



Product Importance Sampling for Light Transport Path Guiding

Sebastian Herholz¹

Oskar Elek²

Jiří Vorba^{2,3}

Hendrik Lensch¹

Jaroslav Křivánek²

¹University Tübingen

²Charles University Prague

³Weta Digital

Motivation

Reference
(4 weeks)



BDPT

(1hr)



BDPT
(1hr)



Vorba2014
(1hr)



Our
(1hr)



Vorba2014
(1hr)



Our
(1hr)


Vorba2014
(1hr)



Light Transport: Rendering Equation

$$L_o = L_e + \underbrace{\int_{\Omega} f_r \cdot L_i \cdot \cos\theta \cdot d\vec{\omega}_i}_{L_R}$$

emission




direct



indirect



combined

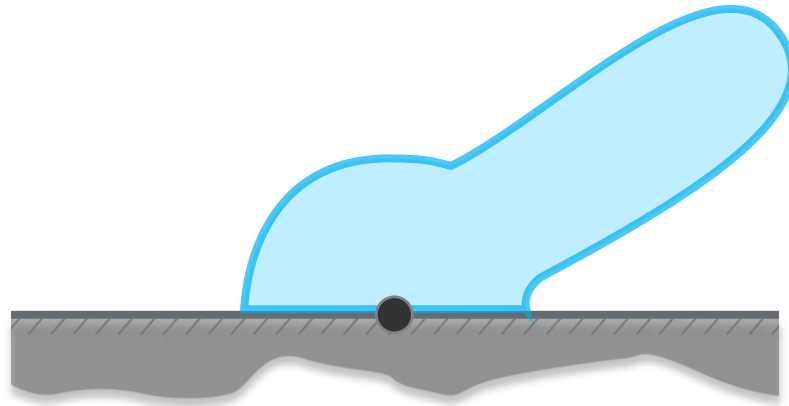


Light Transport: Rendering Equation

$$L_o = L_e + \underbrace{\int_{\Omega} f_r \cdot L_i \cdot \cos\theta \cdot d\vec{\omega}_i}_{L_R}$$




Bidirectional Reflectance Distribution Function (BRDF)



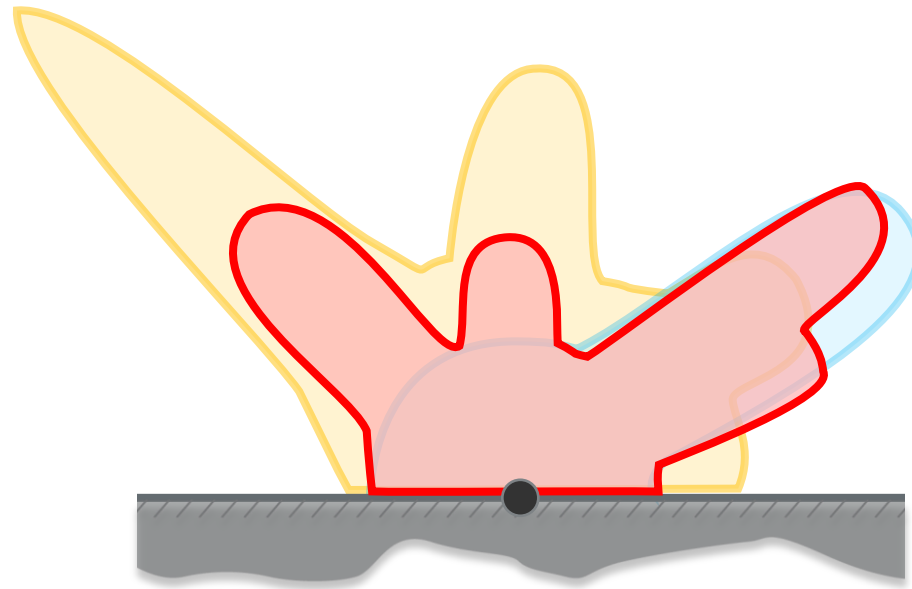
$$L_R = \int_{\Omega_i} \boxed{f_r} \cdot L_i \cdot \cos\theta \cdot d\vec{\omega}_i$$

Incomming Illumination



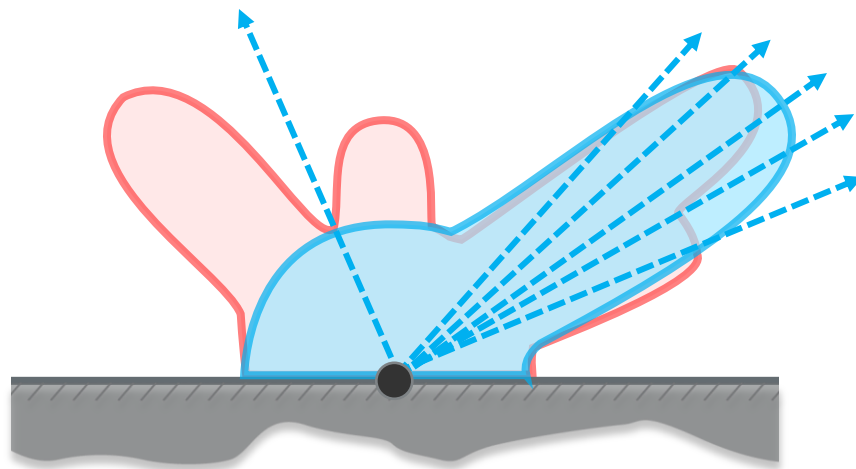
$$L_R = \int_{\Omega} f_r \cdot L_i \cdot \cos\theta \cdot d\vec{\omega}_i$$

Reflectance Integral



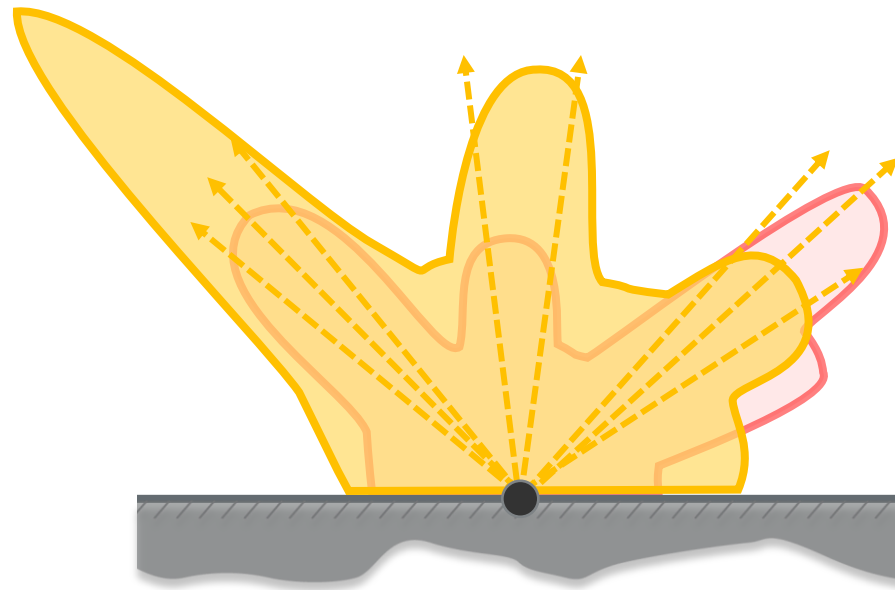
$$L_R = \int_{\Omega} f_r \cdot L_i \cdot \cos\theta \cdot d\vec{\omega}_i$$

BRDF-based Sampling




$$p_{f_r}(\omega_i | \omega_o, x) \propto f_r(x, \vec{\omega}_o, \vec{\omega}_i)$$

Guided Illumination Sampling



$$p_L(\omega_i | \omega_o, x) \propto L_i(x, \vec{\omega}_i) \cdot \cos\theta$$

Optimal (Product) Sampling



$$p_{opt}(\omega_i | \omega_o, x) \propto f_r(x, \vec{\omega}_o, \vec{\omega}_i) L_i(x, \vec{\omega}_i) \cdot \cos\theta$$

Related Work



[CAM08]: Practical product importance sampling for direct illumination



[TCE05]: Importance resampling for global illumination



[CJAMJ05]: Wavelet importance sampling: efficiently evaluating products of complex functions

[JCJ09]: Importance sampling spherical harmonics

Product Importance Sampling

Motivation

Product Importance Sampling

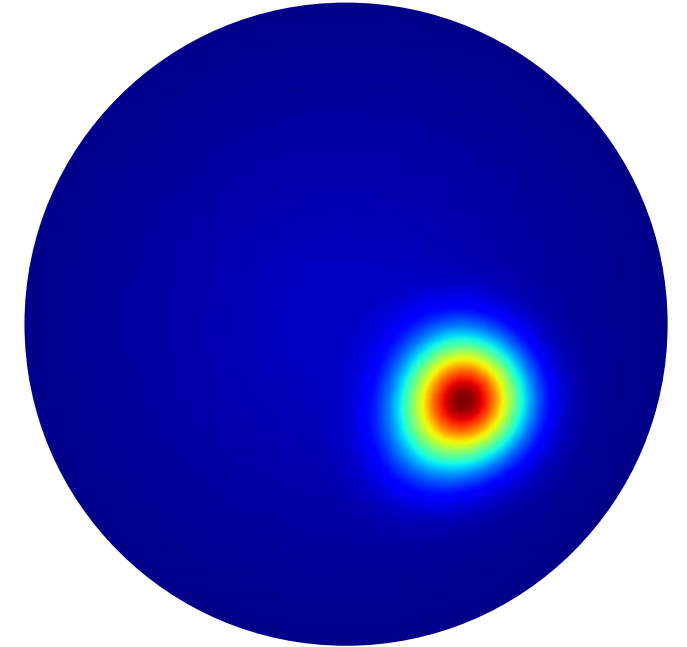

BRDF Fitting

Component Reduction

Results



Future Work

BRDF GMM Representation




$$p_{f_r}(\omega_i | \omega_o, x) \approx G_{f_r}(y, \Theta)$$

Illumination GMM Representation



$$p_L(\omega_i | \omega_o, x) \approx G_L(y, \Theta)$$


Gaussian Mixture Model (GMM)



$$G(y, \Theta) = \sum^K \pi_i N(y, \mu_i, \Sigma_i)$$

$$\Theta = \{\pi_0 \dots, \mu_0 \dots, \Sigma_0 \dots\}$$


Gaussian Mixture Model (GMM)



$$G(y, \Theta) = \sum^K \pi_i N(y, \mu_i, \Sigma_i)$$

$$\Theta = \{\pi_0 \dots, \mu_0 \dots, \Sigma_0 \dots\}$$


Gaussian Mixture Model (GMM)



$$G(y, \Theta) = \sum^K \pi_i N(y, \mu_i, \Sigma_i)$$

$$\Theta = \{\pi_0 \dots, \mu_0 \dots, \Sigma_0 \dots\}$$


Gaussian Mixture Model (GMM)



$$G(y, \Theta) = \sum^K \pi_i N(y, \mu_i, \Sigma_i)$$


$$\Theta = \{\pi_0 \dots, \mu_0 \dots, \Sigma_0 \dots\}$$

Product of two Gaussians is a Gaussian




$$\pi_i N(y, \mu_i, \Sigma_i) \cdot \pi_j N(y, \mu_j, \Sigma_j) = \pi_{ij} N(y, \mu_{ij}, \Sigma_{ij})$$

Product of two Gaussians is a Gaussian




$$\pi_i N(y, \mu_i, \Sigma_i) \cdot \pi_j N(y, \mu_j, \Sigma_j) = \pi_{ij} N(y, \mu_{ij}, \Sigma_{ij})$$

Product of two Gaussians is a Gaussian




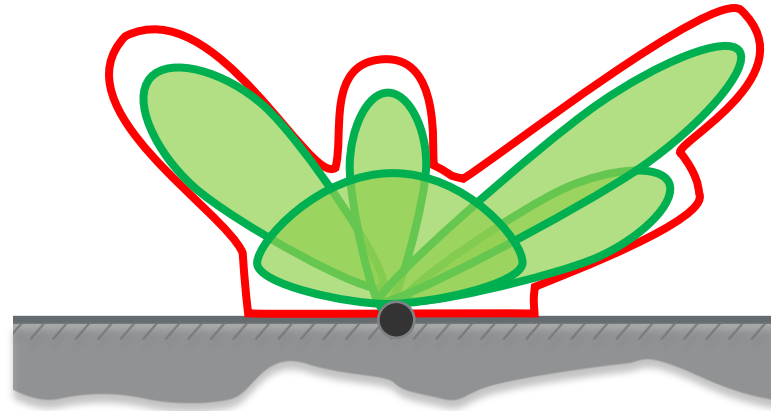
Product of two Gaussians is a Gaussian



$$\pi_i N(y, \mu_i, \Sigma_i) \cdot \pi_j N(y, \mu_j, \Sigma_j) = \pi_{ij} N(y, \mu_{ij}, \Sigma_{ij})$$

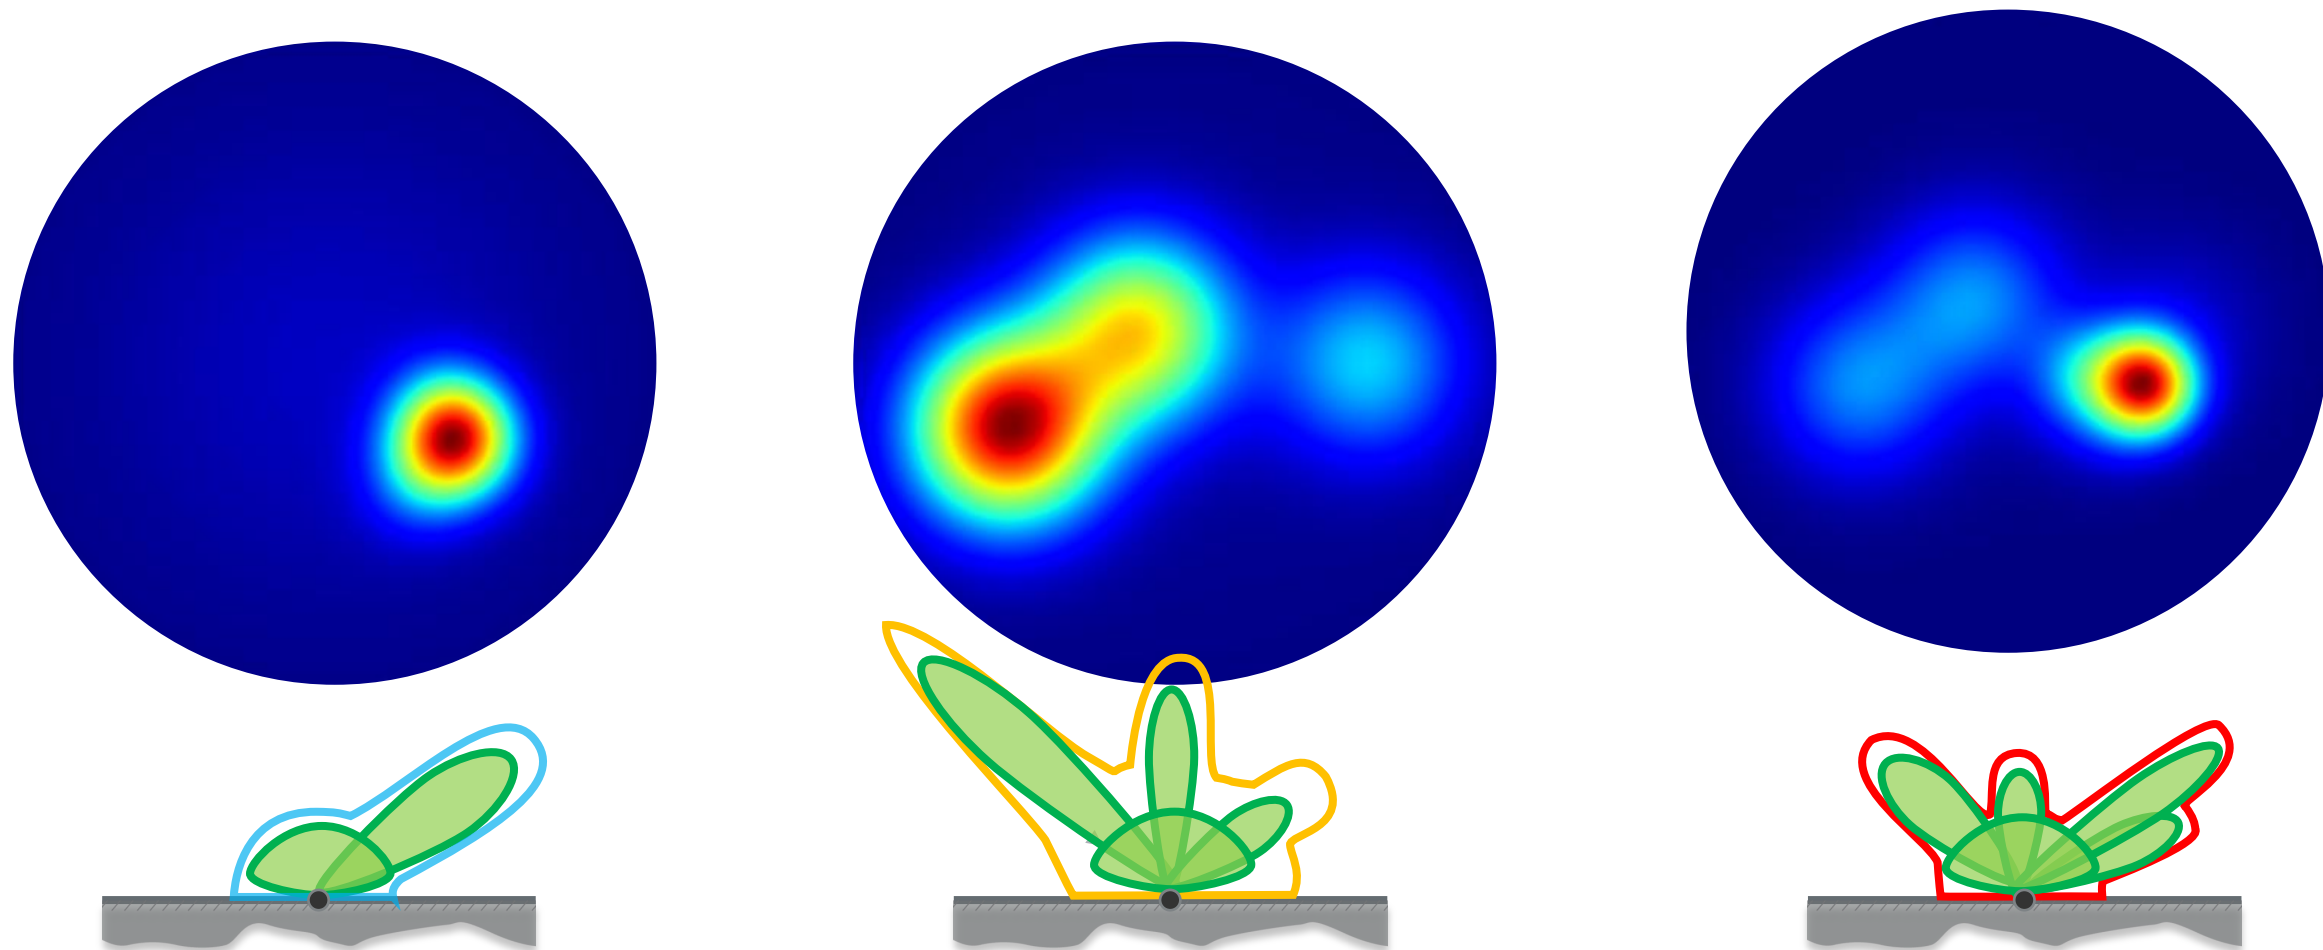
- Full GMM product contains K^2 components

Product GMM Representation



$$p_{f_r} \otimes p_L \approx G_{f_r} \otimes G_L = G_{\otimes}(y, \Theta)$$

Product GMM Representation



Pipeline

Pre-Processing



Rendering



Pipeline

Pre-Processing

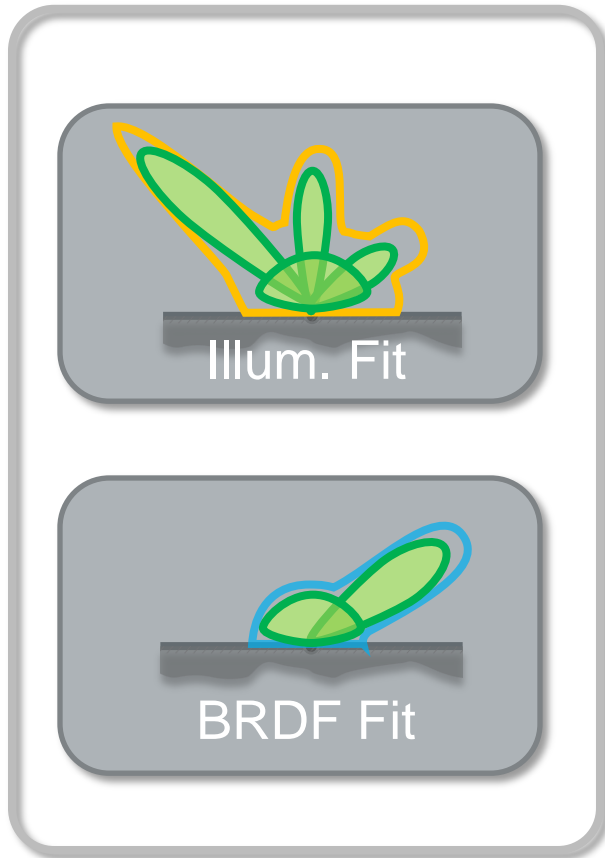


Rendering



Pipeline

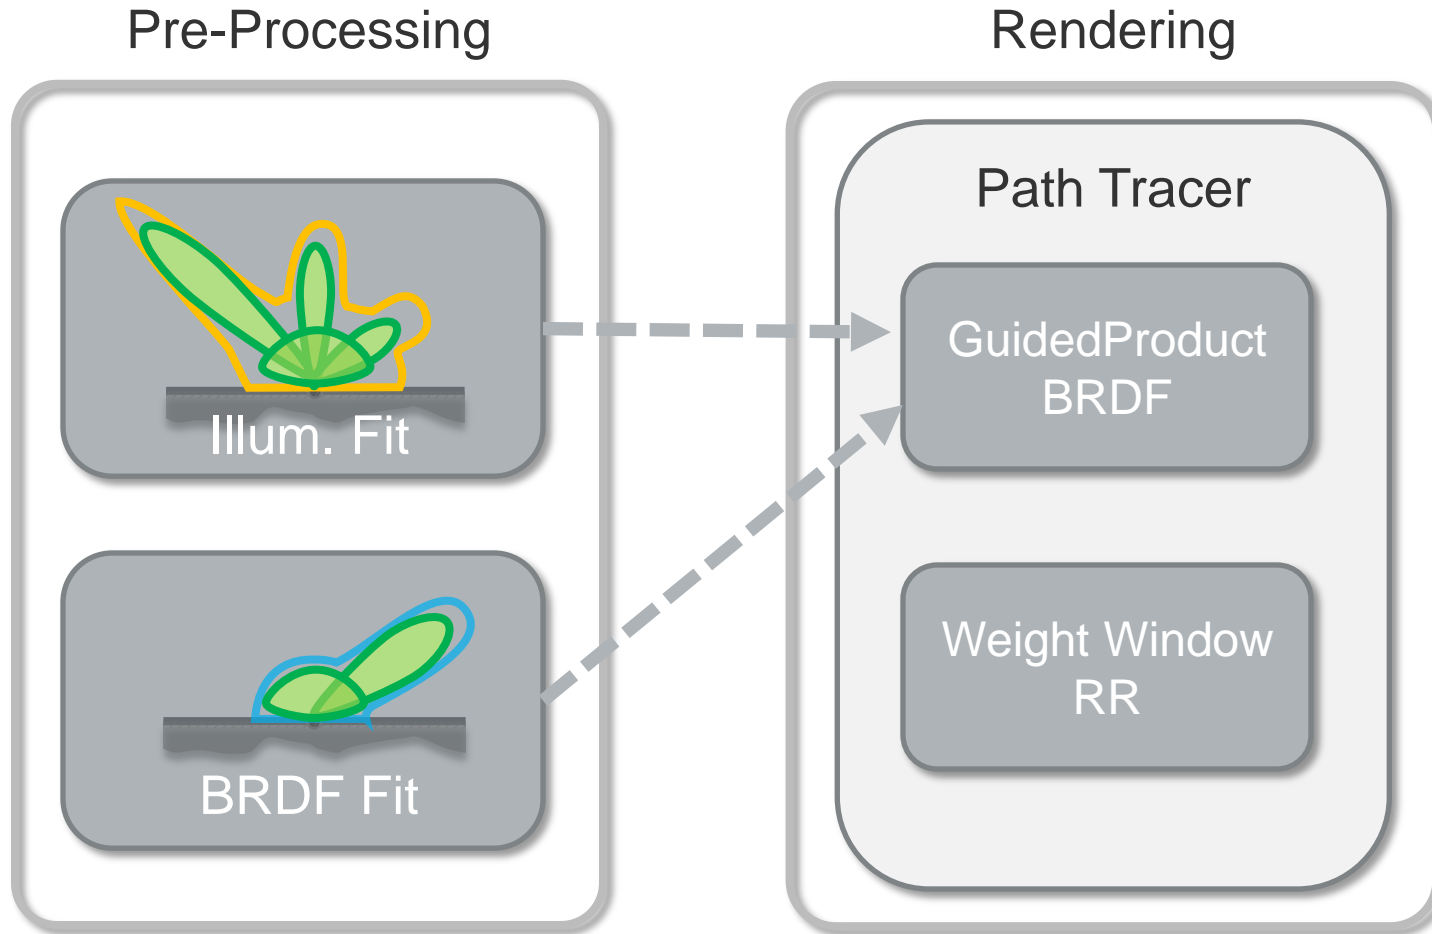
Pre-Processing



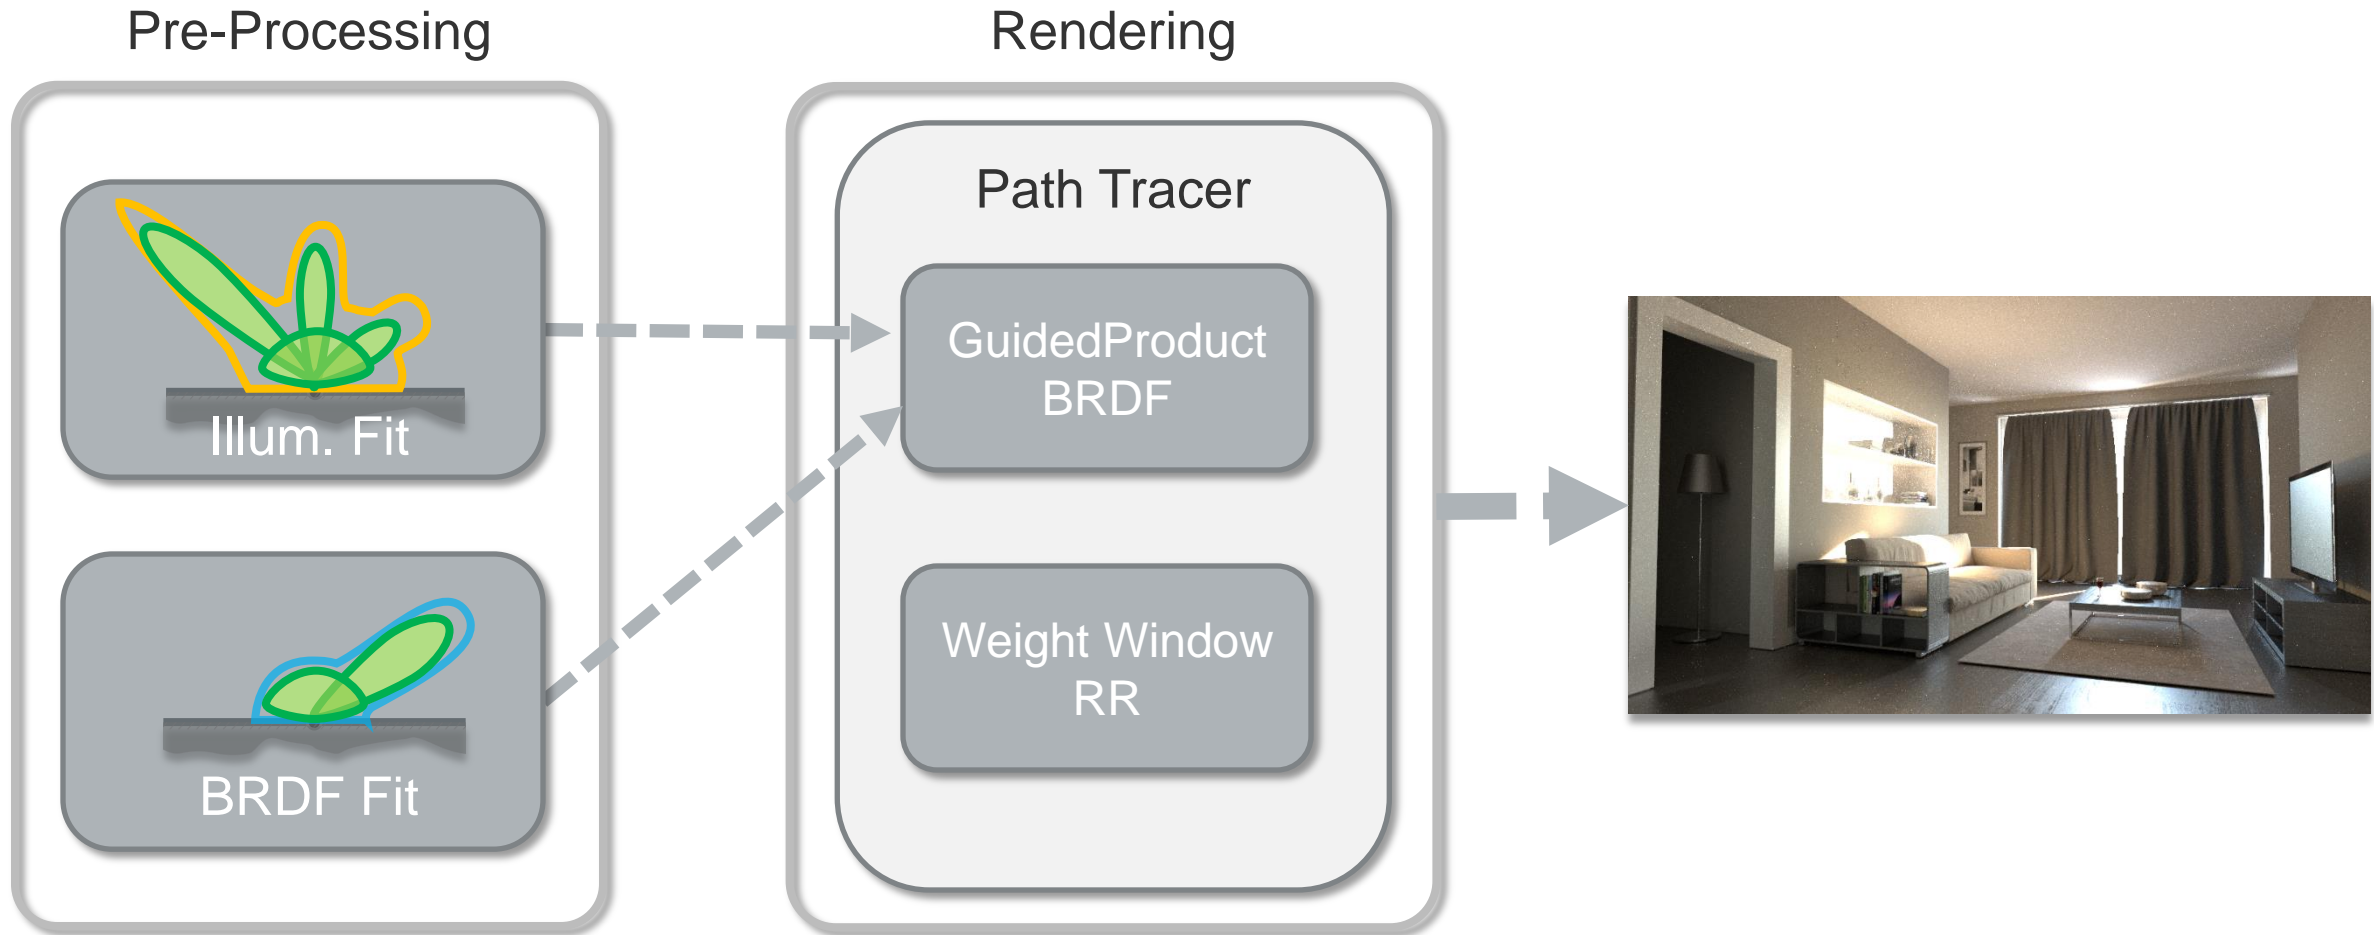
Rendering



Pipeline




Pipeline



Illumination Fit: [Vorba2014]

GMM Illumination caches



[Vorba2014]

On-line Learning of Parametric Mixture Models for Light Transport Simulation

Jiří Vorba^{1*} Ondřej Karlík^{1*} Martin Šik^{1*} Tobias Ritschel^{2†} Jaroslav Krivánek^{1‡}
¹Charles University in Prague ²MPI Informatik, Saarbrücken

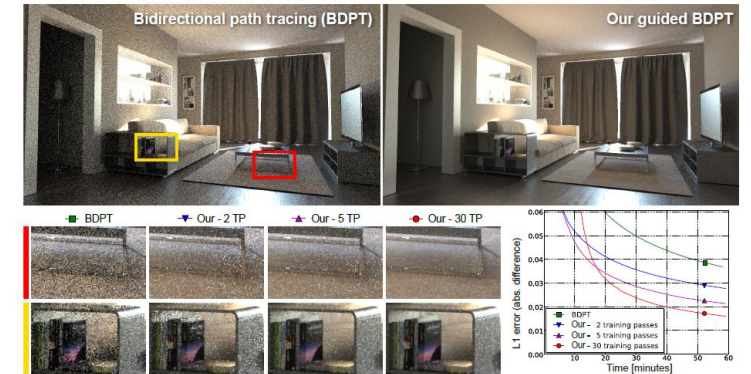


Figure 1: We render a scene featuring difficult visibility with bidirectional path tracing (BDPT) guided by our parametric distributions learned on-line in a number of training passes (TP). The insets show equal-time (1h) comparisons of images obtained with different numbers of training passes. The results reveal that the time spent on additional training passes is quickly amortized by the superior performance of the subsequent guided rendering.

Abstract

Monte Carlo techniques for light transport simulation rely on importance sampling when constructing light transport paths. Previous work has shown that suitable sampling distributions can be recovered from particles distributed in the scene prior to rendering. We propose to represent the distributions by a parametric mixture model trained in an on-line (i.e. progressive) manner from a potentially infinite stream of particles. This enables recovering good sampling distributions in scenes with complex lighting, where the necessary number of particles may exceed available memory. Using these distributions for sampling scattering directions and light emission significantly improves the performance of state-of-the-art light transport simulation algorithms when dealing with complex lighting.

CR Categories: I.3.3 [Computer Graphics]: Three-Dimensional Graphics and Realism—Display Algorithms

Keywords: light transport simulation, importance sampling, parametric density estimation, on-line expectation maximization

Links: [DL](#) [PDF](#) [WEB](#) [CODE](#)

* {jirka.karlik,martin.sik}@cg.mff.cuni.cz
†ritschel1@mpi-inf.mpg.de
‡jaroslav.krivanek@mff.cuni.cz

1 Introduction

Despite recent advances, robust and efficient light transport simulation is still a challenging open issue. Numerous algorithms have been proposed to solve the problem, but certain common lighting conditions, such as highly occluded scenes, remain difficult. Most existing unidirectional and bidirectional methods rely on incremental, local construction of transport sub-paths, which is oblivious to the global distribution of radiance or importance. As a result, the probability of obtaining a non-zero contribution upon sub-path connection in highly occluded scenes is low. This is the main reason why such scenes remain difficult to render. While Metropolis light transport and related methods [Veach and Guibas 1997; Kelemen et al. 2002; Cline et al. 2005] strive for importance sampling on the entire path space, they suffer from sample correlation and are often outperformed by the classic Monte Carlo approaches.

BRDF Fitting and Caching

Motivation

Product Importance Sampling

BRDF Fitting

Component Reduction

Results

Future Work

Fitting BRDF GMM

Weighted EM

- Weighted MAP EM [Vorba2014]
- Sample BRDF (N=512)
- $W_i = \frac{f_r(x, \omega_i, \omega_o)}{p(\omega_o)}$
- Init components using K BRDF samples (QMC sampler)

CERES

- Non-linear optimization
- Init with weighted EM
- Objective function:

$$\sum_i^N \left[1 - \frac{\hat{f}_r(\omega_i)}{G(y|\Theta)} \right]^2$$

Fitting BRDF GMM

Weighted EM


- Weighted MAP EM [Vorba2014]
- Sample BRDF (N=512)
- $W_i = \frac{f_r(x, \omega_i, \omega_o)}{p(\omega_o)}$
- Init components using K BRDF samples (QMC sampler)

CERES

- Non-linear optimization
- Init with weighted EM
- Objective function:


$$\sum_i^N \left[1 - \frac{\hat{f}_r(\omega_i)}{G(y|\Theta)} \right]^2$$

WEM vs CERES




rough conductor $\alpha = 0.3$
 $K = 8$

wEM vs CERES




rough conductor $\alpha = 0.3$
 $K = 8$

wEM vs CERES



rough conductor $\alpha = 0.3$
 $K = 8$


wEM vs CERES



rough conductor $\alpha = 0.15$
 $K = 8$


Caching

- Isotropic



- 512 different elevation angles



- Anisotropic



- 4096 spherical Fibonacci points [KISS15]



Caching

- Isotropic





Caching

- Isotropic





Caching

- Isotropic





Caching

- Isotropic



Caching

- Isotropic



GMM Component Reduction

Motivation

Product Importance Sampling


BRDF Fitting

Component Reduction

Results


Future Work

Component Reduction




$$c_{i,j} = d_{KL}(\{(\pi_i, \mu_i, \Sigma_i), (\pi_j, \mu_j, \Sigma_j)\}, (\pi_{ij}, \mu_{ij}, \Sigma_{ij}))$$


Component Reduction



Component Reduction

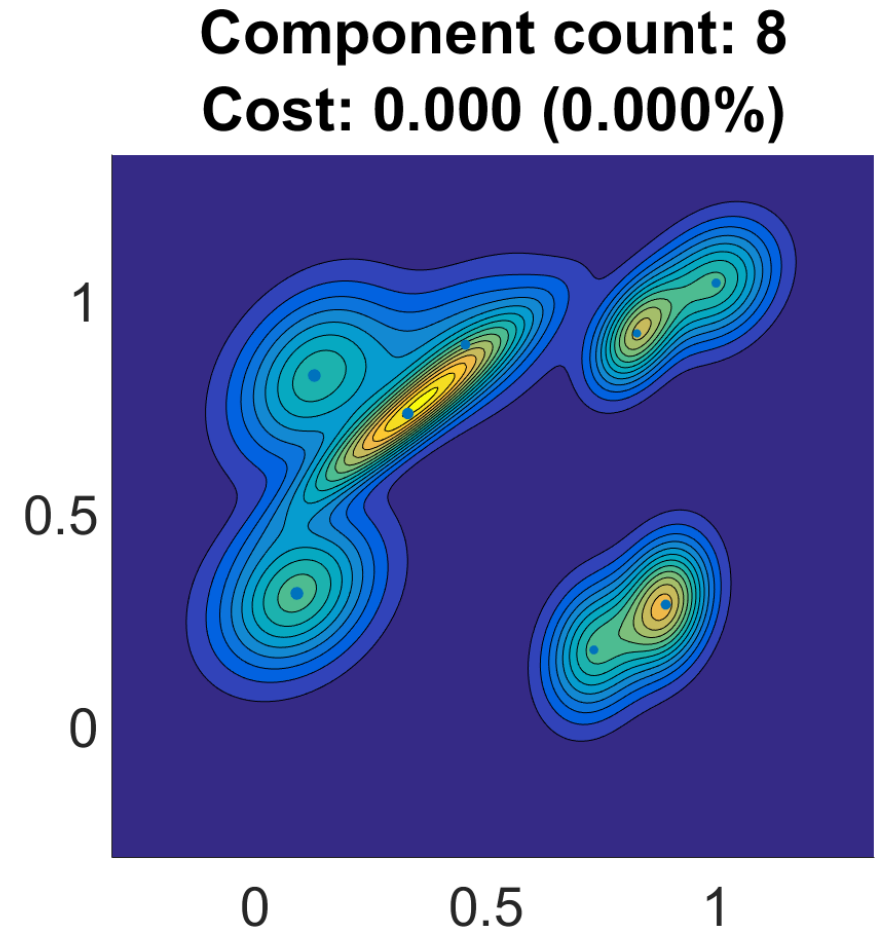



Component Reduction





- Adapted from [Runnals2007], used also in [Jacob2011]
- Kullback-Leibler discrimination: d_{KL}



Component Reduction




Component Reduction




Component Reduction




Component Reduction




Component count: 5
Cost: 0.269 (10.917%)





Component Reduction




Component count: 4
Cost: 0.498 (20.210%)




Component Reduction





Component Reduction





Component count: 2
Cost: 1.443 (58.549%)




Component Reduction



Component Reduction




Component Reduction



Reduction

- BRDF:
 - Full $K = 8$
- Illumination
 - Full $K = 8$
- Product GMM:
 - Avg. $K_{ij} = 64$

Component Reduction



Reduction

- BRDF:
 - Full $K = 8$
 - **Red. avg. $K = 2$**
- Illumination
 - Full $K = 8$
 - **Red. 50% to 4 comp.**
- Product GMM:
 - Avg. $K_{ij} = 64$
 - **Red. avg. $K_{ij} = 12$**

Results



Reference

Reference




SPP / MSE



equal time 1hr

Path Tracer

Reference



4736 / 2.1830

SPP / MSE



equal time 1hr

Path Tracer

Vorba2014


Reference



4736 / 2.1830

882 / 0.0331

SPP / MSE





equal time 1hr

Path Tracer

Vorba2014

Our

Reference



4736 / 2.1830

882 / 0.0331

1128 / 0.0211

SPP / MSE



equal time 1hr

Path Tracer

Vorba2014

Our

Reference



4736 / 2.1830

882 / 0.0331

1128 / 0.0211

SPP / MSE



equal time 1hr



Reference





equal time 1hr

Path Tracer

Vorba2014

Our

Reference



4335 / 0.0081

1528 / 0.0025

1322 / 0.0007



SPP / MSE

Path Tracer

Vorba2014

Our

Reference




4335 / 0.0081

1528 / 0.0025

1322 / 0.0007

SPP / MSE



equal time 1hr







Reference

Path Tracer

Vorba2014

Our

Reference



3120 / 0.9715

1616 / 0.006


712 / 0.007

SPP / MSE



equal time 1hr

equal time 1hr






Path Tracer

Vorba2014

Our

Reference



3120 / 0.9715

1616 / 0.006

712 / 0.007


SPP / MSE

Average Path Lengths

Vorba: 9.8 **Ours: 4.9**

normalized per-path segment contribution

LivingRoom




Average Path Lengths

Vorba: 9.8 **Ours: 4.9**

normalized per-path segment contribution

LivingRoom




Average Path Lengths

Vorba: 9.8 **Ours: 4.9**

normalized per-path segment contribution

LivingRoom




Average Path Lengths

Vorba: 6.5 **Ours: 9.0**

normalized per-path segment contribution


Kitchen



Discussion / Future work

Discussion / Future Work

- Path Length and Russian Roulette
 - Adjoint-driven RR and Splitting [Vorba2016]




Discussion / Future Work

- SVBRDFs
 - Enlarge BRDF caches
 - Direct function transform
BRDF->GMM

Scene	BRDF caching			
	# BRDFs	# Caches	Avg. # comp.	Mem.
LIVINGROOM	41	15k	2.5	7.7 MB
KITCHEN	72	2.5k	1.8	10 MB
JEWELRY	6	1.5k	1.44	0.7 MB

Discussion / Future Work

- Extension to other MC-algorithms
 - BDPT
 - MCMC
 - Gradient domain



GD-PT
(full)



GD-BDPT
(full)

Reference



GD-PT
(estimate)



GD-BDPT
(estimate)



Reference



Our



Reference



Discussion/Future Work

- Optimizing Illumination caches
 - Poorly fitted illumination caches cause inconsistent convergence rates



Product Importance Sampling for Light Transport Path Guiding

Sebastian Herholz¹


Oskar Elek²

Jiří Vorba^{2,3}


Hendrik Lensch¹

Jaroslav Křivánek²


¹University Tübingen



²Charles University Prague



³Weta Digital




Thanks to:

- Martin Šik, Ivo Kondapaneni, Ludvík Koutný, Anton Kaplanyan, Johannes Hanika
- Anonymous reviewers
- You!


Backup

wEM vs CERES


wEM




Ceres




Gradient Domain




Reference









Uncorrected



Corrected



Standard



Preprocessing times:

Table 1: *Timings for the cache fitting stages (in minutes).*

Scene	Illumination	BRDF	
		wEM	Ceres
LIVINGROOM	14.0	0.31	25.5
KITCHEN	20.1	0.44	42.3
JEWELRY	6.1	0.04	6.3

Statistics:

Table 2: Left and middle: *BRDF and illumination caching statistics for the scenes in Fig. 8.* Right: *Overhead of the product sampling relative to illumination-only sampling, without ('naïve') and with the reduction of the BRDF and illumination mixtures.*

Scene	BRDF caching				Illumination caching			Sampling overhead [%]	
	# BRDFs	# Caches	Avg. # comp.	Mem.	# Caches	# Reduced	Mem.	Naïve	Reduced
LIVINGROOM	41	15k	2.5	7.7 MB	82k	57 %	192.9 MB	10.8	7.1
KITCHEN	72	2.5k	1.8	10 MB	107k	62 %	236.9 MB	26.7	9.9
JEWELRY	6	1.5k	1.44	0.7 MB	16k	33 %	19.5 MB	16.5	-1.1

END