

<Special Topic in Rendering> Path Guiding

CS482 Interactive Computer Graphics

Jaeyoon Kim
Ph.D. student @ SGVR Lab
kimjy2630@kaist.ac.kr

Review

Rendering Equation


Monte Carlo Path Tracing

Rendering Equation

- Rendering equation
[Immelt et al. 1986; Kajiya 1986]

L : a radiance
 $\mathbf{x}$: a point
 ω : a direction
 Ω : the hemisphere
 $\mathbf{n}_x$: the normal vector at $\mathbf{x}$
 f_s : the BSDF (material)
 i : inward
 o : outward
 e : self-emitting

$$\underbrace{L_o(\mathbf{x}, \omega_o)}_{\text{Outgoing radiance}} = \underbrace{L_e(\mathbf{x}, \omega_o)}_{\text{Self-emitting radiance}} + \underbrace{\int_{\Omega} L_i(\mathbf{x}, \omega_i)}_{\text{Inward radiance}} \underbrace{f_s(\mathbf{x}, \omega_i, \omega_o)}_{\text{BSDF = Material}} (\omega_i \cdot \mathbf{n}_x) d\omega_i$$




Monte Carlo Ray Tracing

L : a radiance
 $\mathbf{x}$: a point
 $\boldsymbol{\omega}$: a direction
 Ω : the hemisphere
 $\mathbf{n}_{\mathbf{x}}$: the normal vector at $\mathbf{x}$
 f_s : the BSDF (material)
 i : inward
 o : outward
 e : self-emitting

- MC integration of rendering equation.

$$\bullet \langle L_o(\mathbf{x}, \boldsymbol{\omega}_o) \rangle = L_e(\mathbf{x}, \boldsymbol{\omega}_o) + \frac{1}{N} \sum_{k=1}^N \frac{\langle L_i(\mathbf{x}, \boldsymbol{\omega}_{i,k}) \rangle f_s(\mathbf{x}, \boldsymbol{\omega}_{i,k}, \boldsymbol{\omega}_o) (\boldsymbol{\omega}_{i,k} \cdot \mathbf{n}_{\mathbf{x}})}{p(\boldsymbol{\omega}_{i,k})}$$




Monte Carlo Path Tracing

L : a radiance
 $\mathbf{x}$: a point
 ω : a direction
 Ω : the hemisphere
 $\mathbf{n}_x$: the normal vector at $\mathbf{x}$
 f_s : the BSDF (material)
 i : inward
 o : outward
 e : self-emitting

- MC integration of rendering equation.

- Set $N = 1$ for intermediate bounces

- $$\langle L_o(\mathbf{x}, \omega_o) \rangle = L_e(\mathbf{x}, \omega_o) + \frac{\langle L_i(\mathbf{x}, \omega_i) \rangle f_s(\mathbf{x}, \omega_i, \omega_o) (\omega_i \cdot \mathbf{n}_x)}{p(\omega_i)}$$





Path Guiding

Variance in Path Tracing

- $\langle L_o(\mathbf{x}, \boldsymbol{\omega}_o) \rangle = L_e(\mathbf{x}, \boldsymbol{\omega}_o) + \frac{1}{N} \sum_{k=1}^N \frac{\langle L_i(\mathbf{x}, \boldsymbol{\omega}_{i,k}) \rangle f_s(\mathbf{x}, \boldsymbol{\omega}_{i,k}, \boldsymbol{\omega}_o) (\boldsymbol{\omega}_{i,k} \cdot \mathbf{n}_x)}{p(\boldsymbol{\omega}_{i,k})}$
- $$\begin{aligned} Var[\langle L_o(\mathbf{x}, \boldsymbol{\omega}_o) \rangle] &= Var \left[\frac{1}{N} \sum_{k=1}^N \frac{\langle L_i(\mathbf{x}, \boldsymbol{\omega}_{i,k}) \rangle f_s(\mathbf{x}, \boldsymbol{\omega}_{i,k}, \boldsymbol{\omega}_o) (\boldsymbol{\omega}_{i,k} \cdot \mathbf{n}_x)}{p(\boldsymbol{\omega}_{i,k})} \right] \\ &= \frac{1}{N} Var \left[\frac{\langle L_i(\mathbf{x}, \boldsymbol{\omega}_i) \rangle f_s(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) (\boldsymbol{\omega}_i \cdot \mathbf{n}_x)}{p(\boldsymbol{\omega}_i)} \right] \end{aligned}$$
- If $p \propto L_i f_s \cos \theta_i$, $Var[\langle L_o(\mathbf{x}, \boldsymbol{\omega}_o) \rangle] = 0$
 - Shoot more rays to the direction with intense light
- Path guiding: estimation for incident radiance



Path Guiding – 1D example

- MC integration for $\int_0^\pi \sin^2 x \, dx$
 - $\frac{1}{N} \sum_{n=1}^N \frac{\sin^2 x_n}{p(x_n)}$





Path Guiding – 1D example

- MC integration for $\int_0^\pi \sin^2 x \, dx$
 - Sampling pdf p
 1. uniform distribution




Path Guiding – 1D example

- MC integration for $\int_0^\pi \sin^2 x \, dx$
 - Sampling pdf p
 1. uniform distribution
 2. triangle-shaped pdf




Path Guiding – 1D example

- MC integration for $\int_0^\pi \sin^2 x \, dx$
 - Optimal pdf: $\frac{2}{\pi} \sin^2 x \propto \sin^2 x$



Path Guiding Results



Traditional Methods in Path Guiding

Grid-based


MM-based

Tree-based

Variance-aware


Grid-Based Path Guiding [Jensen 1995]

- Represent radiance fields with grid structure





GMM-Based Path Guiding [Vorba et al. 2014]

- Used Gaussian mixture model (GMM) to represent radiance fields




Tree-Based Path Guiding [Müller et al. 2017; Müller 2019]

- Used hierarchical structures
 - k -d tree for spatial subdivision
 - Each spatial leaf node contains a directional quadtree



VMM-Based Path Guiding [Ruppert et al. 2020]

- VMM-based radiance representation
- Variance-based merge/split
- Parallax-aware representation




VMM-Based Path Guiding [Ruppert et al. 2020]

- Von Mises-Fischer (vMF) distribution (spherical Gaussian)

- $\mathcal{V}(\omega|\mu, \kappa) = \frac{\kappa}{4\pi \sinh \kappa} e^{\kappa(\mu \cdot \omega)}$
 - $\mu \in S^2$
 - $\kappa \geq 0$

- vMF mixture model (VMM)

- $\mathcal{V}(\omega) = \sum_{j=1}^K \pi_j \mathcal{V}(\omega|\mu_j, \kappa_j)$
 - $\sum_{j=1}^K \pi_j = 1$




VMM-Based Path Guiding [Ruppert et al. 2020]

- Overall procedure
 - Sample rays under VMM $\mathcal{V}(\omega|\Theta) = \sum_{k=1}^K \pi_k \mathcal{V}(\omega|\mu_k, \kappa_k)$
 - $S = \{s_1, \dots, s_N\}$, $s_n = \{x_n, \omega_n, p(\omega_n|x_n), \tilde{L}_i(x_n, \omega_n)\}$
 - Compute weight for each sample
 - $w_n = \frac{1}{\tilde{\Phi}(x)} \frac{\tilde{L}_i(x_n, \omega_n)}{p(\omega_n|x_n)}$, $\tilde{\Phi}(x) = \frac{1}{N} \sum_{n=1}^N \frac{\tilde{L}_i(x_n, \omega_n)}{p(\omega_n|x_n)}$
 - Compute sufficient statistics for updating VMM parameters
 - $r_k = \sum_{n=1}^N w_n \gamma_k(\omega_n) \omega_n$, $\bar{r}_k = \frac{\|r_k\|}{\sum_{n=1}^N w_n \gamma_k(\omega_n)}$
 - Update VMM parameters from sufficient statistics
 - $\hat{\mu}_k = \frac{r_k}{\|r_k\|}$, $\hat{\kappa}_k \approx \frac{3\bar{r}_k - \bar{r}_k^3}{1 - \bar{r}_k^2}$, $\hat{\pi}_k = \frac{\sum_{n=1}^N w_n \gamma_k(\omega_n)}{\sum_{j=1}^K \sum_{n=1}^N w_n \gamma_j(\omega_n)}$

VMM-Based Path Guiding [Ruppert et al. 2020]

- Parallax-aware representation for incident radiance
 - VMM is shared in each spatial region
 - Parallax causes additional estimation error
- Later, solved in MLP-based methods



Variance-Aware Path Guiding [Rath et al. 2020]

- Recall: If $p \propto L_i f_s \cos \theta_i$, $Var[\langle L_o(\mathbf{x}, \boldsymbol{\omega}_o) \rangle] = 0$
- Can we get exact value of $L_i(\mathbf{x}, \boldsymbol{\omega}_i)$?
 - No!
- How can we compensate estimation error for $L_i(\mathbf{x}, \boldsymbol{\omega}_i)$?

Variance-Aware Path Guiding [Rath et al. 2020]

- Goal: Finding a pdf p^* minimizing variance $Var[\langle L_o(\mathbf{x}, \boldsymbol{\omega}_o) \rangle]$
- From Euler-Lagrange equation,
 - $$p^* \propto \sqrt{E[\langle L_i(\mathbf{x}, \boldsymbol{\omega}_i) \rangle^2]} f_s(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) |\boldsymbol{\omega}_i \cdot \mathbf{n}|$$
$$= \sqrt{E[\langle L_i(\mathbf{x}, \boldsymbol{\omega}_i) \rangle]^2 + Var[\langle L_i(\mathbf{x}, \boldsymbol{\omega}_i) \rangle]} f_s(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) |\boldsymbol{\omega}_i \cdot \mathbf{n}|$$

Machine Learning Approaches in Path Guiding

Normalizing flows

Offline learning


Reinforcement learning

Hierarchical CNN

Implicit representation


Neural Importance Sampling [Müller et al. 2019]

- Learns sampling functions with normalizing flows
 - Piecewise-polynomial coupling transforms




Primary Sample Space [Zheng and Zwicker 2019]

- Learn a bijective warping Ψ_k in primary sample space
 - Φ_k : a mapping from a canonical parametrization of paths with length k , Ω_k over $2(k+1)$ -dimensional hypercube $[0,1]^{2(k+1)}$




(a) Conventional approach, uniform primary sample space (PSS)



(b) Our approach, PSS importance sampling using non-linear warp


Offline Path Guiding [Bako et al. 2019]

- Scene-independent method with supervised learning
- Learns incident radiance fields from local neighbor samples
 - Can be considered as denoising for incident radiance fields




Reinforcement Learning [Huo et al. 2020]

- Define two kinds of actions
 - Refine: subdivide a node
 - Resample: double the pdf value of a node
- Reward: reduced noise after radiance field denoising




Hierarchical CNN [Zhu et al. 2021]

- Used hierarchical CNN to deal with quadtree



MLP-Based Representation [Dong et al. 2023]

- Estimate parameters of VMM through MLP
- Implicit representation helps solving parallax problem



Thank you