9.6 times lower error per sample taken than the power heuristic in this scene.

Multiple Importance Sampling (MIS) is a key technique for achieving robustness of Monte Carlo estimators in computer graphics and other fields. We derive optimal weighting functions for MIS that provably minimize the variance of an MIS estimator, given a set of sampling techniques. We show that the resulting variance reduction over the balance heuristic can be higher than predicted by the variance bounds derived by Veach and Guibas, who assumed only non-negative weights in their proof. We theoretically analyze the variance of the optimal MIS weights and show the relation to the variance of the balance heuristic. Furthermore, we establish a connection between the new weighting functions and control variates as previously applied to mixture sampling. We apply the new optimal weights to integration problems in light transport and show that they allow for new design considerations when choosing the appropriate sampling techniques for a given integration problem.

Authors' addresses: Ivo Kondapaneni, Charles University, Prague; Petr Vévoda, Charles University, Prague, Render Legion, a. s. Prague; Pascal Grittmann, Saarland University, Saarbrücken; Tomáš Skřivan, IST Austria, Vienna; Philipp Slusallek, Saarland University, Saarbrücken, DFKI, Saarbrücken; Jaroslav Křivánek, Charles University, Prague, Render Legion, a. s. Prague.
*Ivo Kondapaneni and Petr Vévoda share the first authorship of this work.

CCS Concepts: • **Mathematics of computing → Probability and statistics**; • **Computing methodologies → Rendering**.

Additional Key Words and Phrases: Monte Carlo integration, Multiple Importance Sampling, combined estimators

ACM Reference Format:
Ivo Kondapaneni, Petr Vévoda, Pascal Grittmann, Tomáš Skřivan, Philipp Slusallek, and Jaroslav Křivánek. 2019. Optimal Multiple Importance Sampling. *ACM Trans. Graph.* 38, 4, Article 37 (July 2019), 14 pages. https://doi.org/10.1145/3306346.3323009

1 INTRODUCTION

Monte Carlo (MC) integration is an essential tool in light transport simulation [Pharr et al. 2016; Veach 1997] and other fields of science and engineering [Kalos and Whitlock 2008]. An inherent problem of MC integration is its slow convergence, which is why numerous variance reduction schemes have been proposed, notably importance sampling. Its extension, known as *multiple importance sampling* (MIS) [Veach and Guibas 1995], is particularly versatile as it enables combining different sampling techniques in a robust way

MIS for direct illumination from enviro lights

Application of MIS to environment light sampling

- Recall: Two sampling strategies for generating the incident direction ω_i

1. **BRDF-proportional sampling -** $p_a(\omega_{in})$
2. **Environment map-proportional sampling -** $p_b(\omega_{in})$

- Plug formulas for $p_a(\omega_{in})$ and $p_b(\omega_{in})$ into the general MIS formulas above

Direct illumination: Two strategies

- Which strategy should we choose?
  - **Both!**
- Both strategies estimate the same quantity $L_{\text{out}}(\mathbf{x}, \omega_{\text{out}})$
  - A mere sum would estimate $2 \times L_{\text{out}}(\mathbf{x}, \omega_{\text{out}})$, which is wrong
- We need a weighted average of the techniques, but **how to choose the weights**? → MIS

MIS weight calculation

MIS weight for a sample direction generated by BRDF lobe sampling

$$w_a(\omega_{\text{in},j}) = \frac{p_a(\omega_{\text{in},j})}{p_a(\omega_{\text{in},j}) + p_b(\omega_{\text{in},j})}$$

PDF for BRDF sampling

PDF with which the direction $\omega_{\text{in},j}$ would have been generated, if we used env map sampling

- Here, we assume one sample from each of the two strategies

MIS for enviro sampling – Algorithm

```
Vec3 omegaInA = generateBrdfSample();
float pdfA = evalBrdfPdf(omegaInA);
float pdfAsIfFromB = evalEnvMapPdf(omegaInA);
float misWeightA = pdfA / (pdfA + pdfAsIfFromB);
Rgb outRadianceEstimate = misWeightA *
        incRadiance(omegaInA) *
        brdf(omegaOut, omegaInA) *
        max(0, dot(omegaInA, surfNormal) / pdfA;

Vec3 omegaInB = generateEnvMapSample();
float pdfB = evalEnvMapPdf(omegaInB);
float pdfAsIfFromA = evalBrdfPdf(omegaInB);
float misWeightB = pdfB / (pdfB + pdfAsIfFromA);
outRadianceEstimate += misWeightB *
        incRadiance(omegaInB) *
        brdf(omegaOut, omegaInB) *
        max(0, dot(omegaInB, surfNormal) / pdfB;
```

MIS applied to enviro sampling

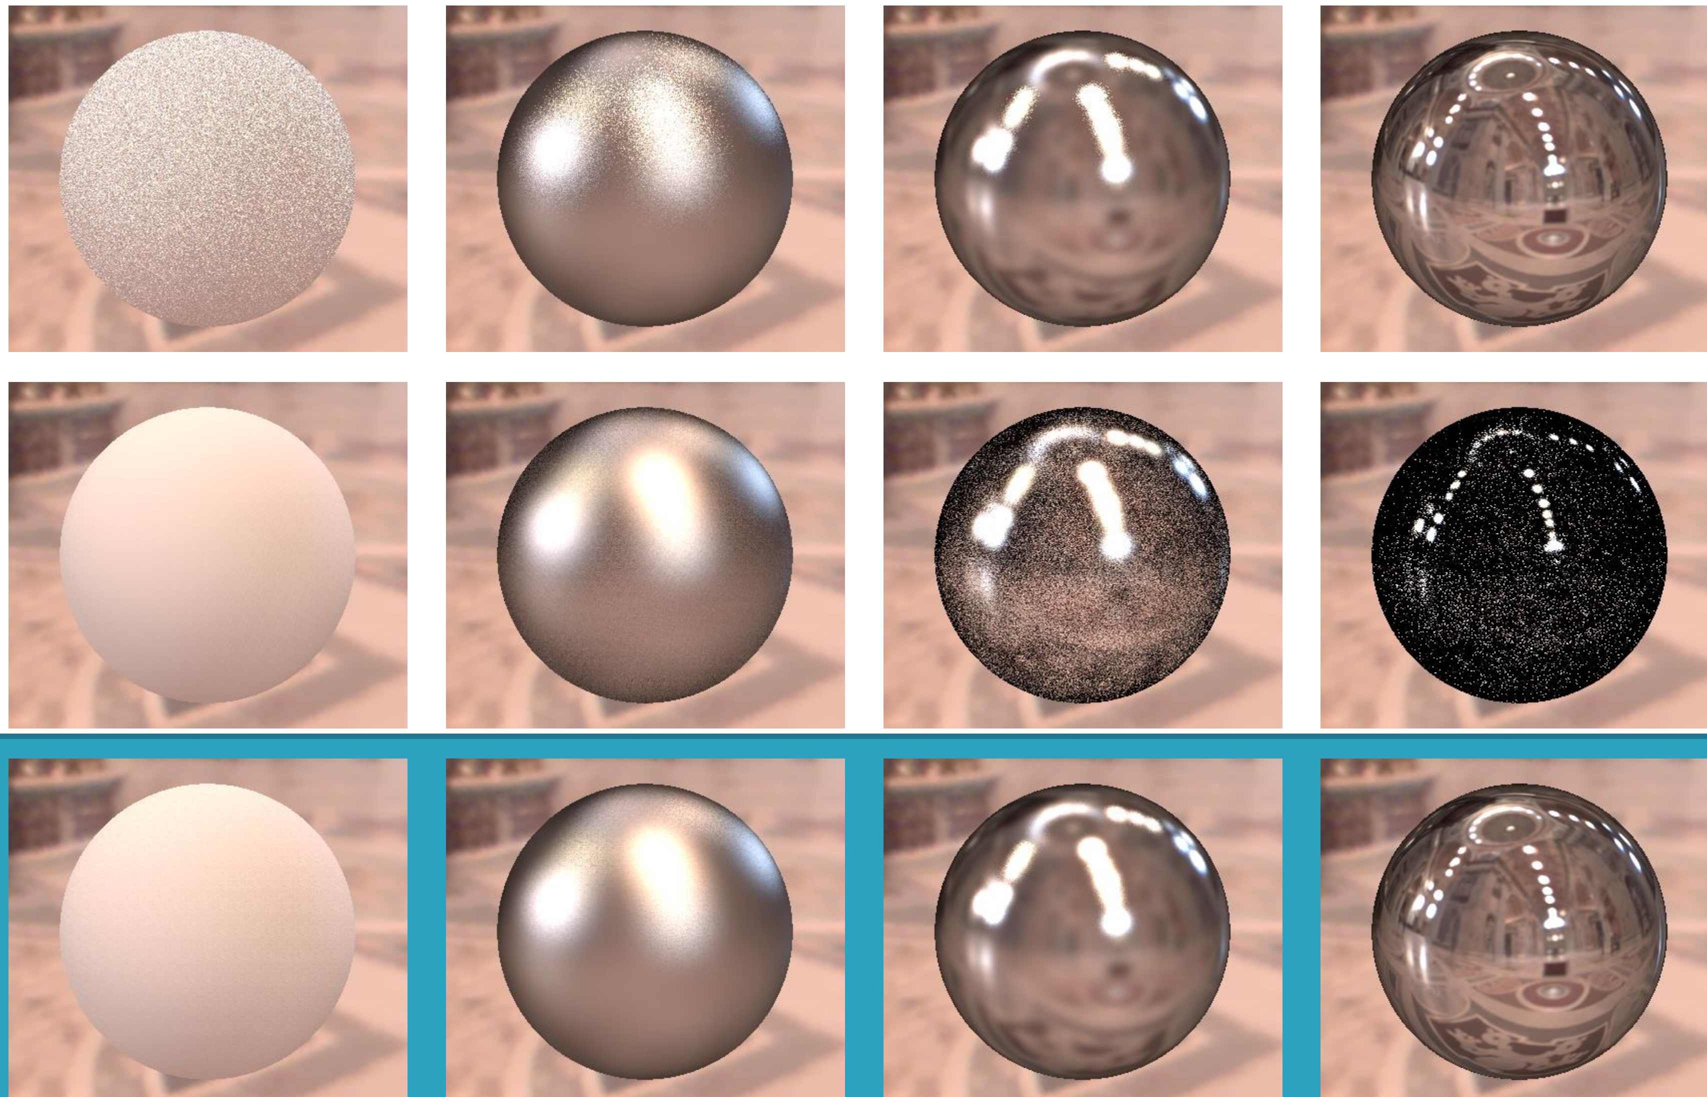

BRDF IS 600 samples

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

MIS for direct illumination from area lights

Area light sampling – Motivation

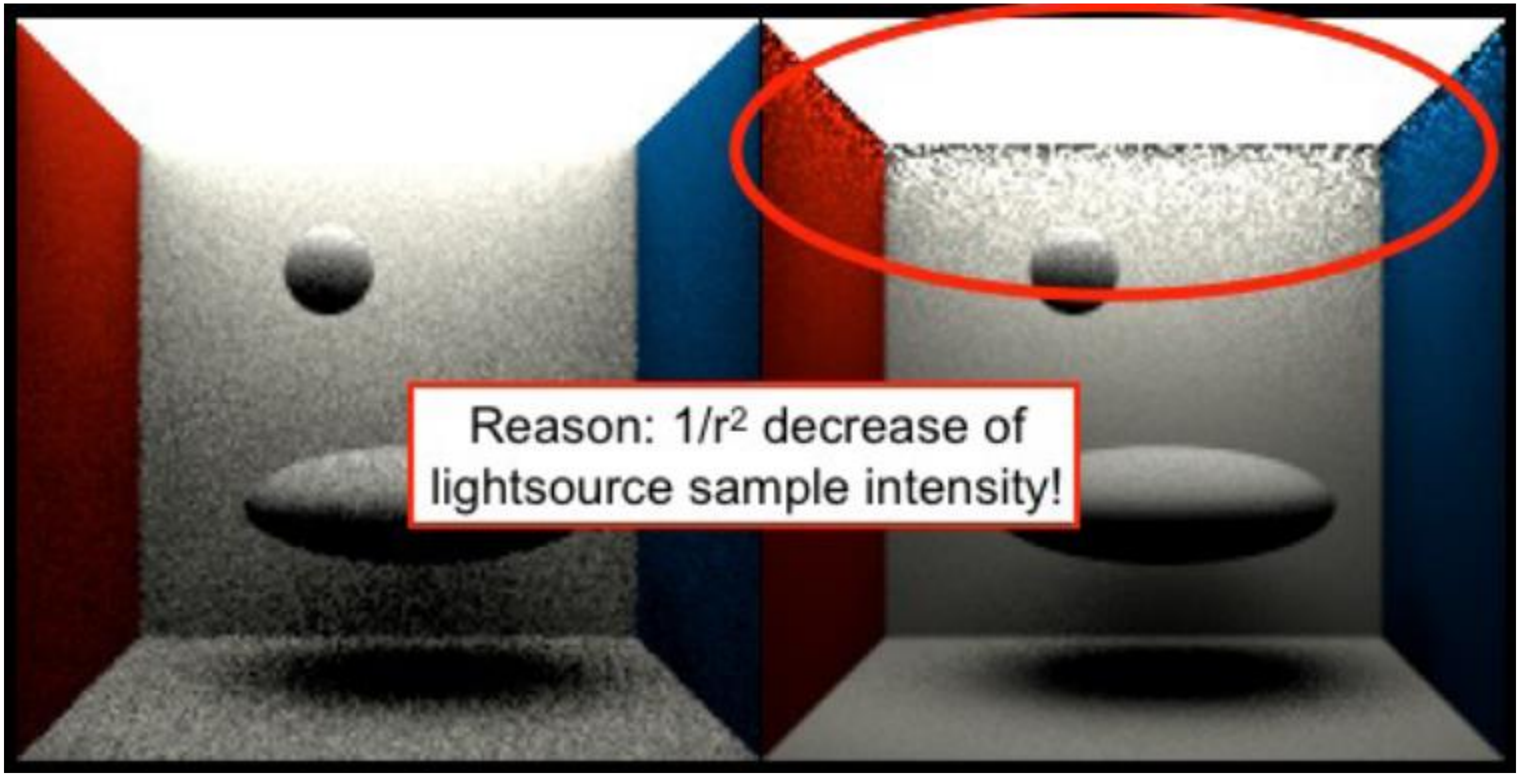

Image: Alexander Wilkie

Sampling technique (pdf) p_a:
BRDF sampling

Sampling technique (pdf) p_b:
Light source area sampling

Recall: Irradiance estimate and G term

- Reformulate the reflection integral (change of variables)

$$E(\mathbf{x}) = \int_{H(\mathbf{x})} L_{\mathrm{i}}(\mathbf{x}, \omega_{\mathrm{i}}) \cdot \cos\theta_{\mathrm{i}} \, \mathrm{d}\omega_{\mathrm{i}}$$

$G(\mathbf{y} \leftrightarrow \mathbf{x})$

$$= \int_{A} L_{\mathrm{e}}(\mathbf{y} \to \mathbf{x}) \cdot V(\mathbf{y} \leftrightarrow \mathbf{x}) \cdot \frac{\cos\theta_{\mathbf{y}} \cdot \cos\theta_{\mathbf{x}}}{\|\mathbf{y} - \mathbf{x}\|^2} \, \mathrm{d}A$$

- PDF for uniform sampling of the surface area:

$$p(\mathbf{y}) = \frac{1}{|A|}$$

MIS-based combination

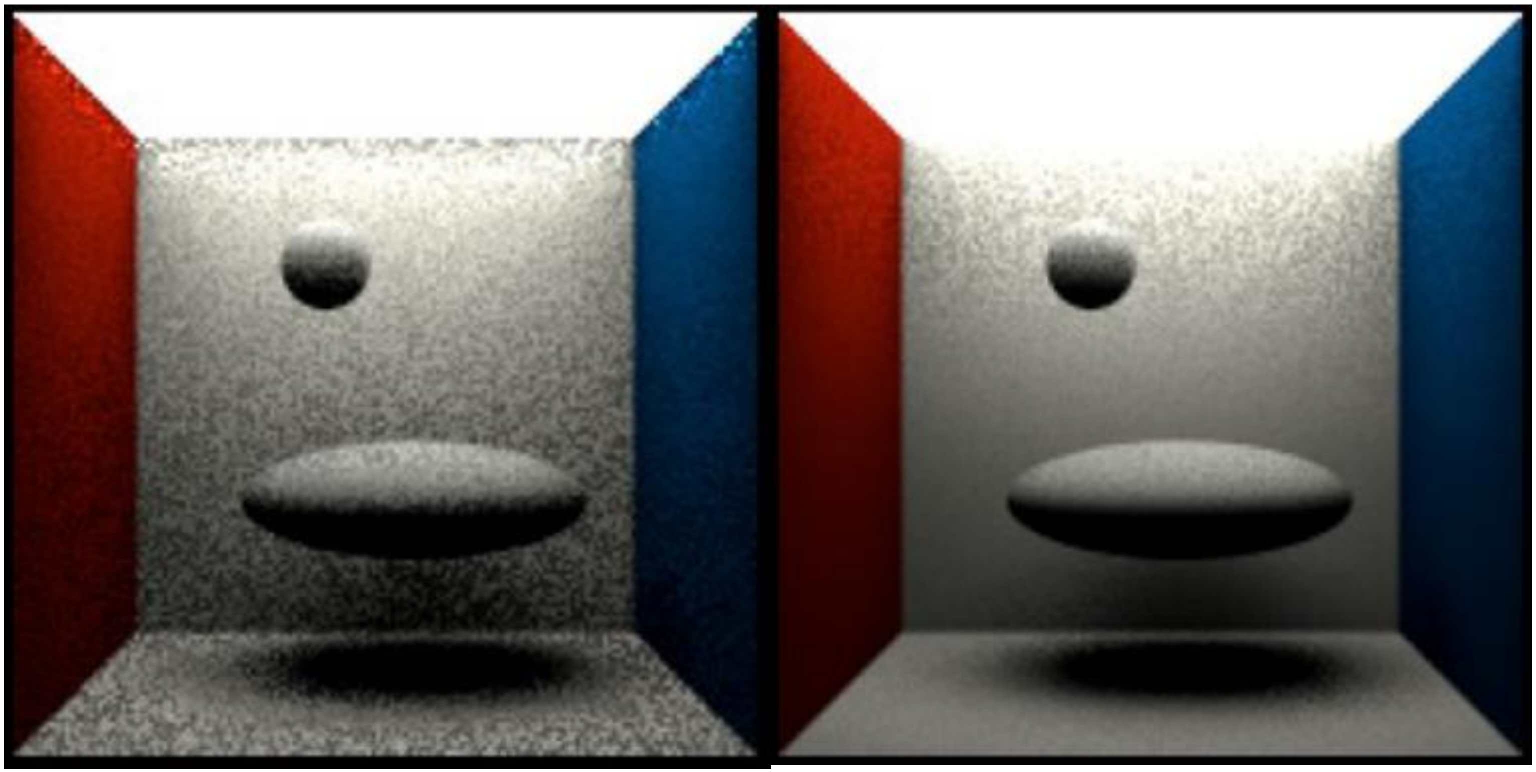

Image: Alexander Wilkie

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

MIS w/ the balance heuristic
Bingo!!!

Area light sampling – Classic Veach's example

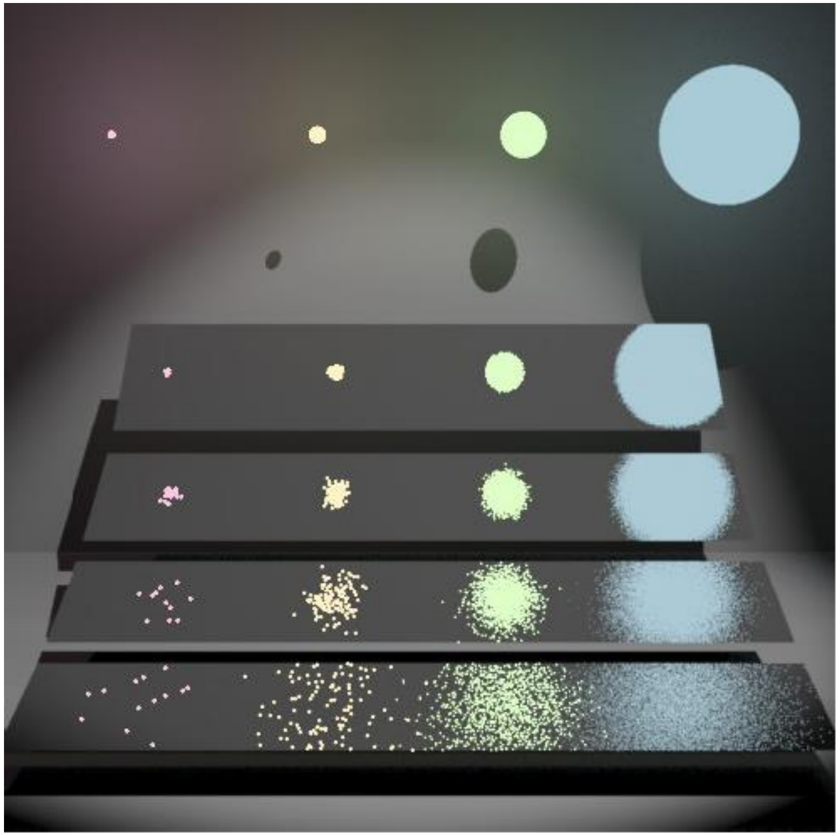

BRDF proportional sampling

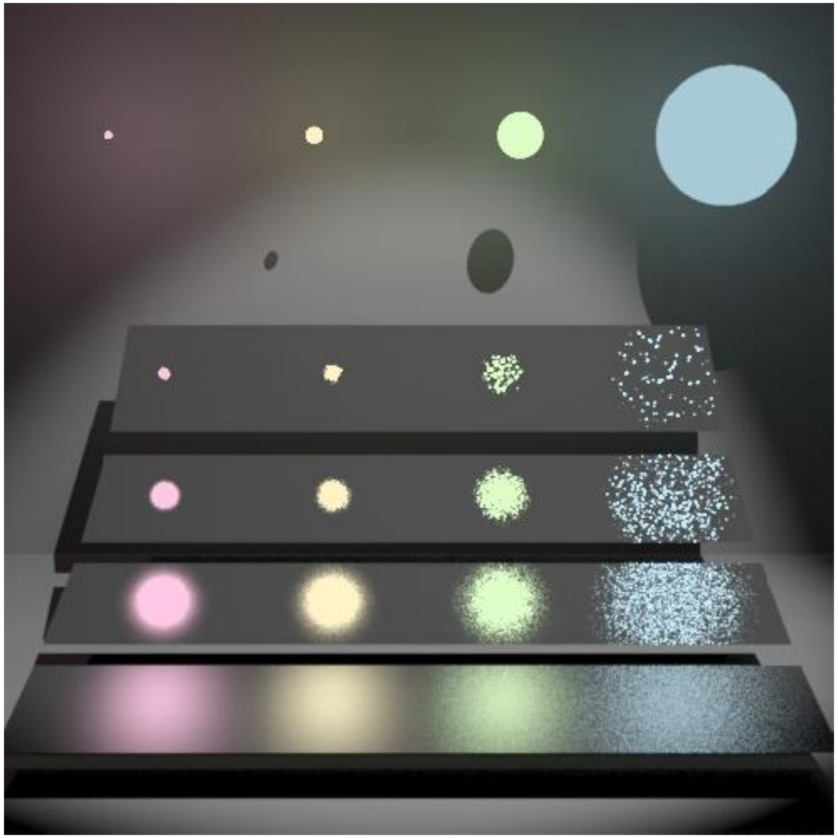

Light source area sampling

Images: Eric Veach

MIS-based combination

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

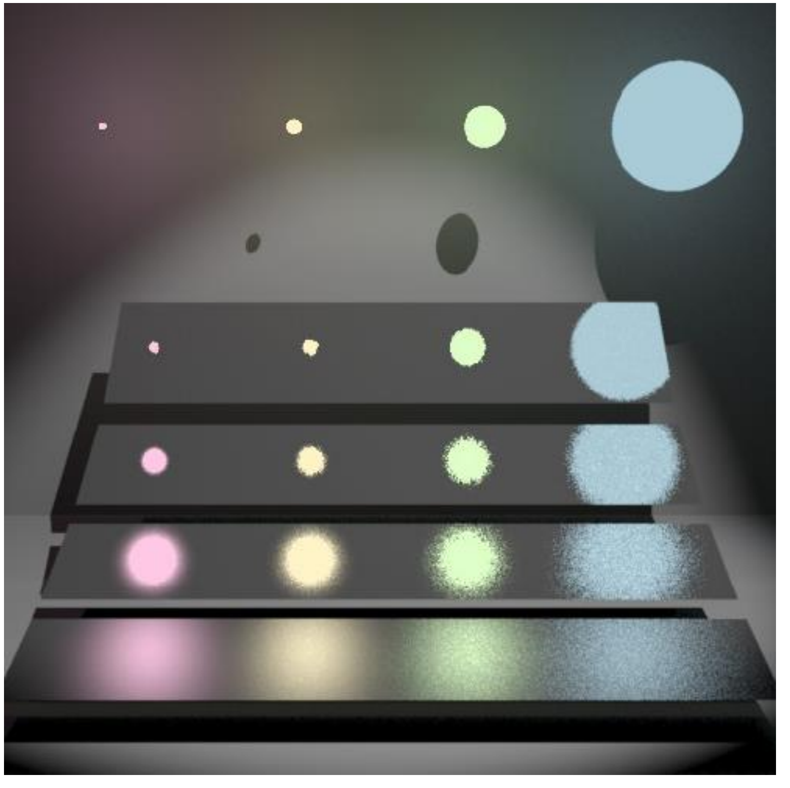

Image: Eric Veach

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

Example PDFs

- **BRDF sampling: $p_a(\omega)$**
  - Depends on the BRDF, e.g. the formulas for physically-based Phong BRDF from the last lecture

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

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

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

Contributions of the sampling techniques

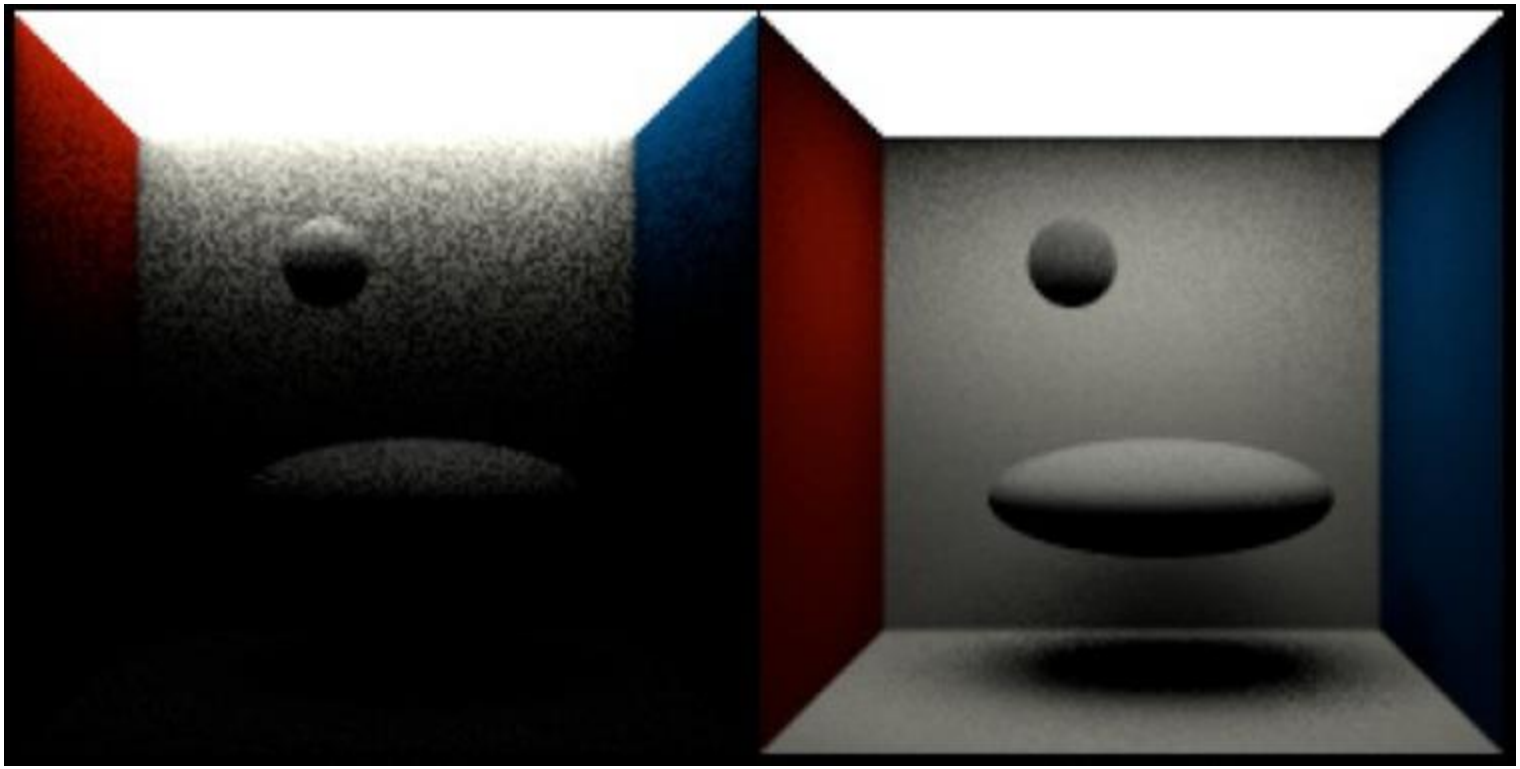

Image: Alexander Wilkie

w_a * **BRDF sampling** w_b * **light source area sampling**

Alternative MIS heuristics

Alternative combination heuristics

- **“Low variance problems”**
- Whenever one sampling technique yields a very low variance estimator, balance heuristic can be suboptimal

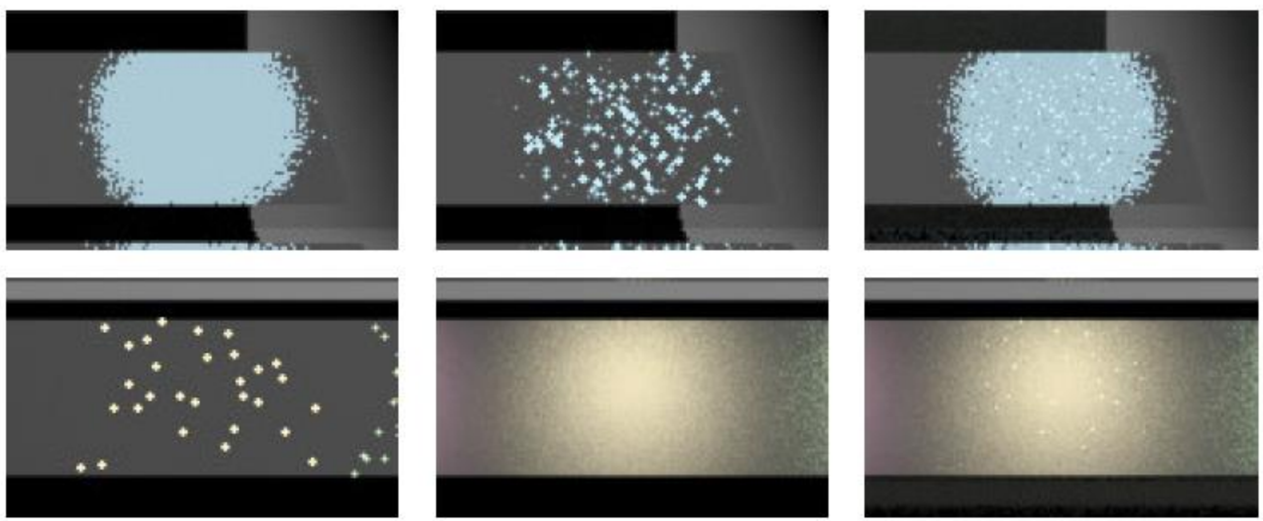

(a) Sampling the BSDF (b) Sampling the lights (c) The balance heuristic

Alternative combination heuristics

- **"Low variance problems"**
- Whenever one sampling technique yields a very low variance estimator, balance heuristic can be suboptimal
- "Power heuristic" or other heuristics can be better in such a case – see next slide

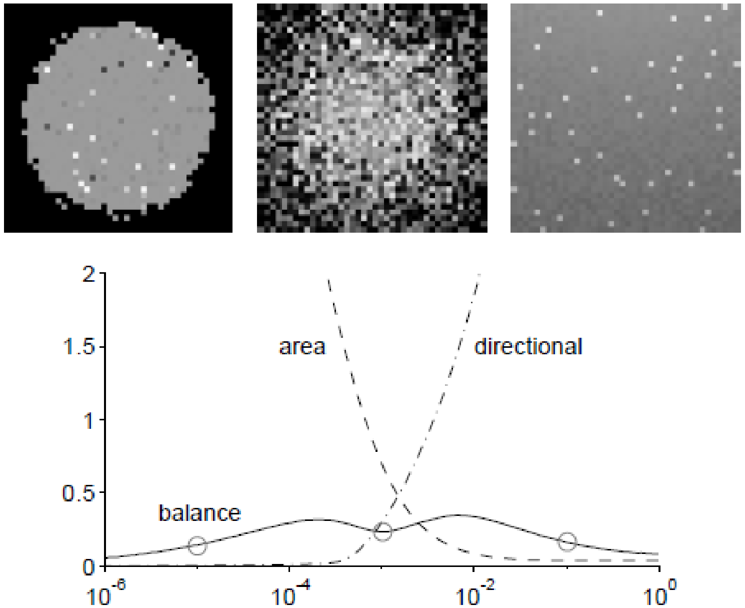

(a) The balance heuristic.

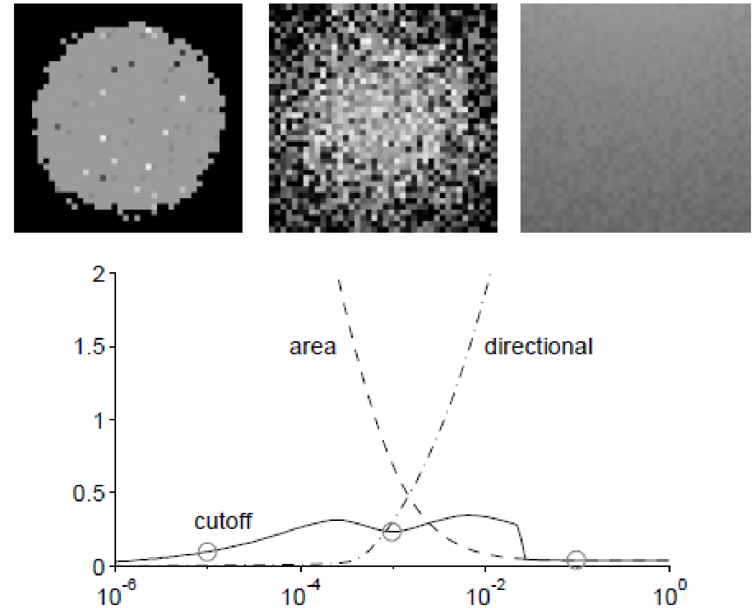

(b) The cutoff heuristic ($\alpha = 0.1$).

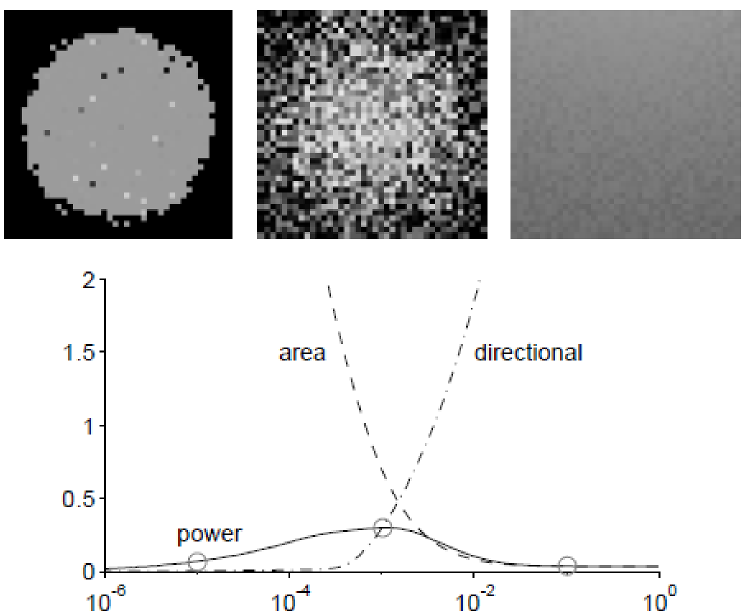

(c) The power heuristic ($\beta = 2$).

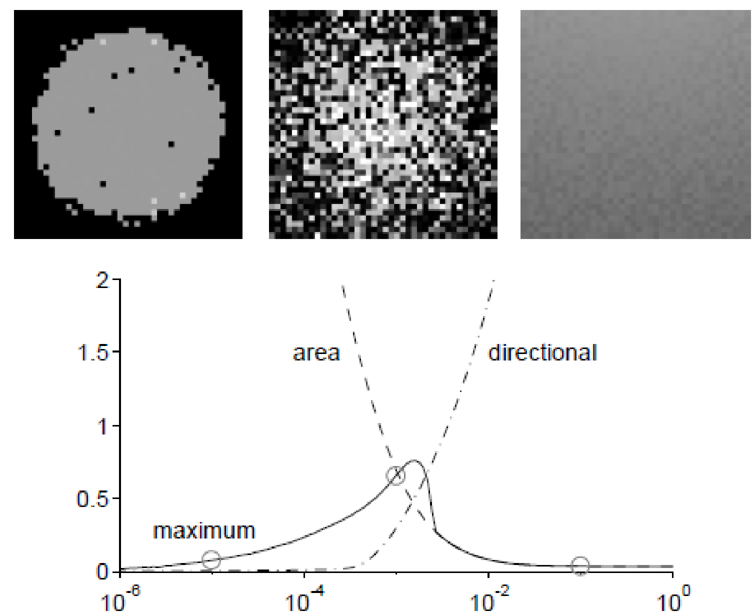

(d) The maximum heuristic.

Other examples of MIS applications

In the following we apply MIS to combine full path sampling techniques for calculating light transport in participating media.

Full transport

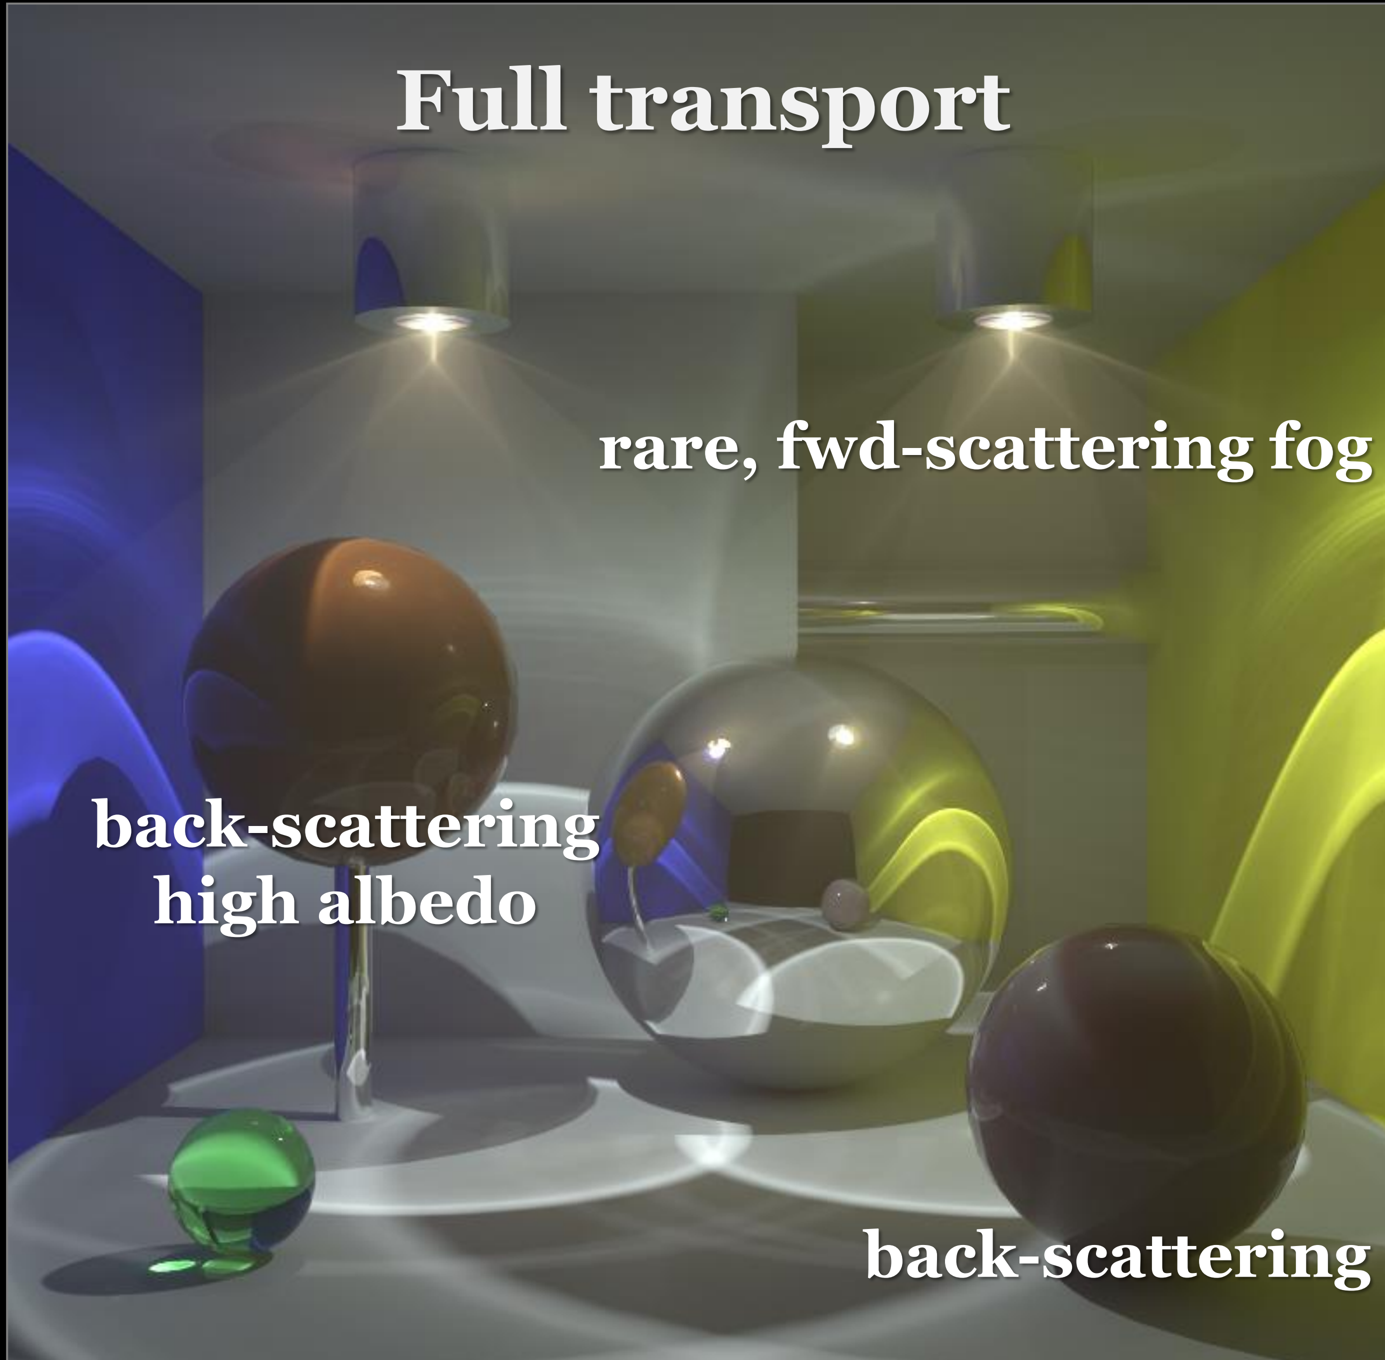

Medium transport only

Previous work comparison, 1 hr

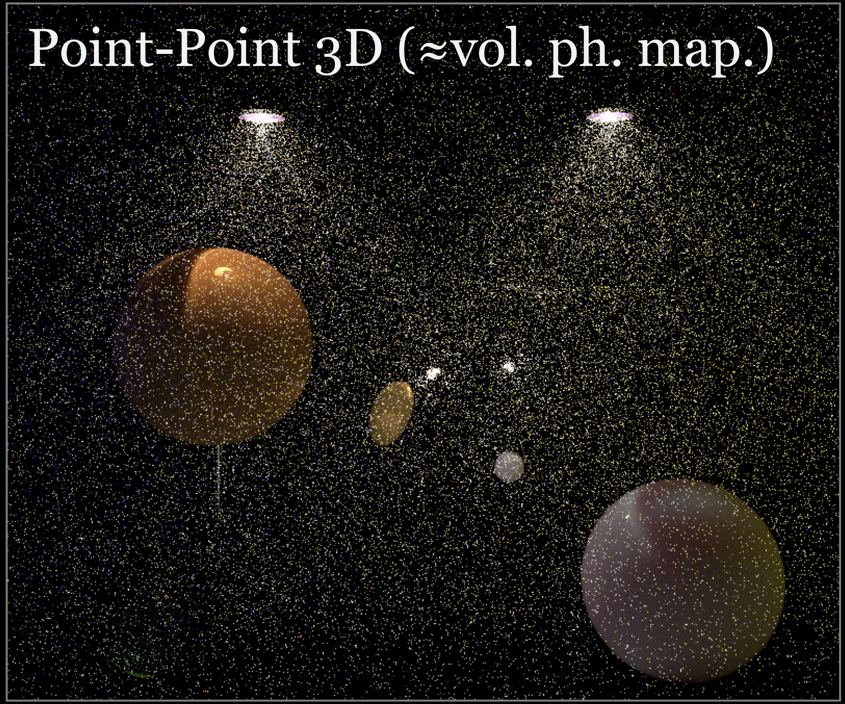

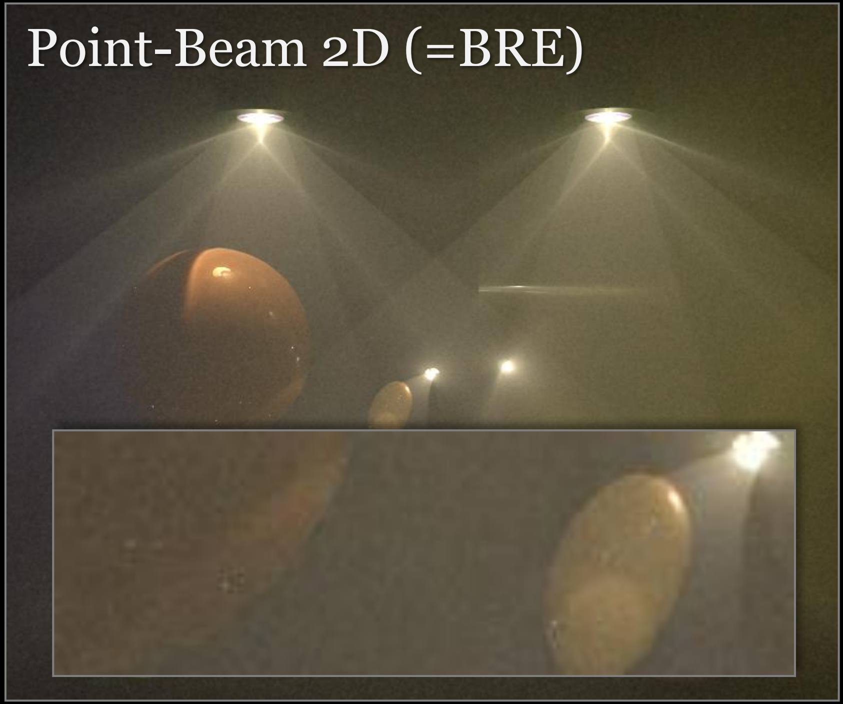

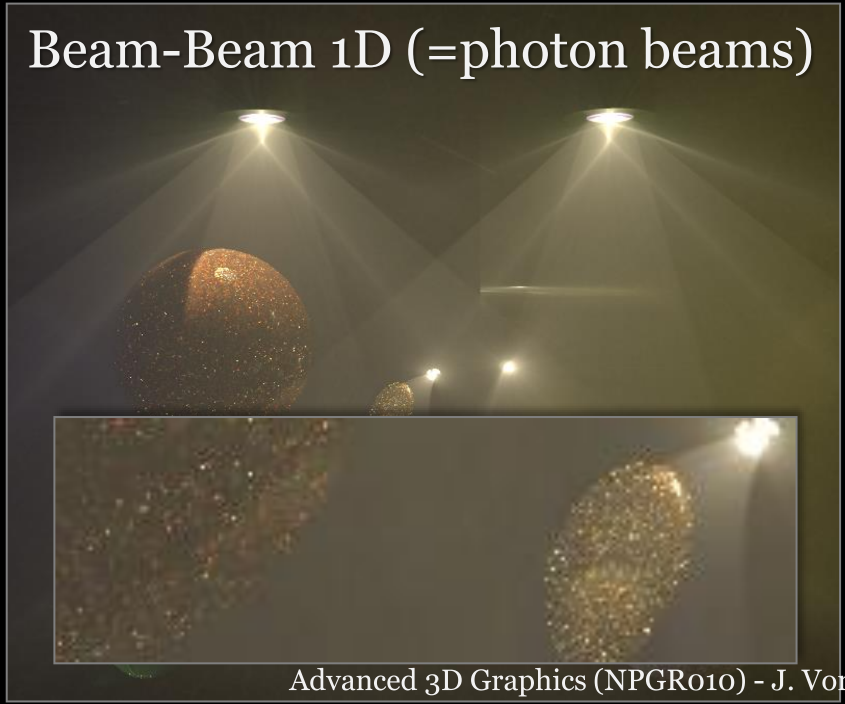

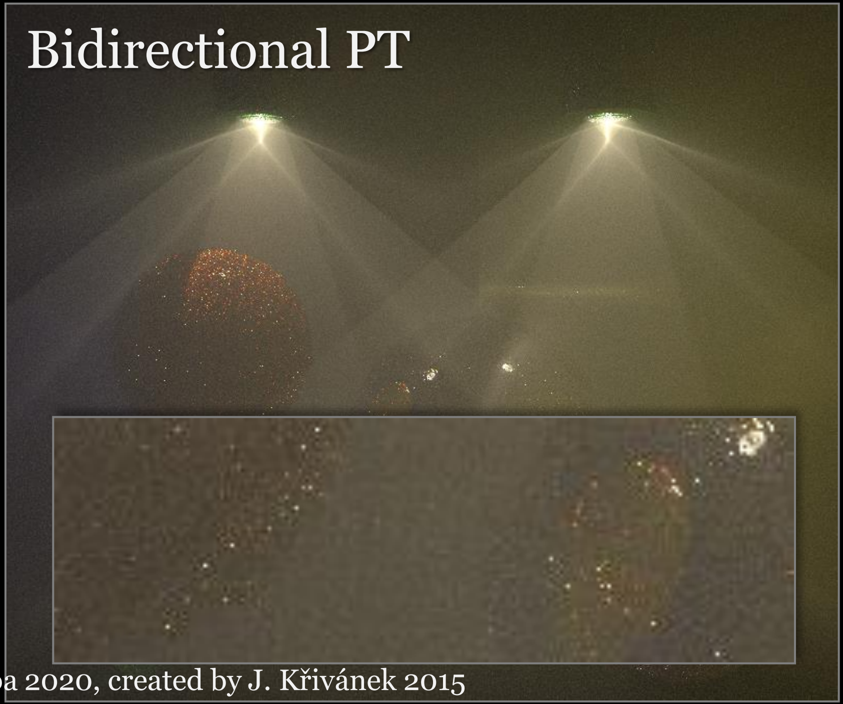

Previous work comparison, 1 hr

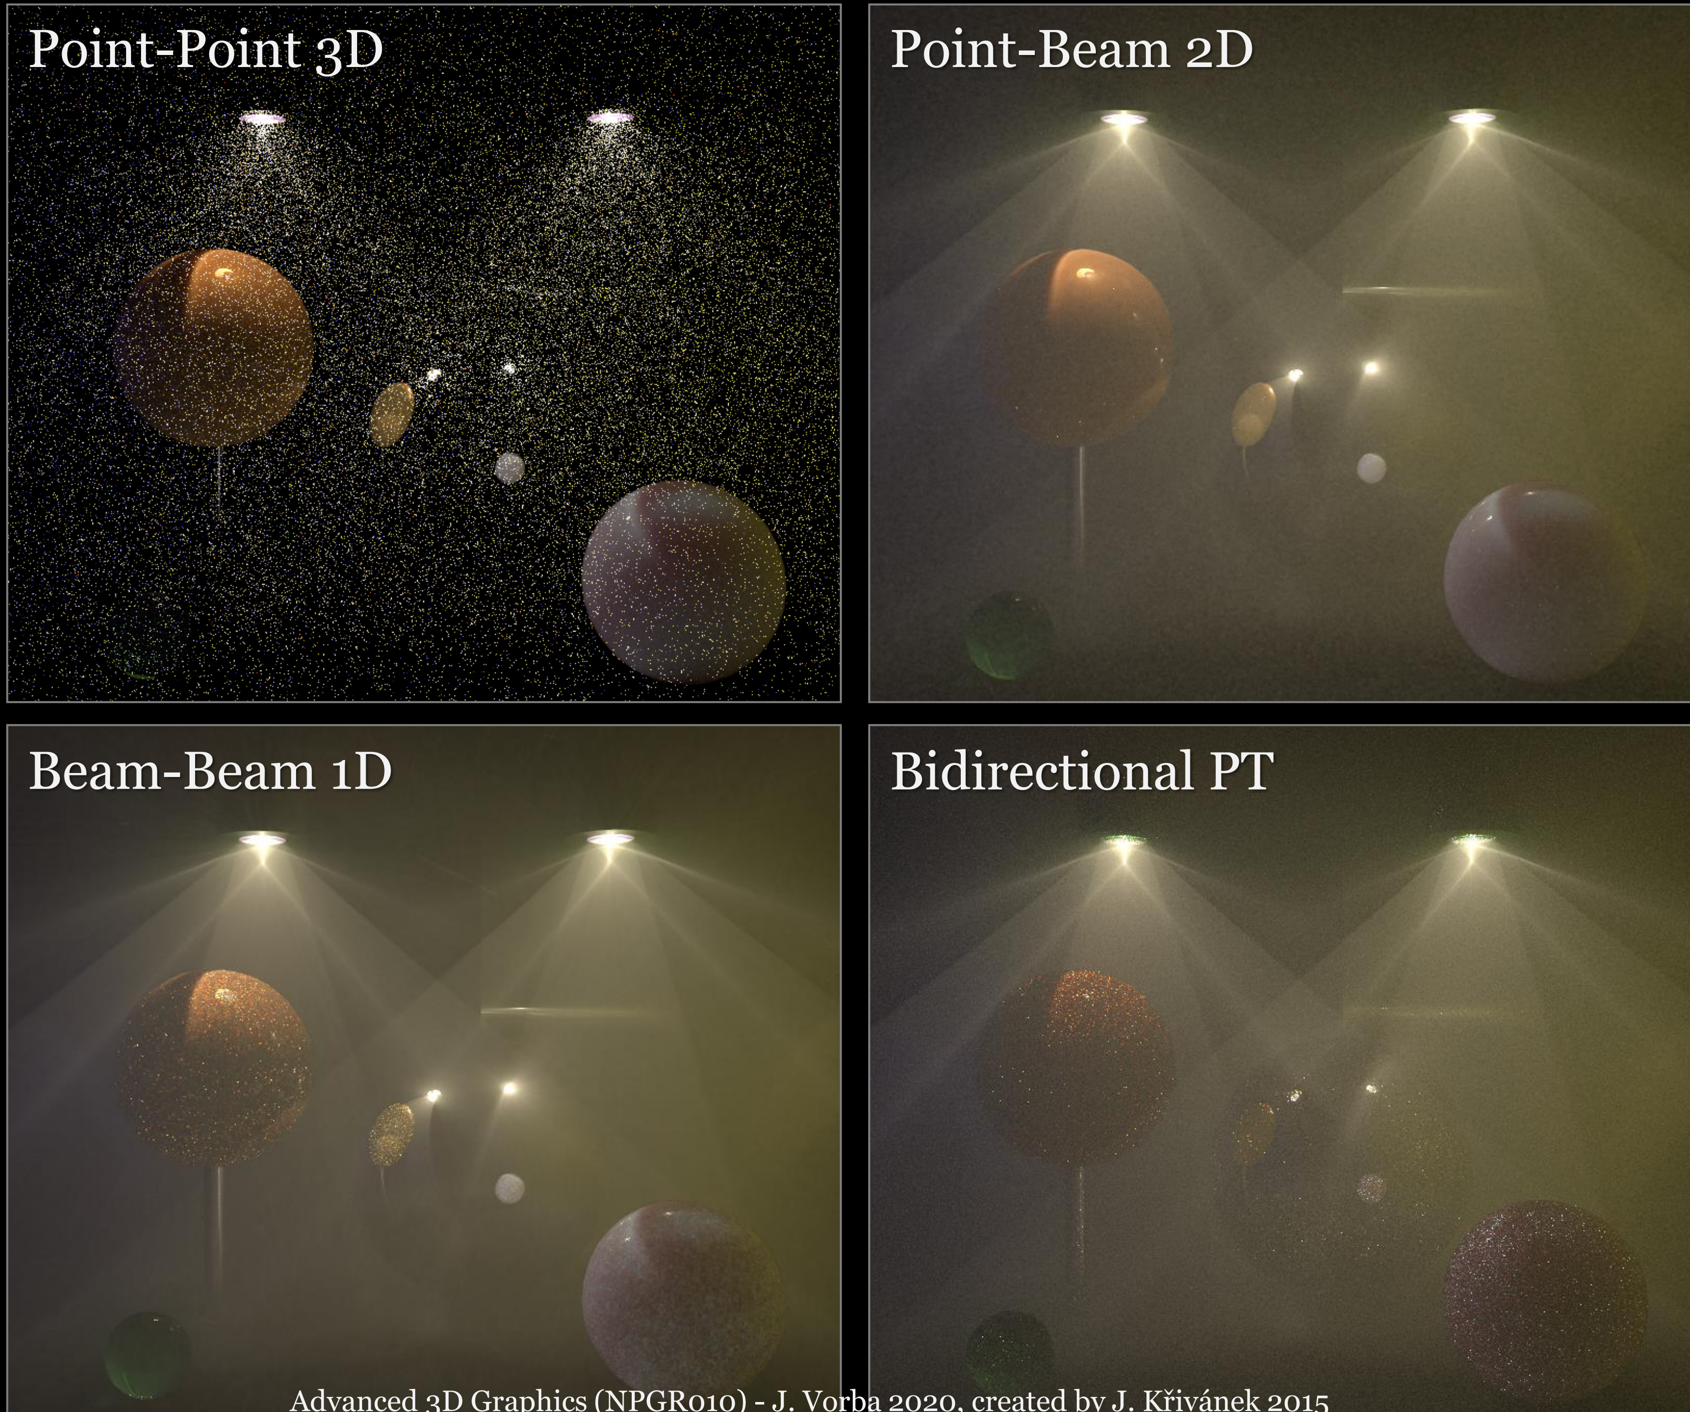

Weighted contributions

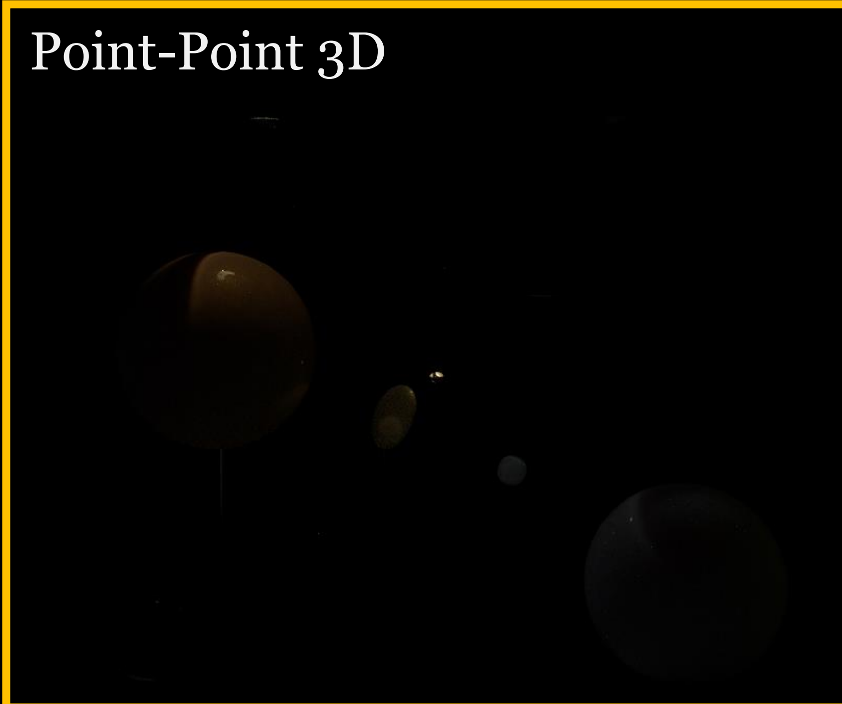

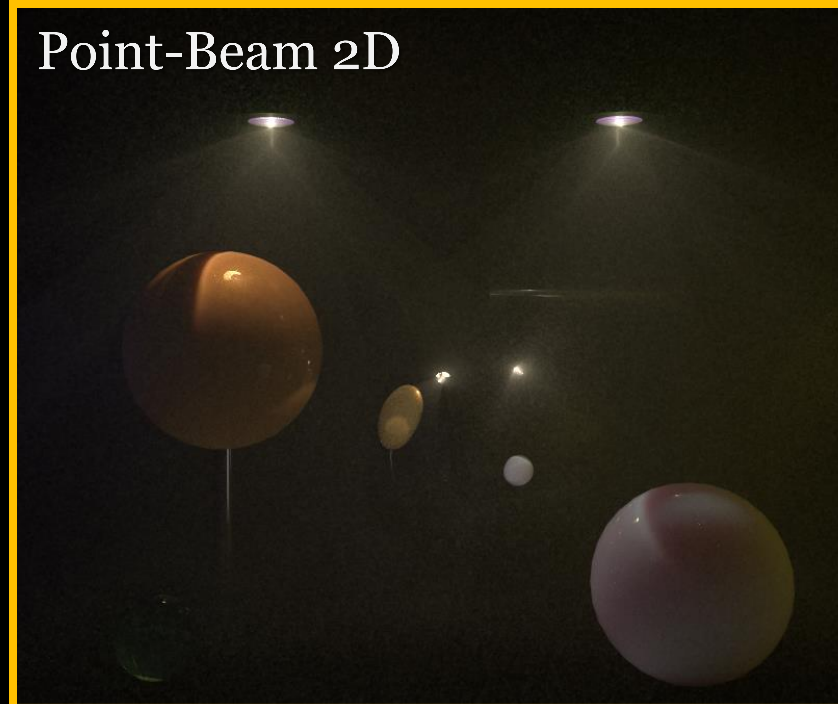

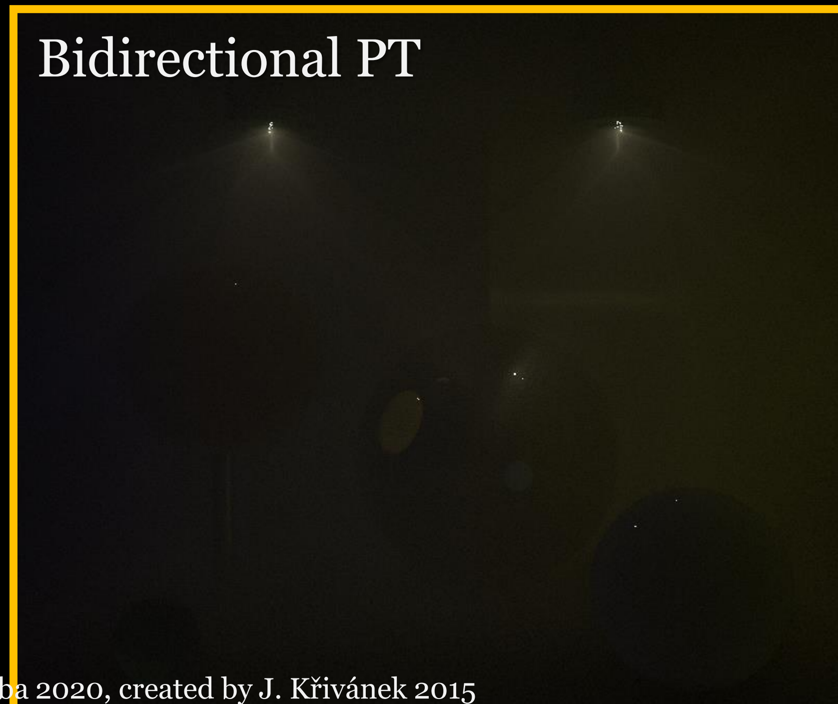

UPBP (our algorithm)
1 hour