Real Time Neural Path Guiding

Bachelor thesis presentation

Dominik Wüst | April 28, 2023

Contents

What is Path Tracing?

Radiative Transport Equation (RTE) [Kajiya 1986]:

$$L_o(x,\omega) = L_e(x,\omega) + \int_\Omega f_r(\omega,x,\omega_i) L_i(x,\omega_i) \cos\theta_i \mathrm{d}\omega_i \quad (1)$$

Properties:

- outgoing radiance L_o at point x in direction ω
- sum of emitted and reflected radiance
- recursive
- no analytical solution

Monte Carlo Path Tracing:

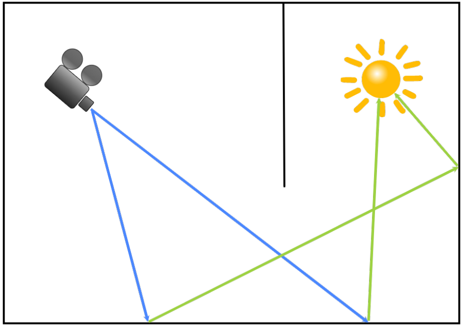

Monte Carlo Estimator for reflection integral [Müller, Rousselle, et al. 2021]:

$$\int_{\Omega} f_r(\omega, x_1, \omega_i) L_i(x_1, \omega_i) \cos\theta_i \mathrm{d}\omega_i \approx \frac{1}{N} \sum_{j=1}^{N} \frac{f_r(\omega, x_1, \omega_j) L_i(x_1, \omega_j) \cos\theta_j}{p(\omega_j)}$$

What is Path Guiding?

Karlsruhe Institute of Technology

Importance sampling of reflection integral:

$$p \propto f_r(\omega, x, \omega_i) L_i(x, \omega_i) \cos\theta_i$$

Goal

Learn spatially varying radiance distribution p

Challenges:

- Mathematical representation of the distribution
  - Gaussian Mixture Model (GMM) [Herholz et al. 2016]
  - Octree [Müller, Gross, and Novák 2017]
- Partition spatial domain
  - Grid (regular) [Huang et al. 2023]
  - Binary tree [Müller, Gross, and Novák 2017]

Realtime Path Guiding

	Offline Rendering	Realtime Rendering	
Time per frame	a few minutes	max 16 ms	temporal reuse
Dynamic content	no	yes	no pre-computation
Hardware	Mostly CPU	GPU	expensive memory access

Reminder: Path guiding is data-driven!

Von Mises-Fisher distribution

Probability Density Function (PDF):

$$f_{\text{vMF}}(\omega; \mu, \kappa) = \begin{cases} \frac{1}{4\pi}, & \text{if } \kappa = 0 \\ \frac{\kappa}{2\pi(1-\exp(-2\kappa))} e^{\kappa\mu^T\omega - 1}, & \text{if } \kappa > 0 \end{cases} \tag{2}$$

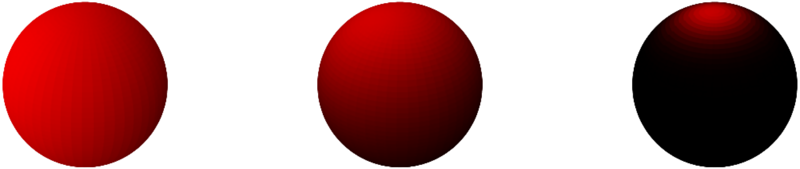

$\kappa = 0.1$ $\kappa = 1$ $\kappa = 10$

Real Time Neural Path Guiding - Concept

Real Time Neural Path Guiding - Concept

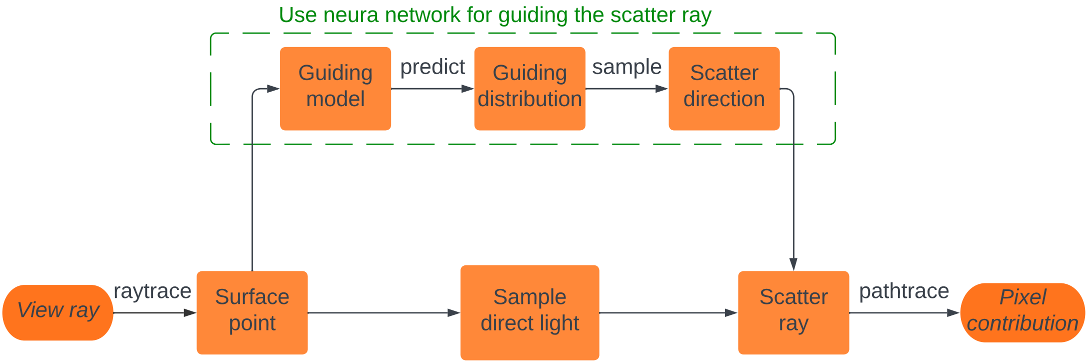

Real Time Neural Path Guiding - Concept

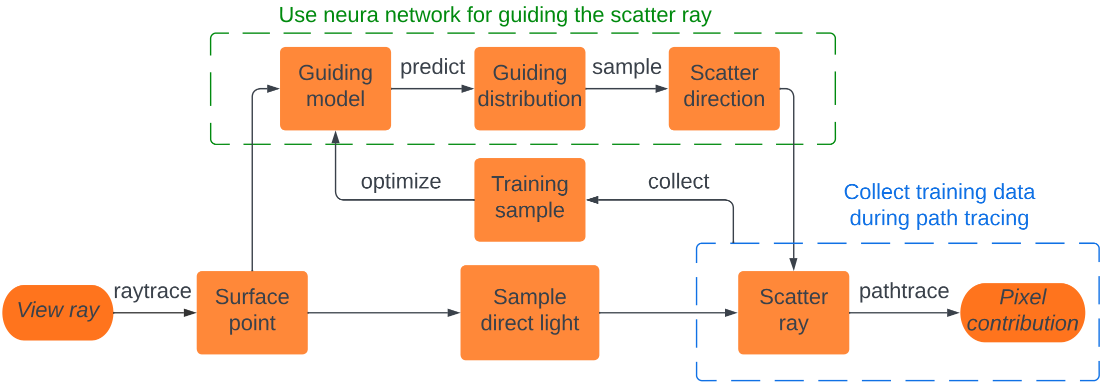

Guiding Model

Multilayer Perceptron (MLP):

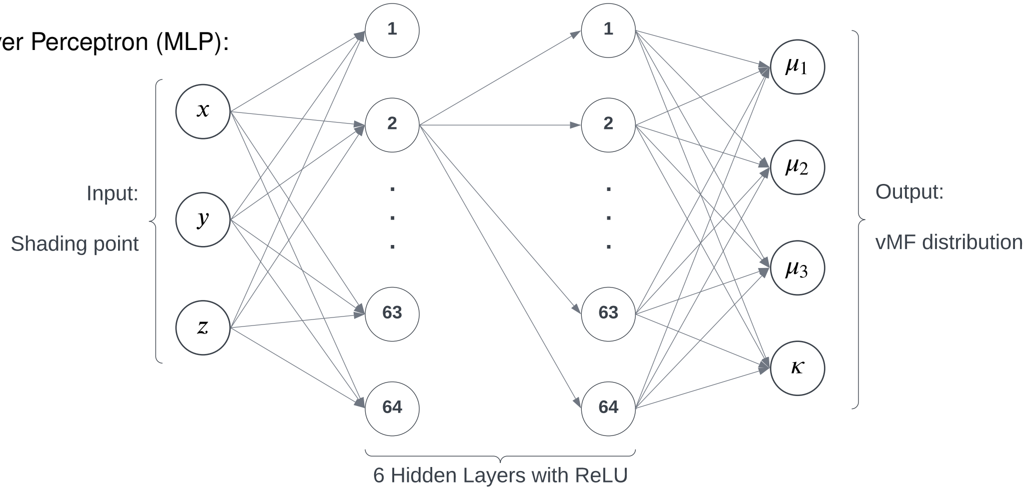

Training

We need to evaluate the fit of our vMF $f_{\text{vMF}}(\omega; \mu, \kappa)$ to the incident radiance $L_i(x, \omega)$!

$$f_L(\omega) := \frac{1}{A} L_i(x, \omega)$$

$A = \int_\Omega L_i(x, \omega_i) \mathrm{d}\omega_i$ is a (unknown) normalization factor
Kullback-Leibler divergence:

$$D_{\text{KL}}(f_L \parallel f_{\text{vMF}}) = \int_\Omega f_L(\omega) \cdot \log \left(\frac{f_L(\omega)}{f_{\text{vMF}}(\omega)} \right) \mathrm{d}\omega$$

One-Sample MC estimate:

$$\begin{aligned} D_{\text{KL}}(f_L(\omega_s) \parallel f_{\text{vMF}}(\omega_s)) &= \frac{f_L(\omega_s)}{p(\omega_s)} \cdot \log \left(\frac{L_i(x, \omega_s)}{f_{\text{vMF}}(\omega_s) \cdot A} \right) \\ &= \frac{f_L(\omega_s)}{p(\omega_s)} \cdot (\log(\underbrace{L_i(x, \omega_s)}_{\text{target}})) - \log(\underbrace{f_{\text{vMF}}(\omega_s) \cdot A}_{\text{prediction}})) \end{aligned}$$

Problem: $L_i(x, \omega_s)$ is noisy
relative squared loss:

$$\mathcal{L}^2_{\text{vMF}}(f_{\text{vMF}}(\omega_s), f_L(\omega_s)) = \frac{\left(L_i(x, \omega_s) - f_{\text{vMF}}(\omega_s) \cdot A\right)^2}{p(\omega_s)^2 \cdot \text{sg}(f_{\text{vMF}}(\omega_s) \cdot A)^2 + \epsilon} \tag{3}$$

Training target:

- Scatter direction ω_s
- PDF $p(\omega_s)$
- Incident radiance $L_i(x, \omega_s)$
- Normalization factor A

First results

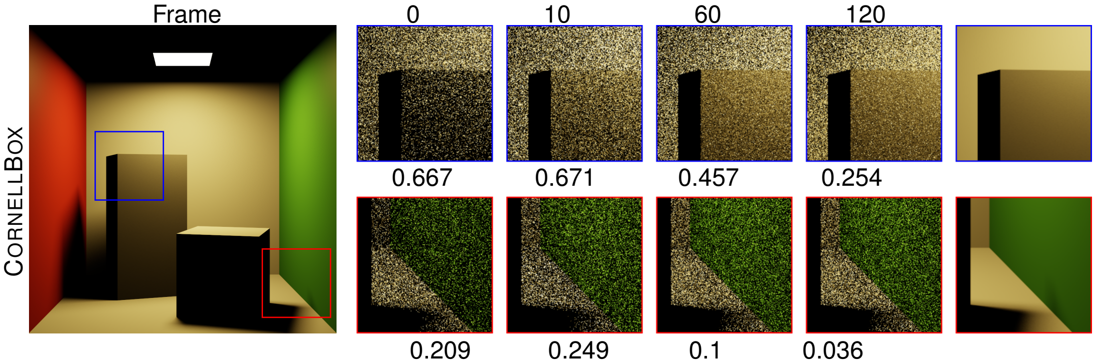

Figure: Improvements after *N* frames of training. We report the MSE for each cutout. The estimate converges faster at surfaces close to the light source (blue cutout).

Artifacts at geometric edges

- Initial setup works well on larger surfaces
- Bad results along edges

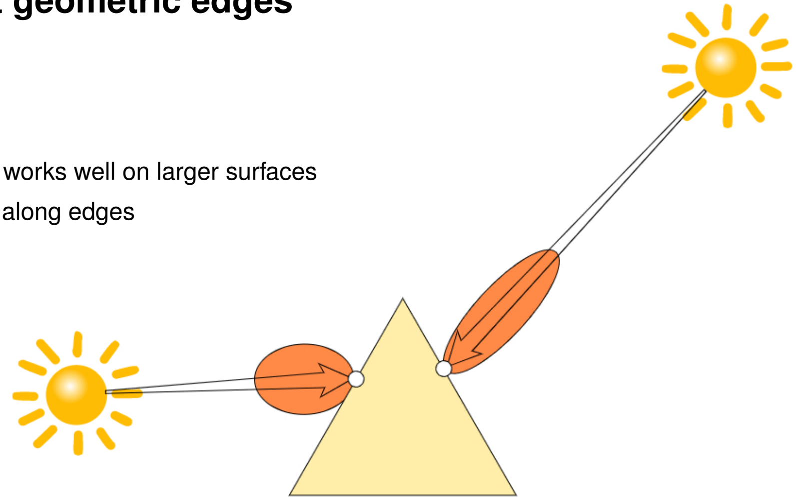

Artifacts at geometric edges

- Initial setup works well on larger surfaces
- Bad results along edges

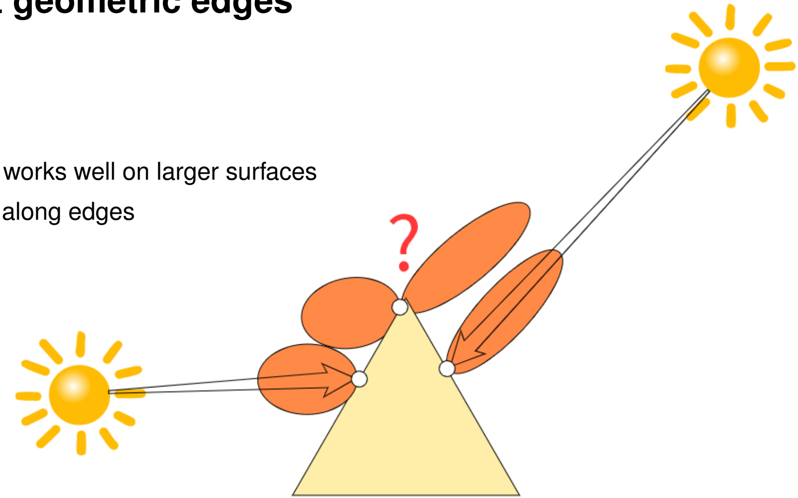

Artifacts at geometric edges

- Initial setup works well on larger surfaces
- Bad results along edges
- Solution: Surface normal as input

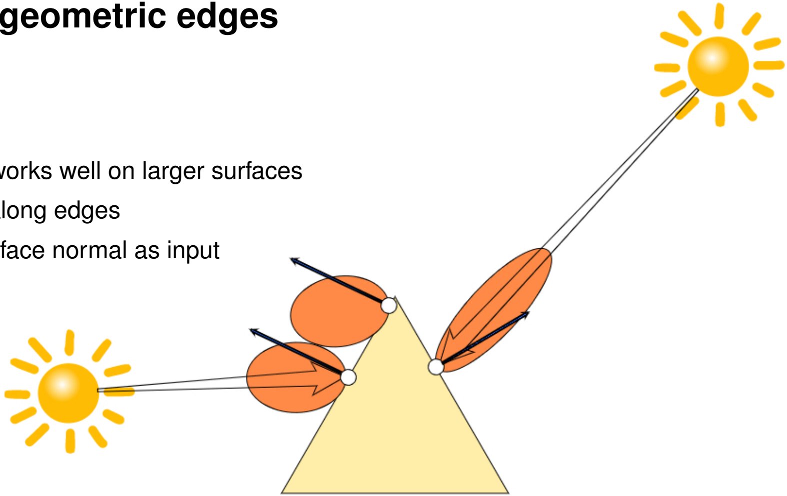

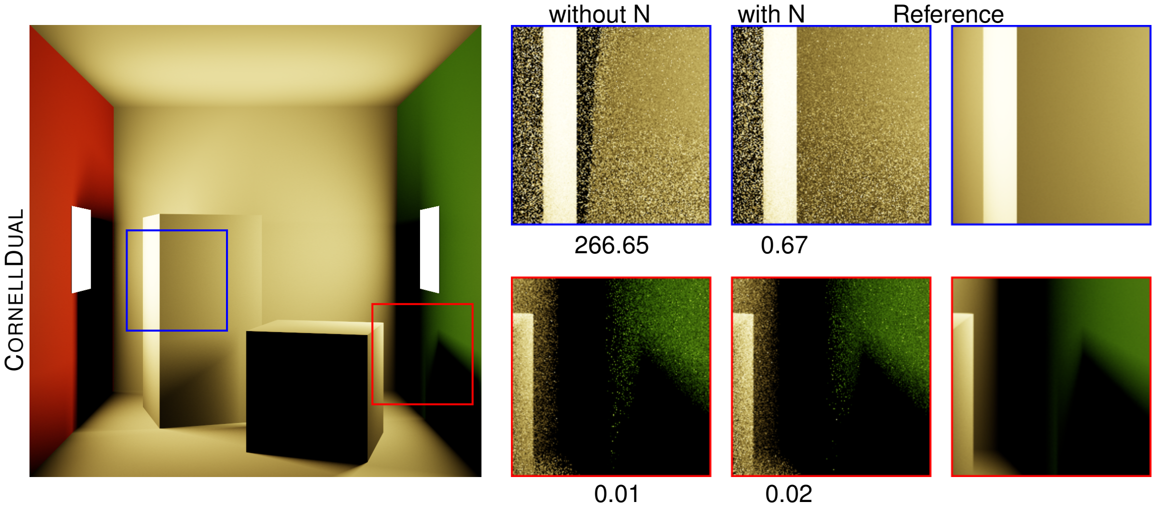

Figure: Comparison of the rendered image when using position only or position and surface normal N as input to the *guiding model*. For both renders we used 16 spp and let the MLP converge to stable state. We report the MSE for each cutout individually.

Computational cost

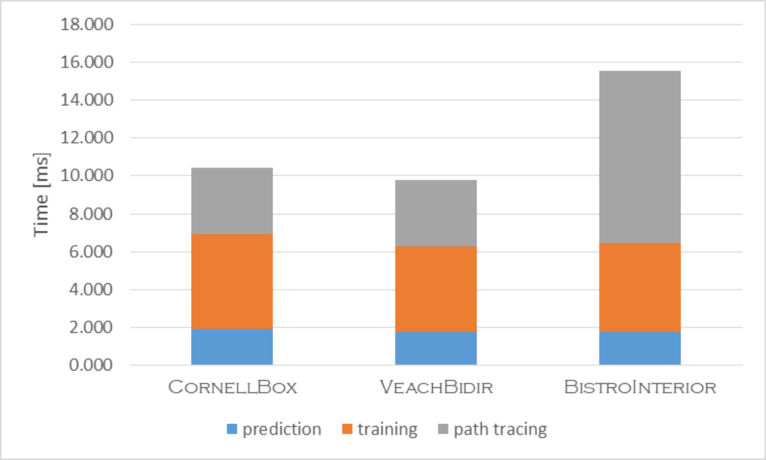

Comparison

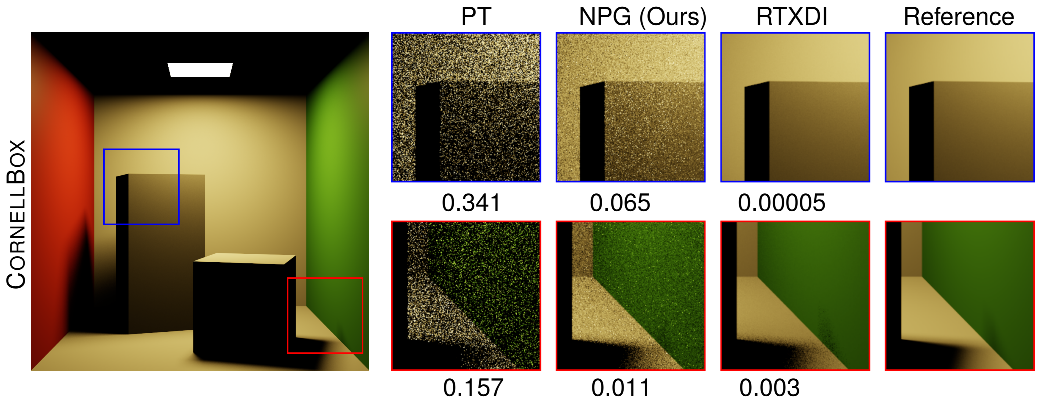

Figure: We compare our NPG to plain path tracing without NEE and to RTXDI which is the current state of the art for realtime direct illumination and report the MSE for each cutout individually. All images were rendered with 16 spp.

Comparison

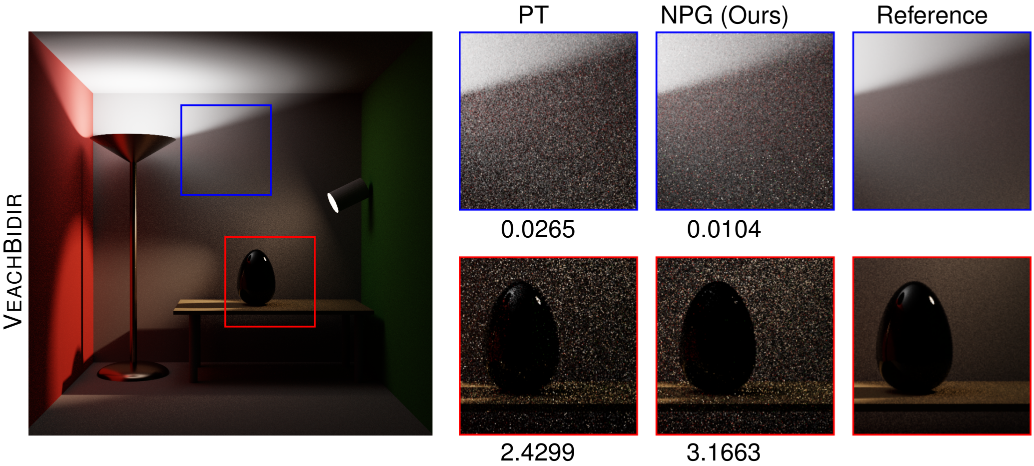

Figure: We compare our NPG to plain path tracing with NEE and BSDF importance sampling.

Limitations of the vMF distribution

CORNELLDUAL

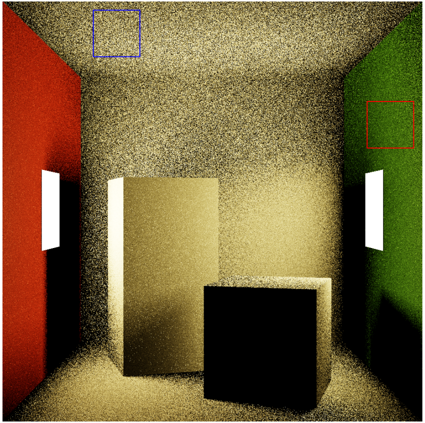

MSE NPG (MSE PT)

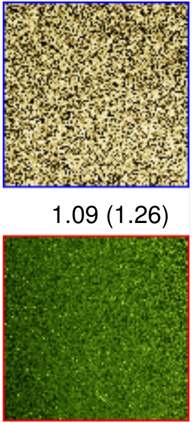

1.09 (1.26)

0.03 (0.1)

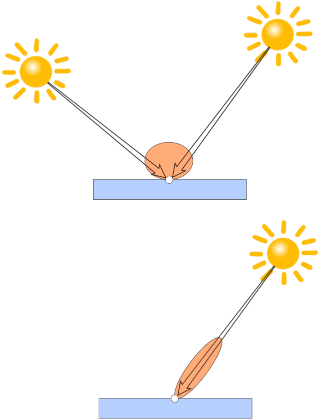

Noisy training targets

Stability issues with noisy targets

Noisy samples are the most interesting

- No path guiding needed if there is only little noise
- Noise is a result of poor importance sampling

Need to find more robust loss function or optimizer!

Conclusion

Realtime path tracing:

- Currently: ray traced reflections and shadows
- First games with path traced global illumination!
- Complex indirect illumination in realtime soon?

Real Time Neural Path Guiding

- predictable computational cost
- adapts to dynamic content
- low memory requirements

Challenge: Learn from noisy targets

Figure: Cyberpunk 2077 Path Tracing [Klotz 2023]

Literatur I

Karlsruhe Institute of Technology

[1] Sebastian Herholz et al. “Product Importance Sampling for Light Transport Path Guiding”. In: *Computer Graphics Forum* 35.4 (2016), pp. 67–77. DOI: `https://doi.org/10.1111/cgf.12950`. eprint: `https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.12950`. URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12950`.

[2] Jiawei Huang et al. “Online Neural Path Guiding with Normalized Anisotropic Spherical Gaussians”. In: *arXiv preprint arXiv:2303.08064* (2023).

[3] James T. Kajiya. “The Rendering Equation”. In: *SIGGRAPH Comput. Graph.* 20.4 (Aug. 1986), pp. 143–150. ISSN: 0097-8930. DOI: `10.1145/15886.15902`. URL: `https://doi.org/10.1145/15886.15902`.

[4] Aaron Klotz. *Cyberpunk 2077 Path Tracing Overdrive Patch Finally Available to Everyone*. https://www.tomshardware.com/news/cyberpunk-277-rt-overdrive-available-to-all. 2023.

Literatur II

Karlsruhe Institute of Technology

[5] Thomas Müller, Markus Gross, and Jan Novák. "Practical Path Guiding for Efficient Light-Transport Simulation". In: *Computer Graphics Forum* 36.4 (2017), pp. 91–100. DOI: `https://doi.org/10.1111/cgf.13227`. eprint: `https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13227`. URL: `https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13227`.

[6] Thomas Müller, Fabrice Rousselle, et al. "Real-time neural radiance caching for path tracing". In: *ACM Transactions on Graphics* 40.4 (Aug. 2021), pp. 1–16. DOI: `10.1145/3450626.3459812`. URL: `https://doi.org/10.1145%2F3450626.3459812`.

Output activations

Mean direction $\mu \in \mathcal{S}^2$:

$$\varphi(\mu_i) = \frac{2}{1 + e^{-\mu_i}} - 1 \qquad \text{for } i = 1, 2, 3$$

- normalize afterwards

Concentration $\kappa \in [0, \infty)$:

$$\varphi_{\text{quadratic}}(\kappa) = \kappa^2 \tag{4}$$

$$\varphi_{\text{exponential}}(\kappa) = e^{\kappa} \tag{5}$$

- high dynamic range of this parameter
- requires numerical stable vMF implementations

Quadratic Exponential

CornellDual

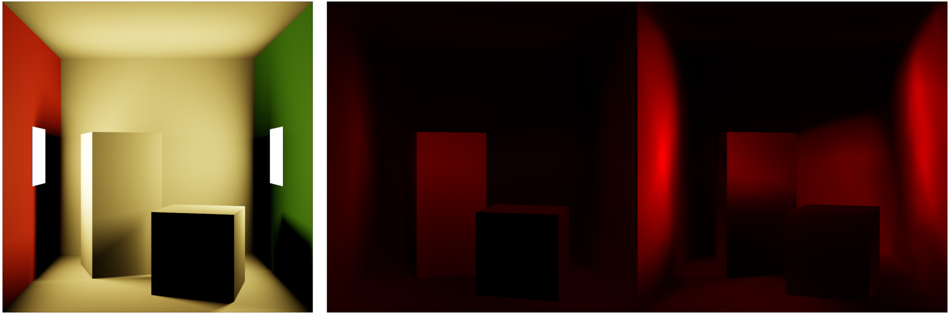

min: 0, max: 82 min: 0, max: 233

References

Limitations for specular materials

Product Importance Sampling:

- Fit vMF to BSDF (pre-computed)
- Predict vMF for radiance distribution
- Compute product of both analytically

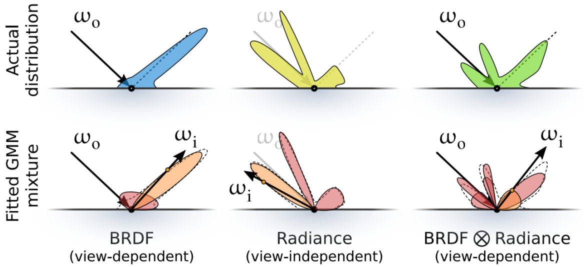

Figure: Product Importance Sampling [Herholz et al. 2016]

References