Neural Two-Level Monte Carlo Real-Time Rendering

Mikhail Dereviannykh, Dmitrii Klepikov, Johannes Hanika, Carsten Dachsbacher

Karlsruhe Institute of Technology

Project Page

















 $\int_{H^+} \frac{L_i(x,\omega_i) f_r(x,\omega_i,\omega_o) \cos \theta_i d\omega_i}{\substack{\text{Incident}\\\text{radiance}}} \underbrace{ \begin{array}{c} \text{Materials}\\ \text{BRDF} \end{array} }$















The Problem:

We can only afford 0.5-16 paths/pixel on consumers GPUs in real-time

16 paths/pixel

64 paths/pixel

1 paths/pixel

14

2048 paths/pixel





Importance sampling by guiding rays towards incident radiance

[Vorba et al. 2019, Müller et al. 2017, Müller at al. 2019, Rath et al. 2020]



Bounces are expensive They introduce additional variance



Tiny Neural Networks predict Outgoing Radiance [Müller et al. 2021]

and the other non-neural works [Greger et al. 1998, Majercik et al. 2019, Krivanek et al . 2005]













[1] Müller et al. 2023





















Project Page

Visualization of NRC

[Müller et al. 2021, Müller et al. 2022]

Ground Truth



Path Tracing NRC [Biased]

Path Tracing [Unbiased]

Path Tracer 1spp [Unbiased]

PT + NIRC (Ours), 1spp [Biased]

Equal Time Comparison

PT + NRC, 1spp [Biased]

Path Tracer 2spp [Unbiased]

PT + NIRC (Ours), 1spp [Unbiased]

Equal Time Comparison

PT + NRC, 2spp [Biased]





Neural Incident Radiance Cache (NIRC)





Neural Incident Radiance Cache (NIRC)










Project Page



NIRC Visualization



Scales pretty well with the depth of MLP (D = 2/4/6)

Two-Level Monte Carlo?

Debiasing?

(De)bias?







(De)bias?



$$\begin{split} L_{o}(x,\omega_{o}) &= L_{e}(x,\omega_{o}) + \int_{H^{+}} L_{i}(x,\omega_{i}) f_{r}(x,\omega_{i},\omega_{o}) \cos\theta_{i} d\omega_{i} \\ \hat{L}_{o}(x,\omega_{o}) &\approx L_{e}(x,\omega_{o}) + \frac{1}{N} \sum_{i=1}^{N} \frac{\prod_{i=1}^{\text{Incident}} M_{\text{aterials BRDF}}}{p(\omega_{i})} f_{r}(x,\omega_{i},\omega_{o}) \cos\theta_{i}} \end{split}$$



(De)bias?







Two-Level Monte CarloRadiance Cache
Integral Estimator
 $\begin{cases} \hat{L}_c(x, \omega_o, \hat{\theta}) \approx \frac{1}{N_c} \sum_{i=1}^{N_c} \frac{NRC}{NN_i(x, \omega_i, \hat{\theta}) f_r(x, \omega_i, \omega_o) \cos \theta_i} \\ p(\omega_i) \end{cases}$

Final biased estimator:
$$\hat{L}_o(x, \omega_o, \hat{\theta}) \approx L_e(x, \omega_o) + \hat{L}_c(x, \omega_o, \hat{\theta})$$







Variance of Two-Level Monte Carlo



$$V(\langle \hat{L}_o \rangle) = V(\langle \hat{L}_c \rangle) + V(\langle \hat{L}_r \rangle) = V(\langle \hat{L}_c \rangle) + (\hat{L}_o(x, \omega_o) - L_e(x, \omega_o) - \hat{L}_c(x, \omega_o, \hat{\theta}))^2$$

Final unbiased estimator: $\hat{L}_o(x, \omega_o, \hat{\theta}) = L_e(x, \omega_o) + \hat{L}_c(x, \omega_o, \hat{\theta}) + \hat{L}_r(x, \omega_o, \hat{\theta})$



Variance of Two-Level Monte Carlo



$$V(\langle \hat{L}_o \rangle) = V(\langle \hat{L}_c \rangle) + V(\langle \hat{L}_r \rangle) = V(\langle \hat{L}_c \rangle) + (\hat{L}_o(x, \omega_o) - L_e(x, \omega_o) - \hat{L}_c(x, \omega_o, \hat{\theta}))^2$$

The variance stays the same...

Variance of Two-Level Monte Carlo



samples to estimate L_c

Project Page















Results

1. Can be used to predict incident radiance or the visibility term





- 1. Can be used to predict incident radiance or the visibility term
- 2. Can be efficiently sampled w/o additional ray tracing and iNGP is amortized







- 1. Can be used to predict incident radiance or the visibility term
- 2. Can be efficiently sampled w/o additional ray tracing and iNGP is amortized
- 3. Decrease variance for unbiased estimator

For 1920x1080 neural samples, RTX 4080: NIRC: ~ 0.5-1.5ms NRC: ~ 5-10ms (with tracing costs)





- 1. Can be used to predict incident radiance or the visibility term
- 2. Can be efficiently sampled w/o additional ray tracing and iNGP is amortized
- 3. Decrease variance for unbiased estimator
- 4. Biased estimator can save some indirect bounces!

For 1920x1080 neural samples, RTX 4080: NIRC: ~ 0.5-1.5ms NRC: ~ 5-10ms (with tracing costs)





Saving on bounces: Bias vs Indirect Bounces





First Level Monte Carlo Control Variates GT Integrand SH NIRC SS NIRC NCV vMF 20MB + 50KB 68.2MB + 55KB 20MB + 2MB 69.1MB VRAM: 71MB $\mathbb{V}_{rel}(\langle F_r \rangle)$: 1.13 1.43 1 35 2614 3 39 (Neural) Control Variates vs Two-Level MC



PT NVC NIRC Neural Visibility Cache (NVC)



NIRC with **LODs** for the Multi-Level MC

The distribution of Neural Samples in a light path







Please check out our paper



Thanks for attention!

Biased NIRC\NRC\PT 1spp: ~35-40ms Unbiased NIRC: 65ms

Please check out our paper



Limitations



NIRC is often not accurate enough for metallic surfaces

Limitations



Lower\Higher Variance after application of 2-Level MC estimator with the NIRC

Similarities with Control Variates (CV)

CVs **limits the mathematical basis** which could be used for approximating the integrand in each point (SHs, vMFs, Normalizing flows) [Pantaleoni 2020][Müller et. al. 2020]

We **can not use** expressive MLPs with nonlinearities for CVs directly :(

analytically calculated with CVs

Final unbiased estimator:
$$\hat{L}_o(x, \omega_o, \hat{\theta}) \approx L_e(x, \omega_o) + \tilde{L}_c(x, \omega_o, \hat{\theta}) + \hat{L}_r(x, \omega_o, \hat{\theta})$$



Comparisons with Control Variates (CVs)



70

Without the sky illumination



NIRC Convergence





Neural Visibility Cache


LODs of NIRC



Distribution of Neural Samples per light path



Figure 7: We explore the effect of redistribution of the number of neural samples between the first and second vertex of a light path, denoted as N_c^1 and N_c^2 , respectively. The experiments show the Sponza and Country Kitchen scenes rendered with up to 5 vertices in a light path without Russian roulette in order to match the render time. We observe that the optimal allocation of neural samples between the first and second vertex is highly scene-dependent, influenced by, for example, material properties and illumination.