

# <Recent Advances in Rendering> Monte Carlo Noise Reduction

CS482 – Interactive Computer Graphics

TA: Kyu Beom Han

[qbhan@kaist.ac.kr](mailto:qbhan@kaist.ac.kr)

**SGVR** Lab




# Today's Content

- **Reviews on Monte Carlo(MC) ray tracing and MC noise**
- Image-space MC noise reduction
- Learning-based MC noise reduction


# Why Monte Carlo (Rendering) Noise Reduction?

- High complexity (2D v.s. 3D)
- Complex compatibility
  - Renderer type
  - Asset type
  - HW type

3D asset



Graphics Pipeline




Noisy Image

Clean Image



# Review - Rendering Equation

$$\underbrace{L_o(x, \vec{\omega}_o)}_{\text{Outgoing Radiance}} = \underbrace{L_e(x, \vec{\omega}_o)}_{\text{Emitting Radiance}} + \int_{\Omega} \underbrace{f_r(x, \vec{\omega}_i, \vec{\omega}_o)}_{\text{Material Property (e.g., BRDF)}} \underbrace{L_i(x, \vec{\omega}_i)}_{\text{Incoming Radiance}} (\vec{\omega}_i \cdot \vec{n}) d\vec{\omega}_i$$






# Review – MC Ray Tracing

- For fast convergence, we need to...

- Shoot more samples (Large  $N$ )
- Find a good pdf  $p(\vec{w}_i^k) \sim f_r(x, \vec{w}_i^k, \vec{w}_o) L_i(x, \vec{w}_i^k) (\vec{w}_i^k \cdot \vec{n})$

$$L_o(x, \vec{w}_o) = L_e(x, \vec{w}_o) + \int_{\Omega} f_r(x, \vec{w}_i, \vec{w}_o) L_i(x, \vec{w}_i) (\vec{w}_i \cdot \vec{n}) d\vec{w}_i$$

$$\sim L_e(x, \vec{w}_o) + \frac{1}{N} \sum_{k=1}^N \frac{f_r(x, \vec{w}_i^k, \vec{w}_o) L_i(x, \vec{w}_i^k) (\vec{w}_i^k \cdot \vec{n})}{p(\vec{w}_i^k)}$$



# Review – MC Ray Tracing and MC Noise

- Shooting few samples per pixel (spp) leads to noisy radiance estimation




2.5s



8.9s




18m



# Review - Metropolis Light Transport (MLT)


- Using advanced sampling technique (Metropolis-Hasting algorithm) to generate valid (important) samples.
- Beneficial for scenes with complex geometry and indirect lighting.






# Review - Bidirectional Path Tracing (BDPT)

- Combining rays traced from the camera and light sources
- Beneficial for scenes with complex geometry and indirect lighting



## Review - Irradiance Caching


- Caching irradiance (and its gradient) of the points visible from camera
- Intuition: Indirect lighting is mostly smooth  $\rightarrow$  Sparse computation is enough





# Review - Irradiance Caching

- Caching irradiance (and its gradient) of the points visible from camera
- Intuition: Indirect lighting is mostly smooth → Sparse computation is enough

Irradiance Caching  
(Constant Extrapolation)



Irradiance Caching + Gradients  
(Linear Extrapolation)



--- actual irradiance      — extrapolated irradiance      ● irradiance cache point

## Review - Photon Mapping

- Shoot photons from the light source and save information (energy, position, direction, etc.) (a)
- Use K-nearest photons for estimating the radiance of the query point (b)




Figure 3: The Museum scene



Figure 4: Direct visualization of the global photon map in the Museum scene



# Review - Photon Mapping

- Shoot photons from the light source and save information (energy, position, direction, etc.) (a)
- Use K-nearest photons for estimating the radiance of the query point (b)




Figure 3: The Museum scene



Figure 4: Direct visualization of the global photon map in the Museum scene




# Content

- Reviews on Monte Carlo(MC) ray tracing and MC noise
- **Image-space MC noise reduction**
- Learning-based MC noise reduction


# Image-space MC Noise Reduction

- Efficiently dealing noise on image-space, similar to general image denoising
- Reducing working space from N-dim path-space to 2-dim image space




# General Image Denoising Algorithms for MC Rendering


- Efficiently dealing noise on image-space, similar to general image denoising
- Reducing working space from N-dim path-space to 2-dim image space



Non-local Means Filter



Bilateral Filter



Block-Matching 3D (BM3D)

# General Image Denoising Algorithms for MC Rendering


- Efficiently dealing noise on image-space, similar to general image denoising
- Reducing working space from N-dim path-space to 2-dim image space
- Filter weights determined based on similarity in RGB, G-buffers



RGB



Albedo



Normal



Depth

$$w_{ij} = \exp\left[-\frac{1}{2\sigma_p^2} \sum_{1 \leq k \leq 2} (\bar{p}_{i,k} - \bar{p}_{j,k})^2\right] \times \text{Pixel position}$$

$$\exp\left[-\frac{1}{2\sigma_c^2} \sum_{1 \leq k \leq 3} \alpha_k (\bar{c}_{i,k} - \bar{c}_{j,k})^2\right] \times \text{RGB}$$

$$\exp\left[-\frac{1}{2\sigma_f^2} \sum_{1 \leq k \leq m} \beta_k (\bar{f}_{i,k} - \bar{f}_{j,k})^2\right], \text{ G-buffers (Albedo, normal, depth, etc.)}$$


# Content

- Reviews on Monte Carlo(MC) ray tracing and MC noise
- Path-space MC noise reduction
- Image-space MC noise reduction
- **Learning-based MC noise reduction**
  - Image-space
  - Sample-space
  - Path Guiding
  - Post-post processing




# Deep-learning Era for Image-space Denoising

- Various neural networks (MLP, ConvNets, Transformers, etc.) and training strategies (supervised, self-supervised, unsupervised, etc.) are introduced during the last decade
- Reduce design biases of traditional denoising filters




# Conventional Configuration for Learning-based Methods

- Training a neural network to predict the clean image based on the input noisy image and auxiliary features (e.g., G-buffers)




# Deep-learning for Image-space MC Noise Reduction

- Estimating parameters from cross-bilateral filters using MLP and a large dataset
  - Input : G-buffers, world position, visibility, mean/standard/mean deviation, gradients, spp




# Predicting Kernel Weights using CNN

- Robust training by training the network to predict the denoising kernels (KPCN) instead of denoised pixel value (DPCN)
  - Reduces the search space (pixel radiance : 0 ~ unlimited, kernel weights: 0~1)




(a) Diffuse



(b) Specular





# Decompose to Diffuse and Specular

- Train each denoising CNNs to deal with separate lighting effects
  - Diffuse: Geometry dependent, Smooth & low range
  - Specular: View dependent, High range






# Kernel-predicting Convolutional Network (KPCN)

| Ours                                                                               | Input (32 spp)                                                                      | RDFC (log)                                                                          | APR (log)                                                                             | NFOR (log)                                                                            | LBF-RF (log)                                                                          | Ours                                                                                  | Ref. (1K-4K spp)                                                                      |
|------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
|    |    |    |    |    |    |    |    |
| relative $\ell_2$<br>1 - SSIM                                                      | 19.21e-3<br>0.354                                                                   | 1.67e-3<br>0.043                                                                    | 2.66e-3<br>0.058                                                                      | 1.29e-3<br>0.034                                                                      | 2.15e-3<br>0.051                                                                      | 1.16e-3<br>0.032                                                                      |                                                                                       |
|    |    |    |    |    |    |    |    |
| relative $\ell_2$<br>1 - SSIM                                                      | 18.88e-3<br>0.271                                                                   | 1.54e-3<br>0.026                                                                    | 1.95e-3<br>0.028                                                                      | 1.24e-3<br>0.019                                                                      | 2.67e-3<br>0.038                                                                      | 0.93e-3<br>0.016                                                                      |                                                                                       |
|  |  |  |  |  |  |  |  |
| relative $\ell_2$<br>1 - SSIM                                                      | 9.28e-3<br>0.090                                                                    | 2.44e-3<br>0.023                                                                    | 3.35e-3<br>0.030                                                                      | 2.12e-3<br>0.019                                                                      | 4.69e-3<br>0.027                                                                      | 2.16e-3<br>0.019                                                                      |                                                                                       |

# Adversarial Training for Direct Pixel Denoising (AdvMCD)


- Jointly train the denoising networks and critic networks
- The critic networks are trained to guess whether the input image is clean or noisy (denoised)
- Denoising networks are trained to fool the critic network





# Feature-guided Self-attention (AFGSA)


- Stack multiple transformer blocks that creates self-attention map from input image and auxiliary features



Adversarial Monte Carlo Denoising with Conditioned Auxiliary Feature Modulation, Yu et al., ToG 2021

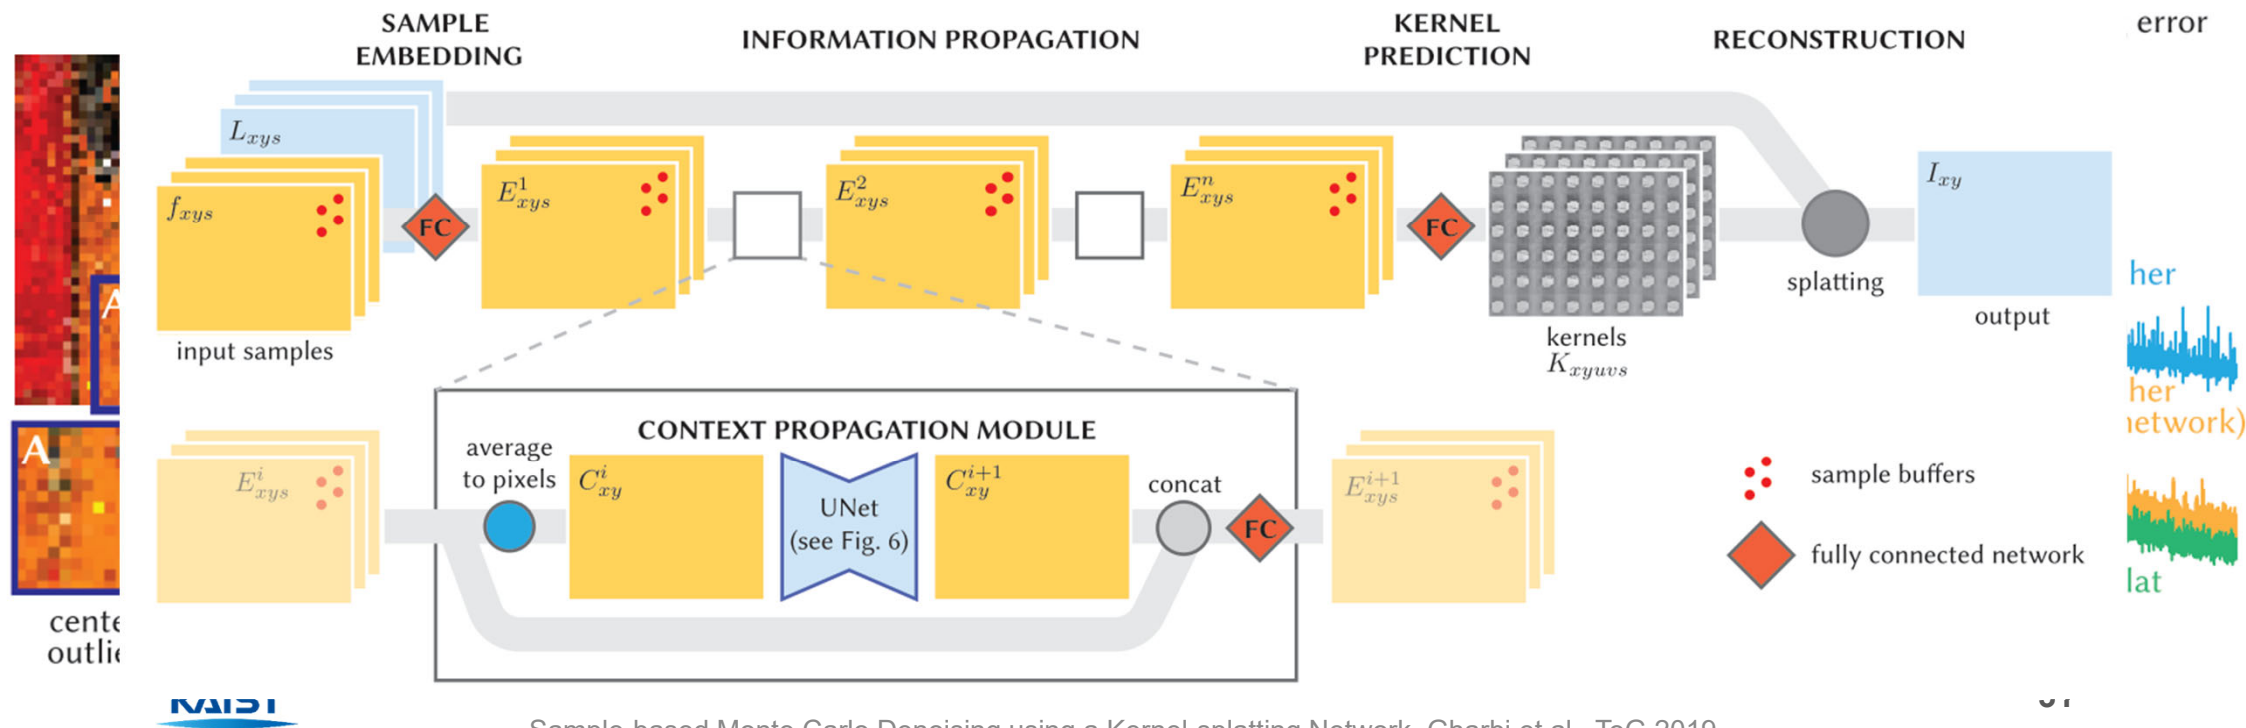
# Jumping from Image-space to Sample-space

- Ray tracing allows to naturally generate blurring effects
- How to reduce the noise while preserving these effects?




# Splatting Kernel for Samples

- Conventional kernels: Gathers nearby pixels (samples) with assigned weights
  - Denoised Pixel : Is the  $i$ \_th sample of my  $j$ \_th neighbor an outlier?
- Splatting Kernels: Pixels (samples) contributes to nearby pixels with assigned weights
  - Noisy Pixel (sample) : Am I an outlier to my  $j$ \_th neighbor?
- Intuitive & permutation invariant






# Sample-based Monte Carlo Denoising (SBMC)




# Path-space Features for Denoising

- Multi-bounce features are useful for reconstructing complex lighting details
- High-dimensionality harms the training of neural network

- [Gharbi 2019; Lin 2021]




## Reconstruction Network




# Manifold Learning for Path-space Features

- Embed path features to low-dimensional space




# Manifold Learning for Path-space Features

- Use pixel colors as pseudo-labels
- Embed path features based on pixel-color similarity using contrastive learning





# Manifold Learning for Path-space Features






# Neural Radiance Caching

- Solving rendering equation via **Radiance-predicting Neural Network  $L_\theta$**

$$L_o(x, \vec{w}_o) = L_e(x, \vec{w}_o) + \int_{\Omega} f_r(x, \vec{w}_i, \vec{w}_o) L_i(x, \vec{w}_i) (\vec{w}_i \cdot \vec{n}) d\vec{w}_i$$

$$\sim L_e(x, \vec{w}_o) + \mathbf{L}_\theta(x, \vec{w}_o)$$




# Neural Radiance Caching

Train the neural network → Cache, Estimate the radiance → Interpolate


$$L_o(x, \vec{w}_o) = L_e(x, \vec{w}_o) + \int_{\Omega} f_r(x, \vec{w}_i, \vec{w}_o) L_i(x, \vec{w}_i) (\vec{w}_i \cdot \vec{n}) d\vec{w}_i$$

$$\sim L_e(x, \vec{w}_o) + L_{\theta}(x, \vec{w}_o)$$




# Self-training for Neural Radiance Cache

- Minimize the loss between the **calculated radiances** and the **estimated radiances** of the preceding vertices
- $Loss = relL2(L_1, L_\theta(y_1, \omega_1)) + relL2(L_2, L_\theta(y_2, \omega_2)) + relL2(L_3, L_\theta(y_3, \omega_3))$
- Trace a short rendering path ( $x_0x_1x_2$ ) where we used the cached(estimated) radiance in vertex  $x_2$  for rendering



# Post-processing the Denoiser (Post-post Processing)

- Denoising models trained on certain noise level is biased to the noise level & dataset
- Cannot show consistent performance throughout noise levels



(a) DRAGON



# Combining Biased and Unbiased Estimates

- Path Traced Result  $X$  : Noisy but Unbiased (Bias  $\downarrow$ , Variance  $\uparrow$ )
- Denoised Result  $Y$ : Smooth but Biased (Bias  $\uparrow$ , Variance  $\downarrow$ )
- James-Stein Estimator shrinks  $X$  towards  $Y$  as  $\delta(X, Y) = Y + \left(1 - \frac{(p-2)\sigma^2}{\|X-Y\|^2}\right) (X - Y)$ 
  - $p$  : Dimension of estimation (3 = RGB channel),  $\sigma$ : variance of radiance
- Performs better than sample mean (in our case,  $X$ ) if  $p \geq 3$


$$MSE = (BIAS)^2 + VARIANCE$$

Leaving some space on BIAS, James-Stein Estimator reduces the VARIANCE by shrinking the points to be dense


# Combining Biased and Unbiased Estimates

- Path Traced Result  $X$  : Noisy but Unbiased (Bias  $\downarrow$ , Variance  $\uparrow$ )
- Denoised Result  $Y$ : Smooth but Biased (Bias  $\uparrow$ , Variance  $\downarrow$ )
- James-Stein Estimator shrinks  $X$  towards  $Y$  as  $\delta(X, Y) = Y + \left(1 - \frac{(p-2)\sigma^2}{\|X-Y\|^2}\right) (X - Y)$ 
  - $p$  : Dimension of estimation (3 = RGB channel),  $\sigma$ : variance of radiance
- Performs better than sample mean (in our case,  $X$ ) if  $p \geq 3$


Radius = 0.5



Radius = 1.0



Radius = 2.0




James-Stein Estimator shows less MSE error

Grey – Sampled points on radius sphere with center (1, 1, 1)

Red – James-Stein estimator applied on sampled points

# Combining Biased and Unbiased Estimates

- Path Traced Result  $X$  : Noisy but Unbiased (Bias  $\downarrow$ , Variance  $\uparrow$ )
- Denoised Result  $Y$ : Smooth but Biased (Bias  $\uparrow$ , Variance  $\downarrow$ )
- James-Stein Estimator shrinks  $X$  towards  $Y$  as  $\delta(X, Y) = Y + \left(1 - \frac{(p-2)\sigma^2}{\|X-Y\|^2}\right) (X - Y)$ 
  - $p$  : Dimension of estimation (3 = RGB channel),  $\sigma$ : variance of radiance
- Performs better than sample mean (in our case,  $X$ ) if  $p \geq 3$




<Intuition>


Balancing the bias and variance of between path traced result and denoised result

# Neural James-Stein Combiner

- Small U-Net to estimate weights for James-Stein Combiner




# Neural James-Stein Combiner





# Neural James-Stein Combiner



# What We Covered

- Image-space MC noise reduction
- Learning-based MC noise reduction
  - Image-space methods
  - Sample- & Path-space methods
  - Post-post processing