

<Recent Advances in Rendering> Monte Carlo Noise Reduction

CS380 – Introduction to Computer Graphics

TA: Kyu Beom Han

qbhan@kaist.ac.kr

SGVR Lab



Today's 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

Review - Rendering Equation

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

Outgoing Radiance ← $L_o(x, \vec{\omega}_o)$ ← Emitting Radiance $L_e(x, \vec{\omega}_o)$

Material Property (e.g., BRDF) ← $f_r(x, \vec{\omega}_i, \vec{\omega}_o)$ ← Incoming Radiance $L_i(x, \vec{\omega}_i)$



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



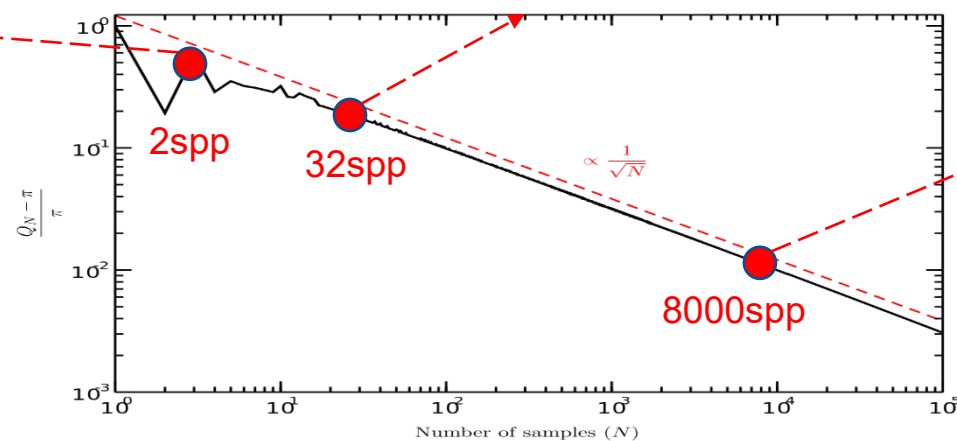
2.5s



8.9s



18m




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


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




(a) Light tracing




(b) Path tracing
▼
eye point



(c) PT with DI




(d) BDPT
↔
light source



(e) BDPT with MIS

Review - Irradiance Caching

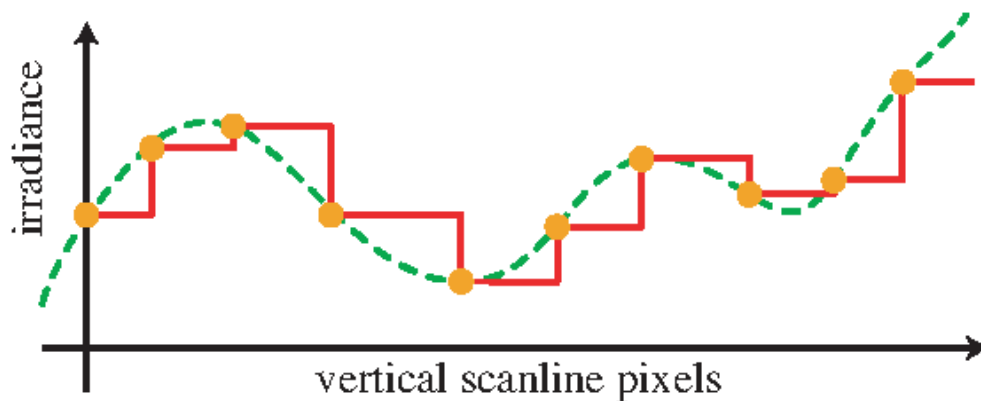
- Caching irradiance (and its gradient) of the points visible from camera
- Intuition: Indirect lighting is mostly smooth \rightarrow Sparse computation is enough



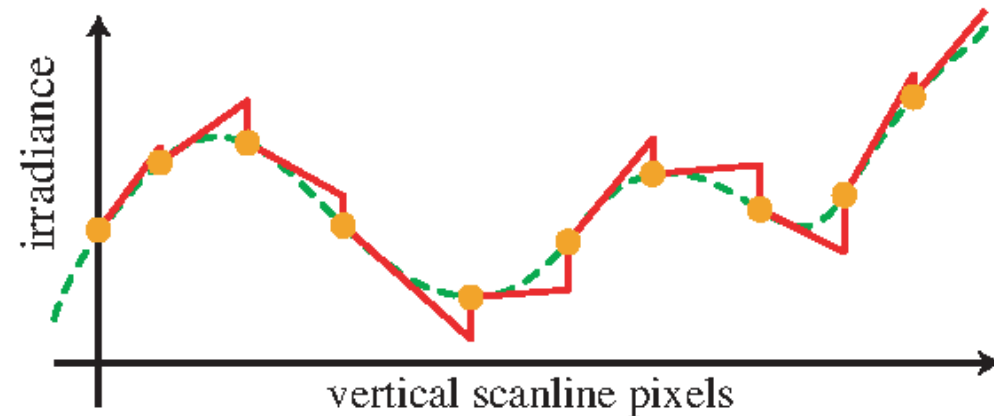
Review - Irradiance Caching

- 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)





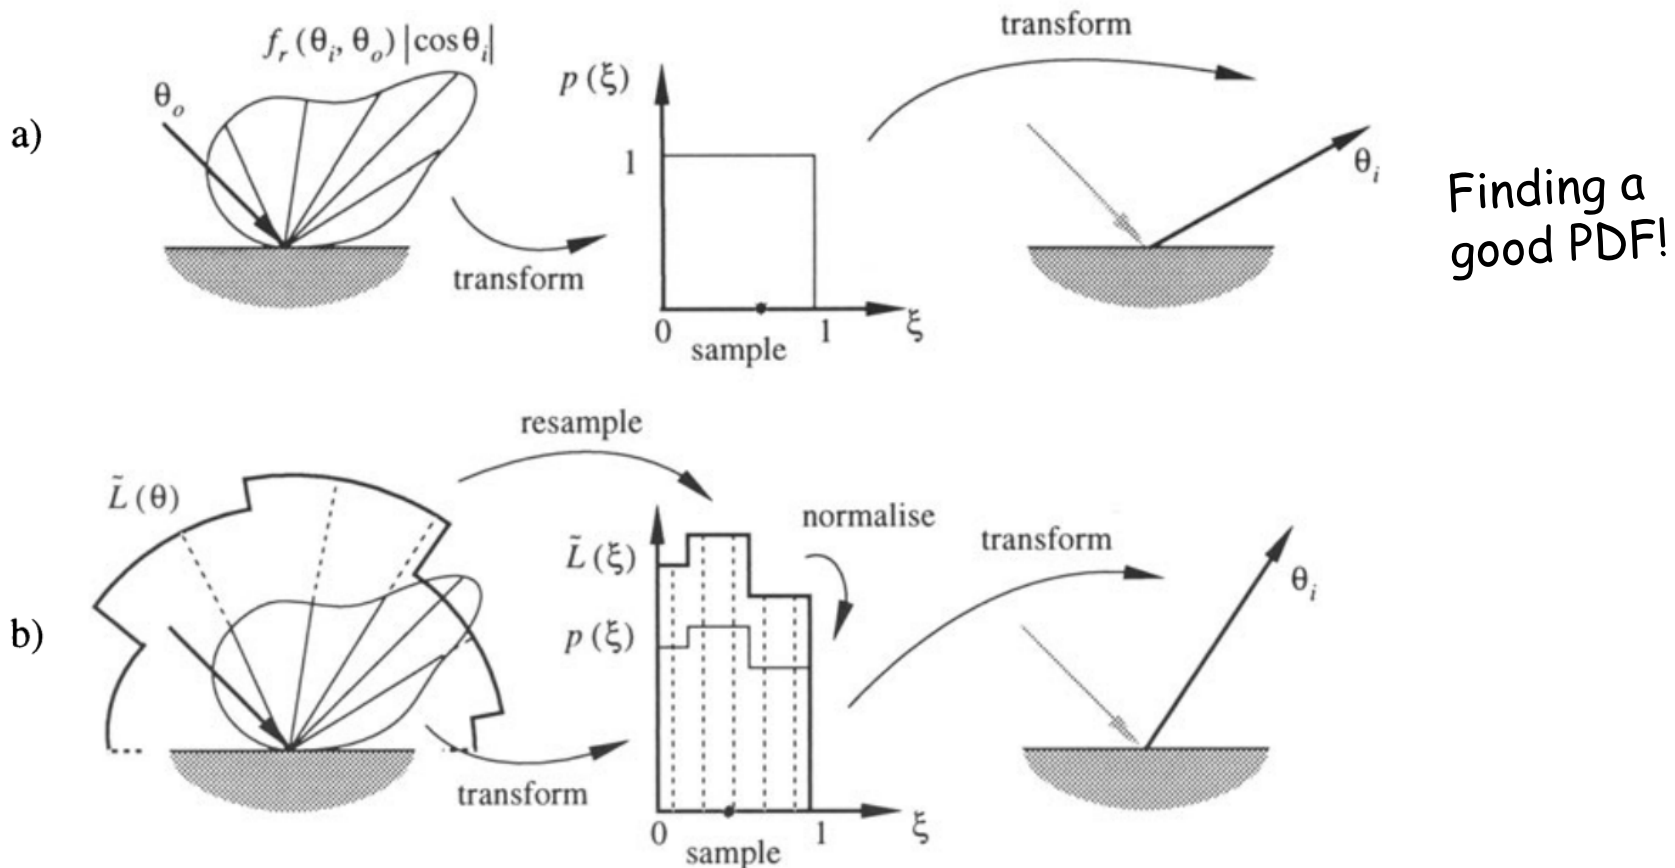
Figure 3: The Museum scene



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

Path Guiding

- Guiding samples to certain position or direction
- Giving higher probability to position/direction with higher radiance (or any other metric)



Path Guiding

- PDFs stored on various grid- or tree-like structures
 - PDF as 2D map (θ, ϕ)


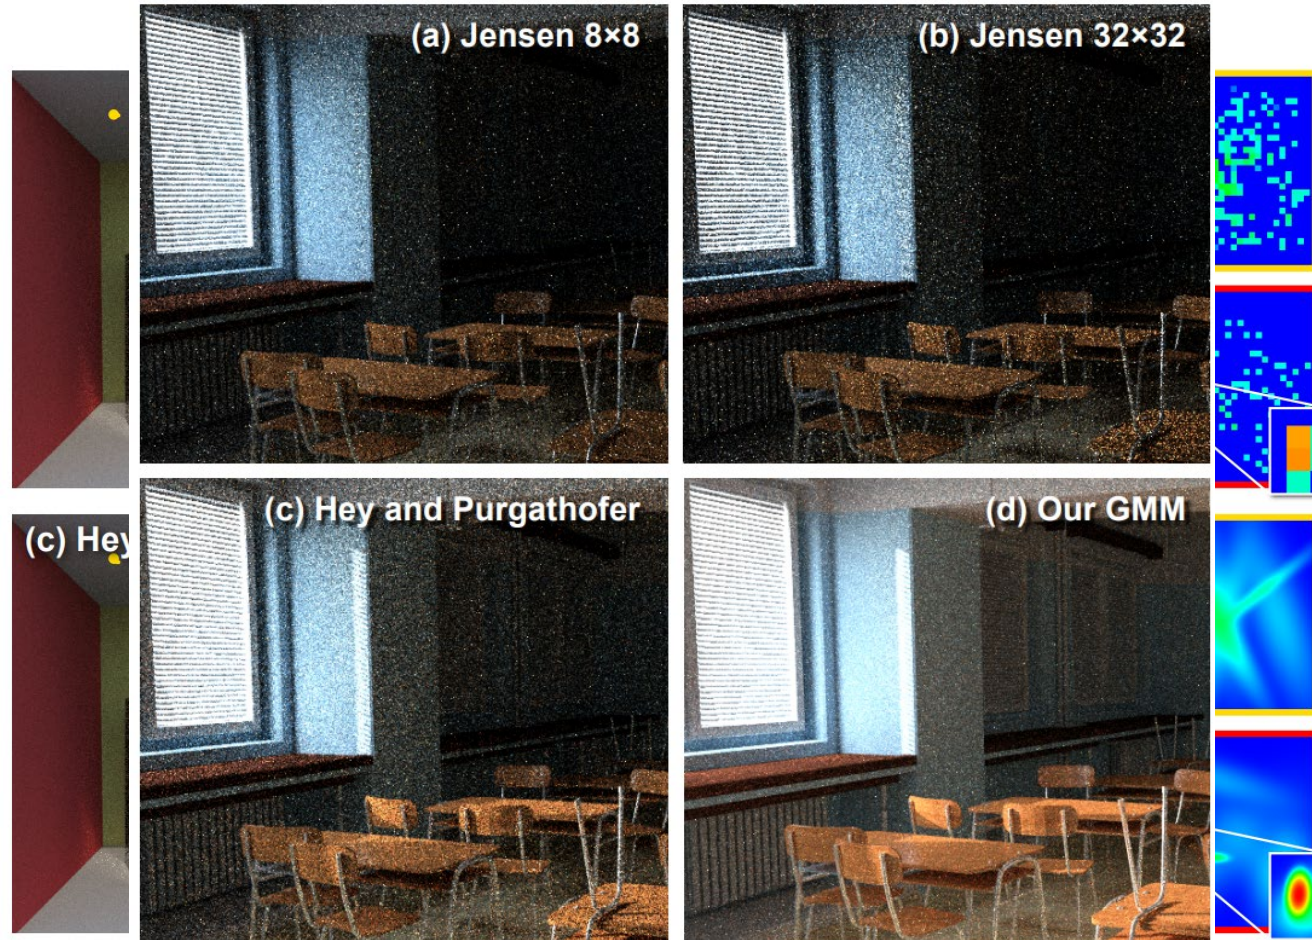


Figure 2: Color-mapped scalar values of predicate p at 2 different refinement stages.


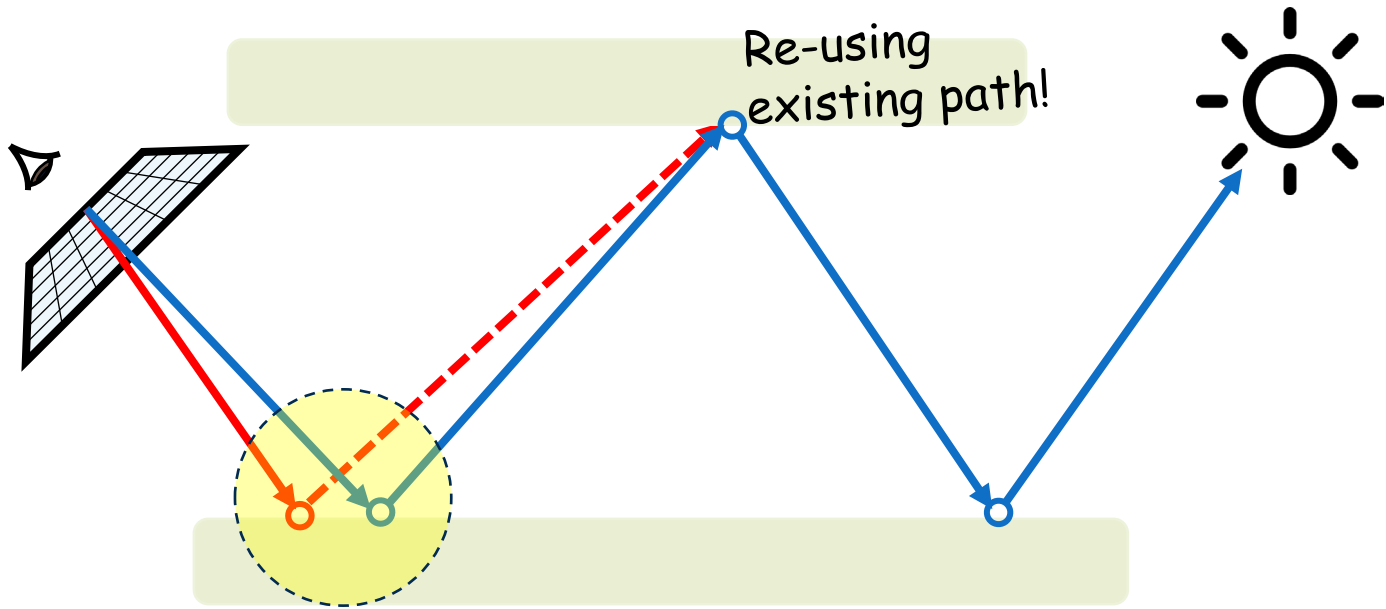
Path Guiding

- PDFs stored on various grid- or tree-like structures
 - PDF as 2D map (θ, ϕ)
 - Advanced structures (e.g., mixture models) also possible



Path-space Filtering (or Path Reuse)

- Estimating an appropriate denoising filter (kernel) to be applied on each bounce of samples
 - $\bar{L}(y_k) = \sum_{j=1}^n L(x_j) \cdot \frac{w(x_j)}{p(x_j)}, x_i \in B(y_k)$
 - $\bar{L}(\cdot)$: prefiltered radiance, $L(\cdot)$: prefiltered radiance, $p(\cdot)$: probability of sampled vertex, $w(\cdot)$: weight calculated by balance heuristic




Path-space Filtering (or Path Reuse)

- Estimating an appropriate denoising filter (kernel) to be applied on each bounce of samples
- Can involve various indirect illumination (dashed lines)

Clean Image

15 spp w/o reuse

15 spp w/ reuse




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 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



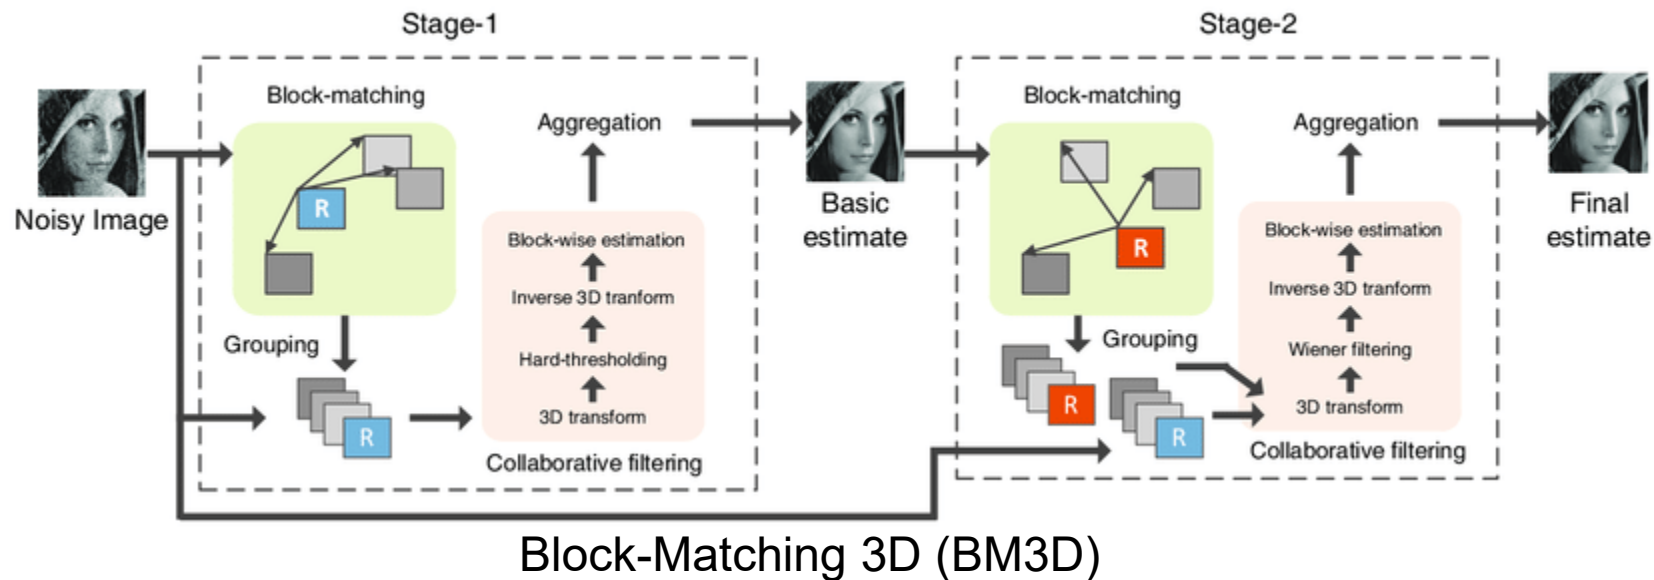
General Image Denoising Algorithms for MC Rendering

- 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

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

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

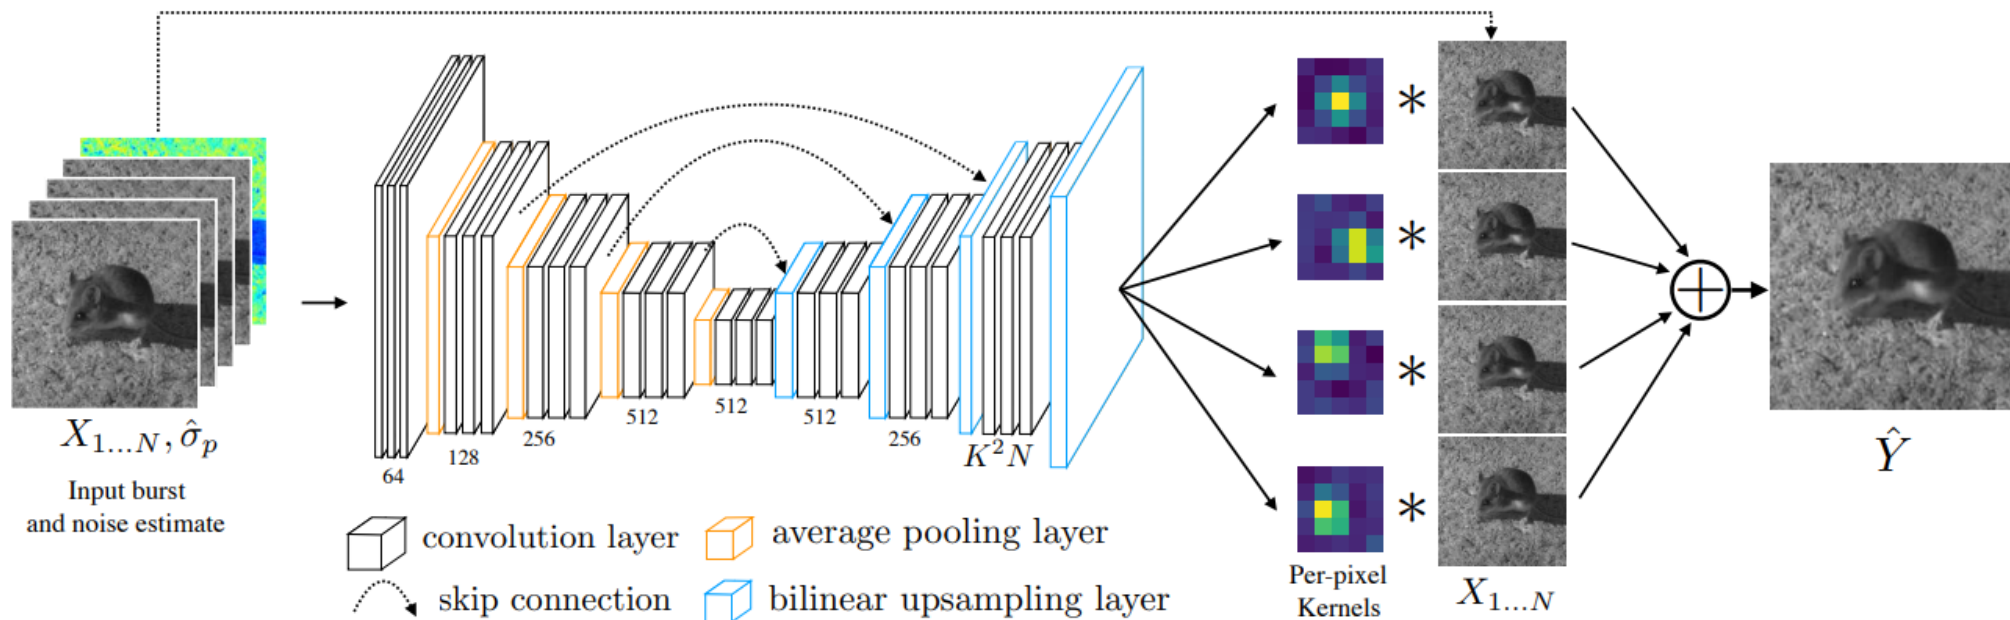
$$\exp\left[-\frac{1}{2\sigma_f^2} \sum_{1 \leq k \leq m} \beta_k (\bar{f}_{i,k} - \bar{f}_{j,k})^2\right], \text{ G-buffers (Albedo, normal, depth, etc.)}$$

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

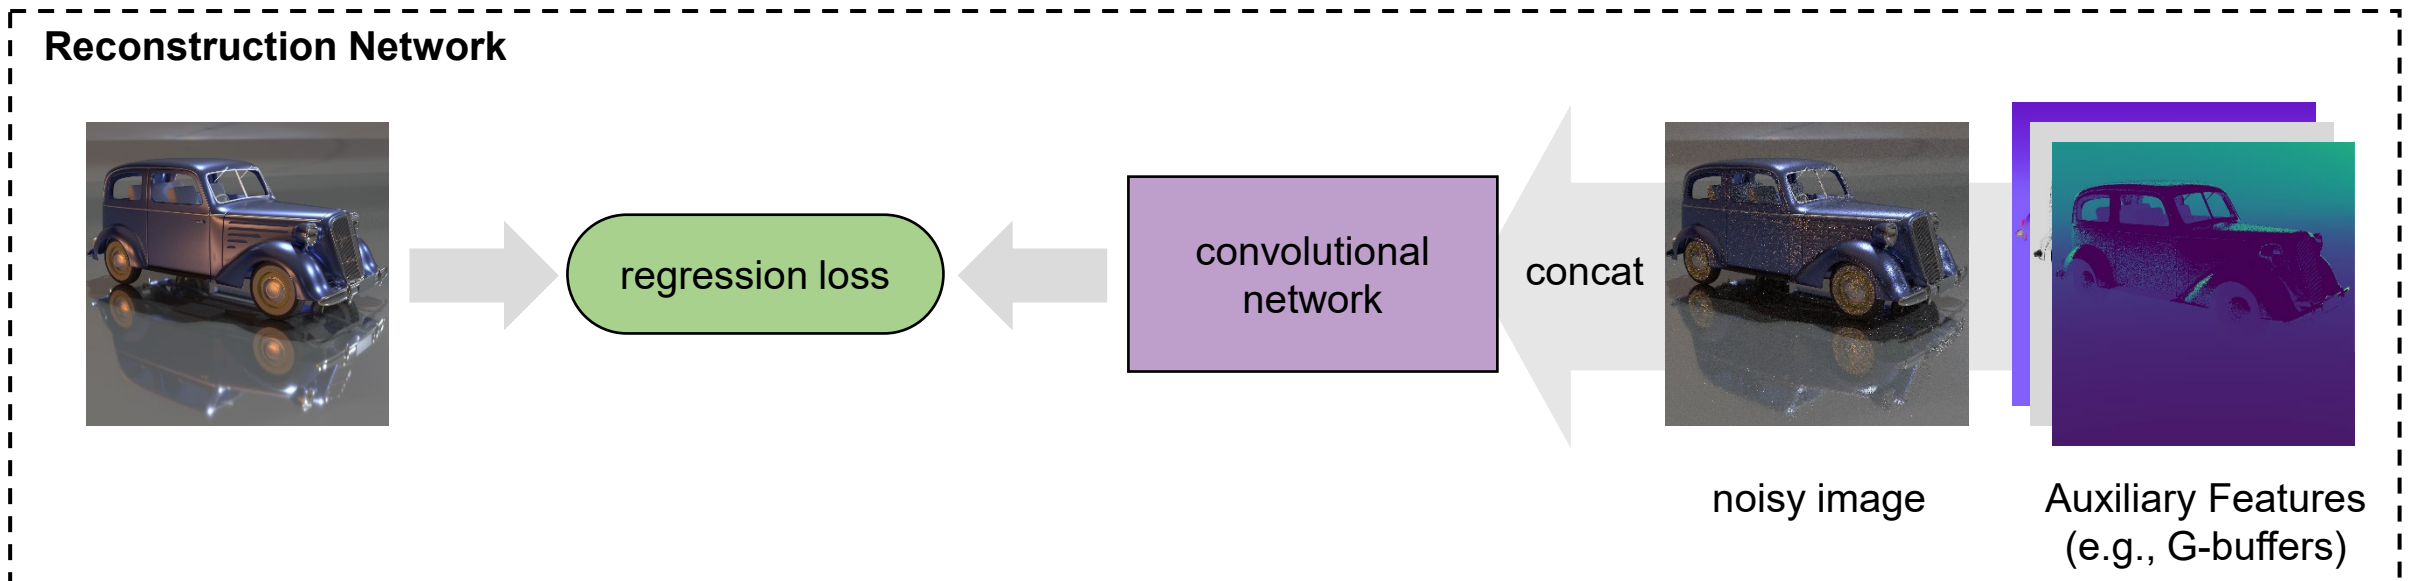
Deep-learning Era for Image-space Denoising

- 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)



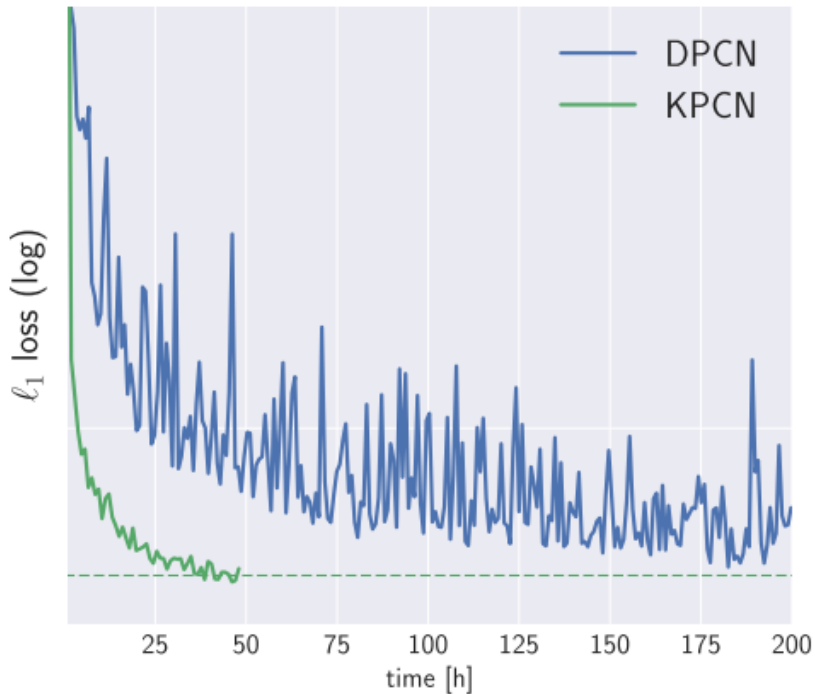
Deep-learning for Image-space MC Noise Reduction

- 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)



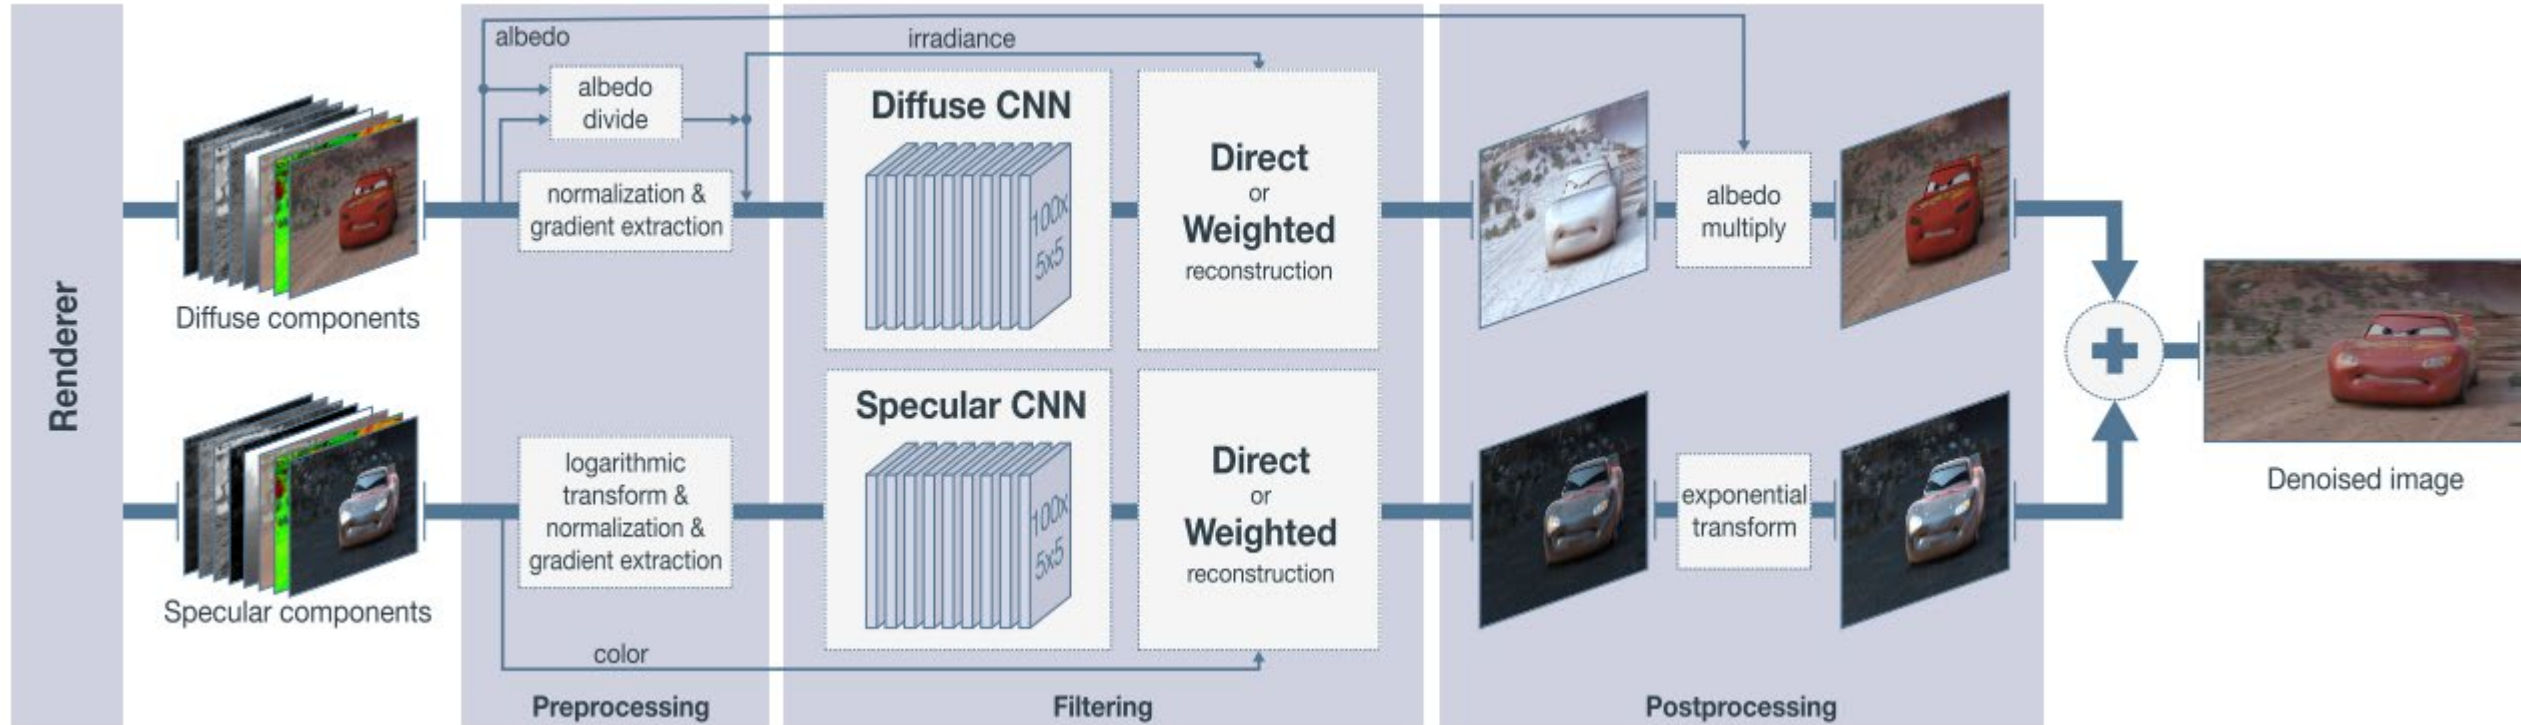
(a) Diffuse



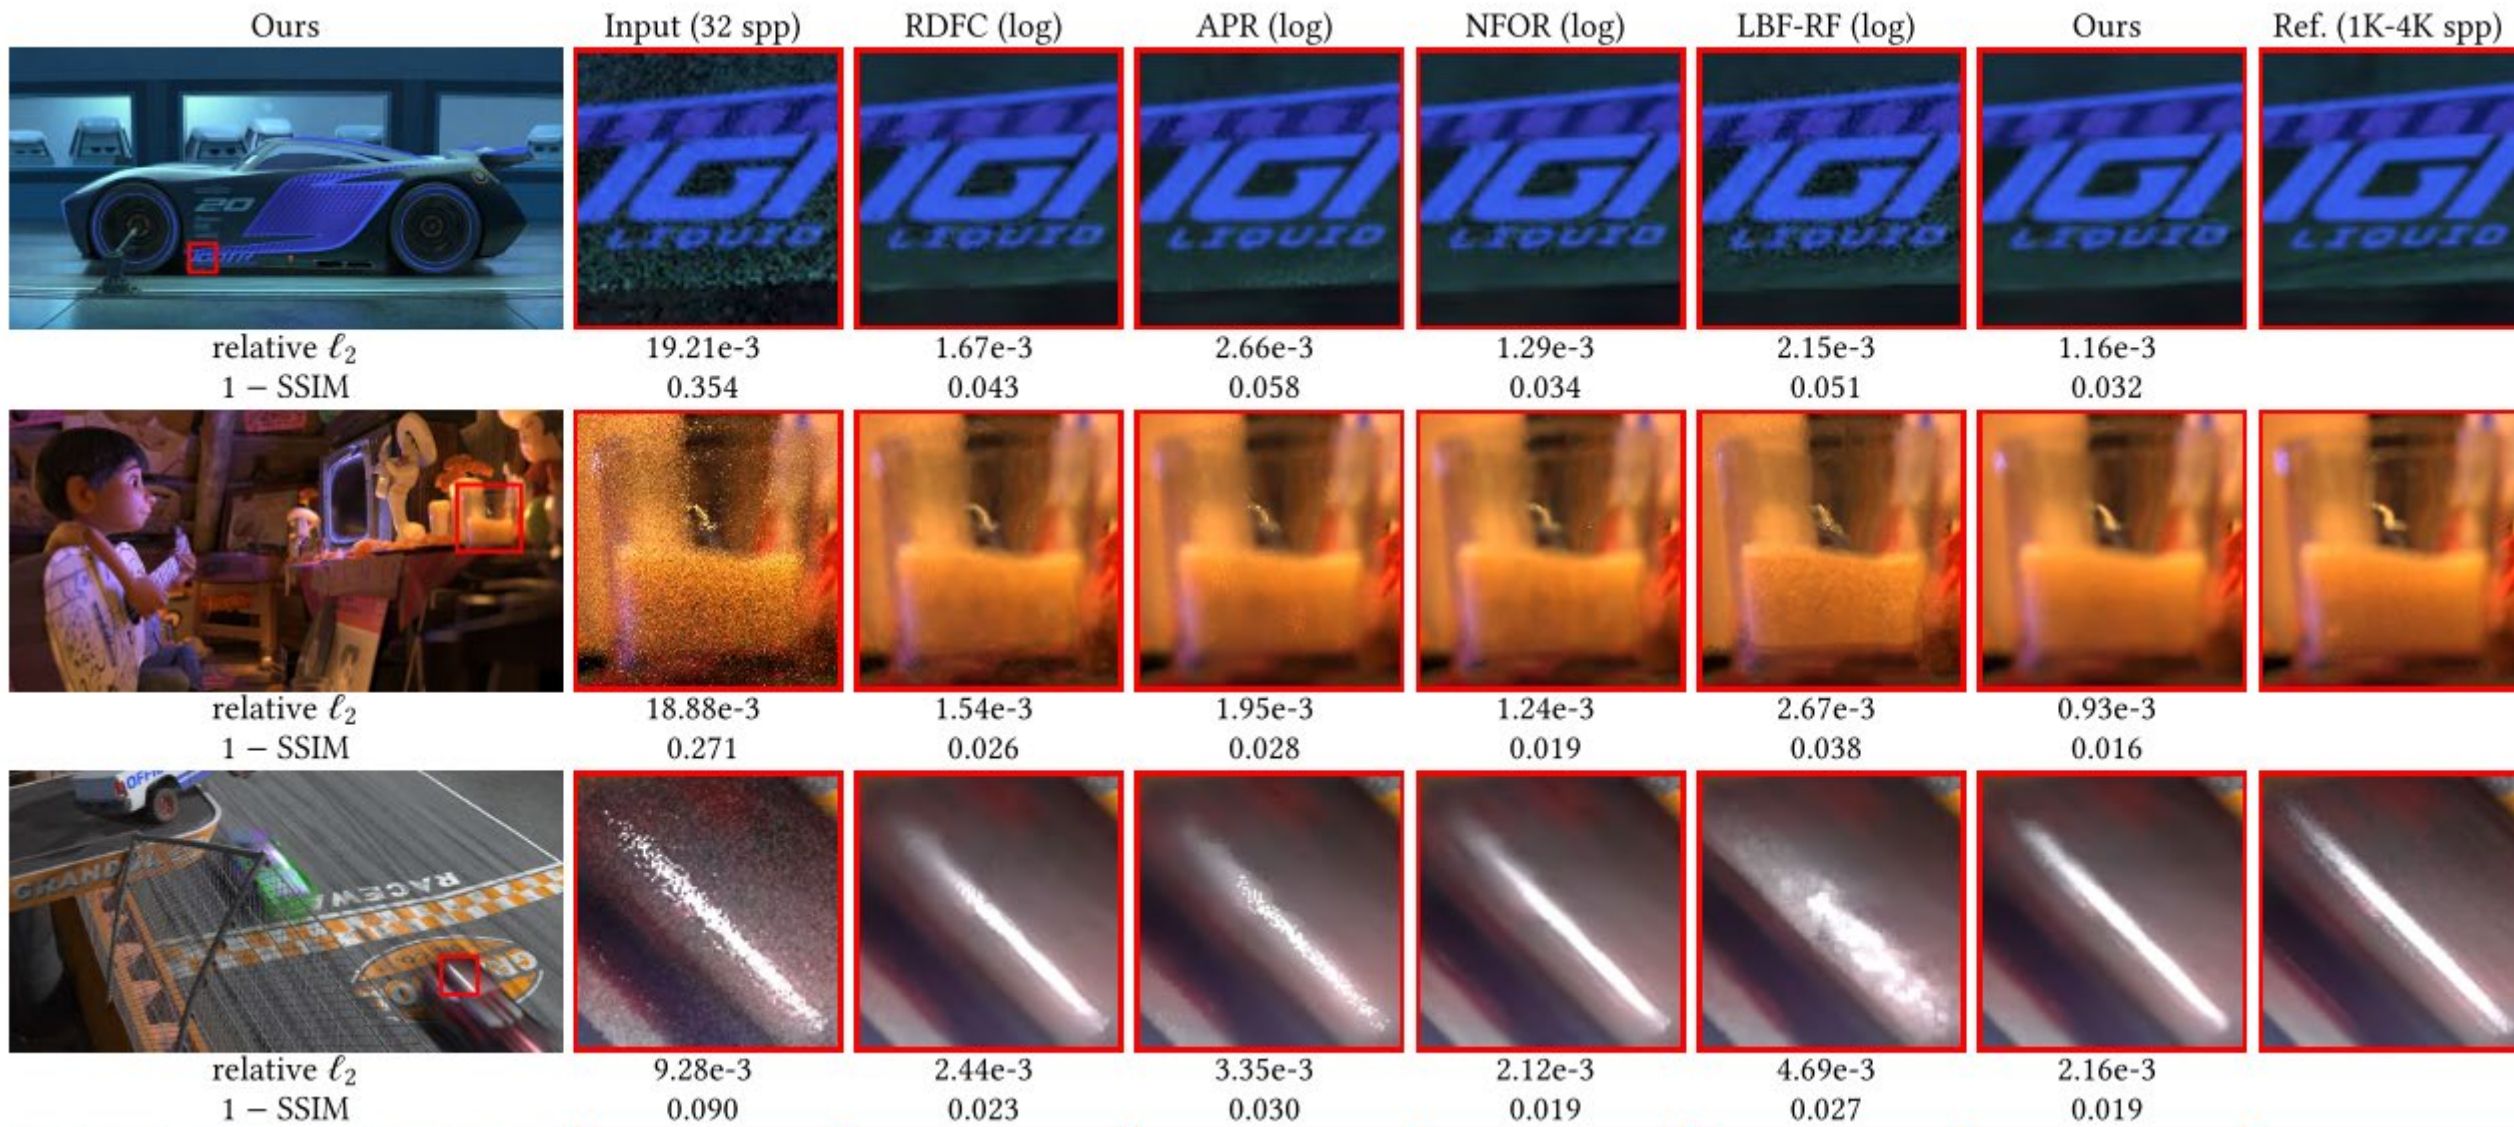
(b) Specular

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

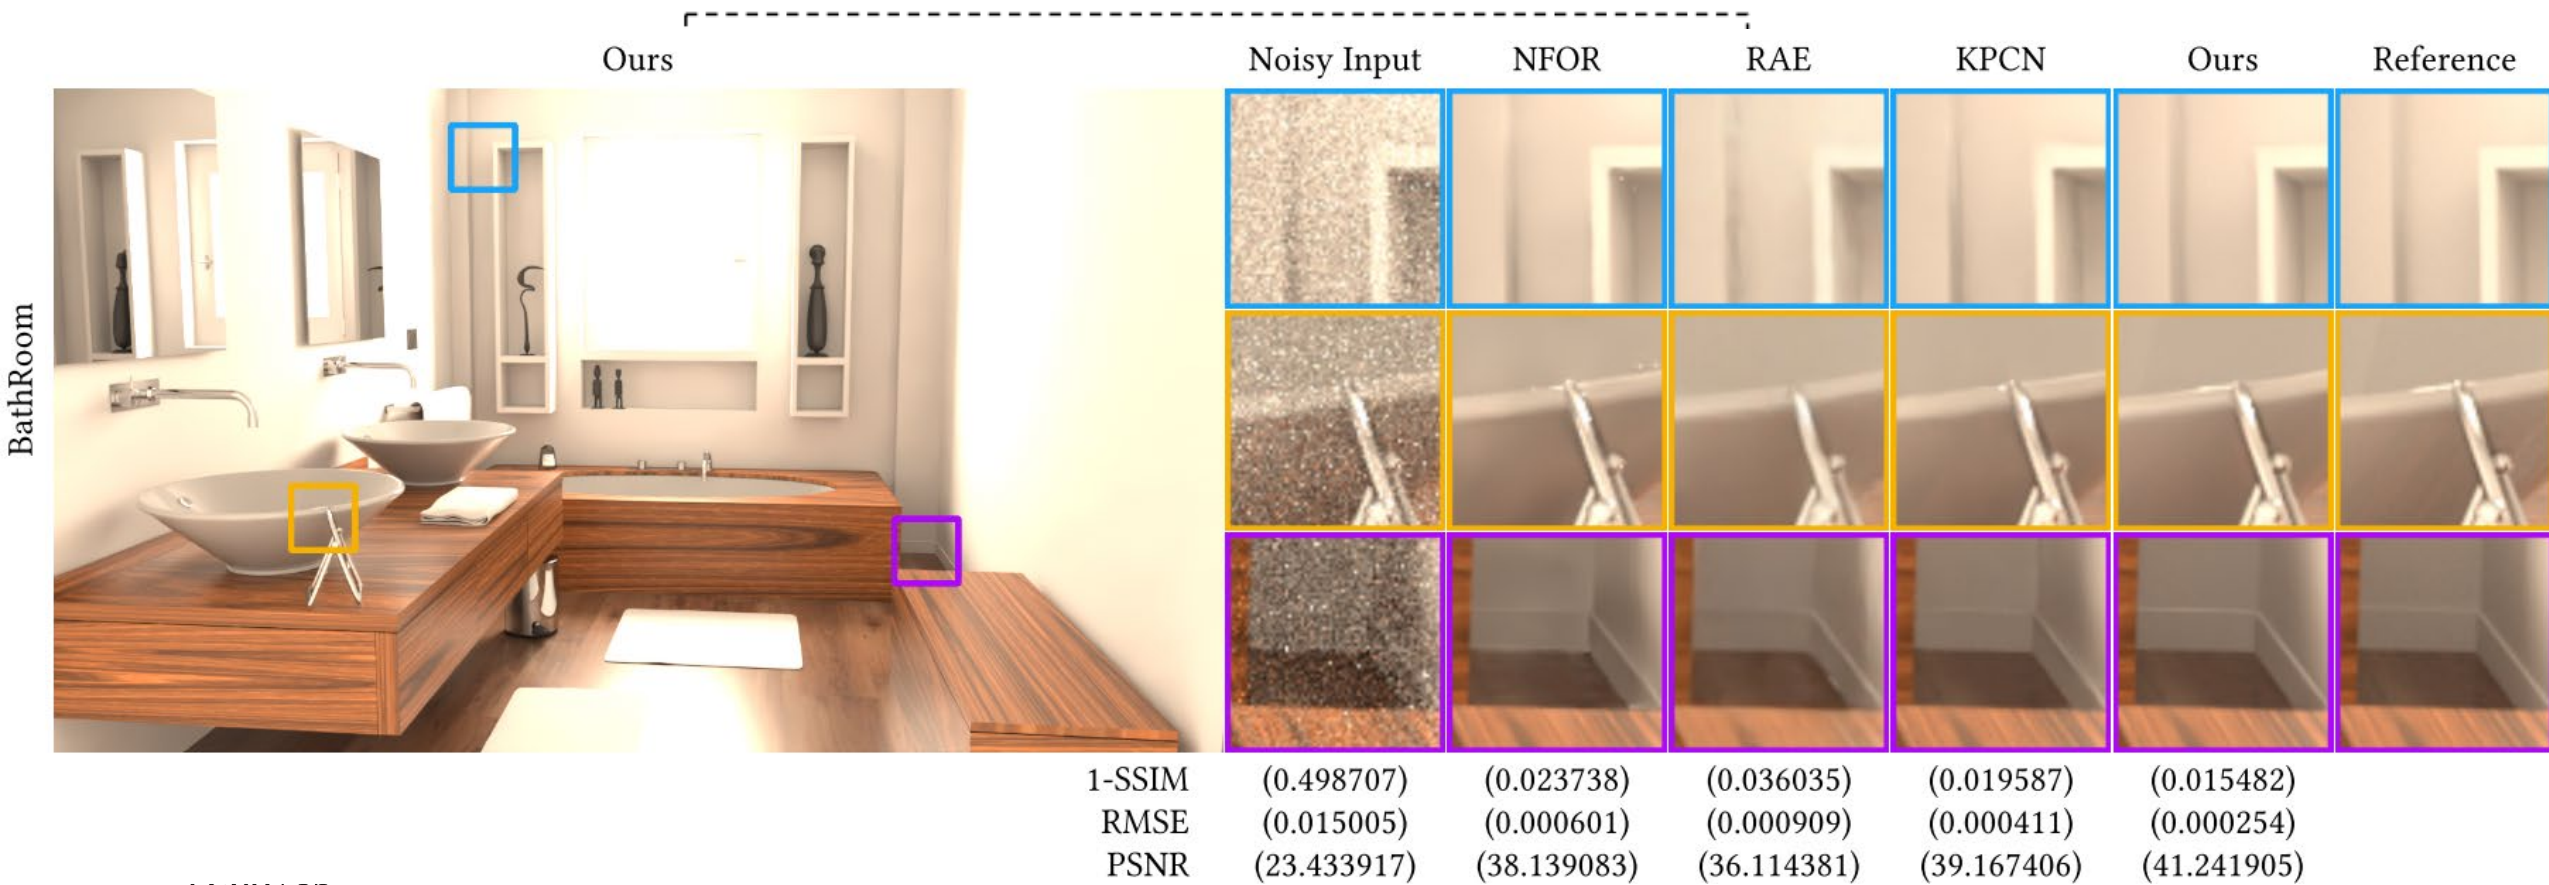


Kernel-predicting Convolutional Network (KPCN)




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




Jumping from Image-space to Sample-space


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



(a) scene




(b) input 16spp



(c) depth histogram



(d) inset 16spp



(e) reference



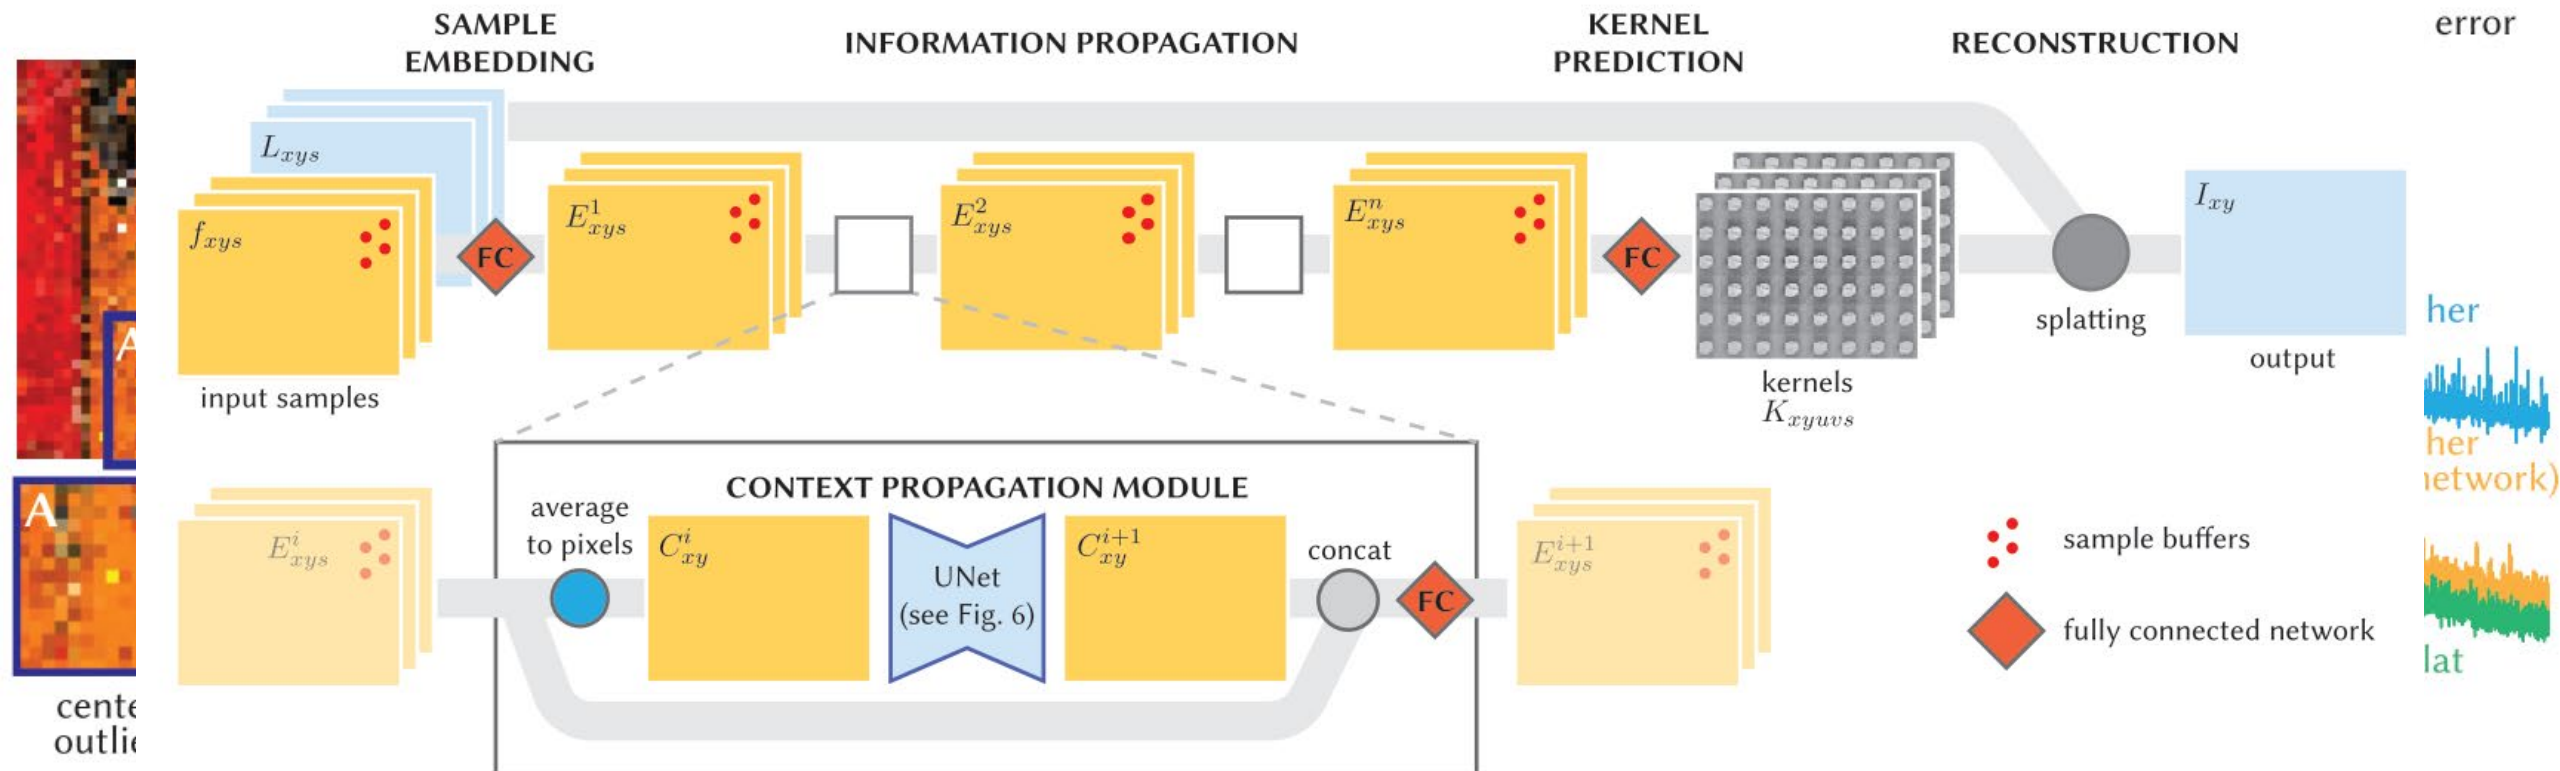
(f) [Sen 2012]



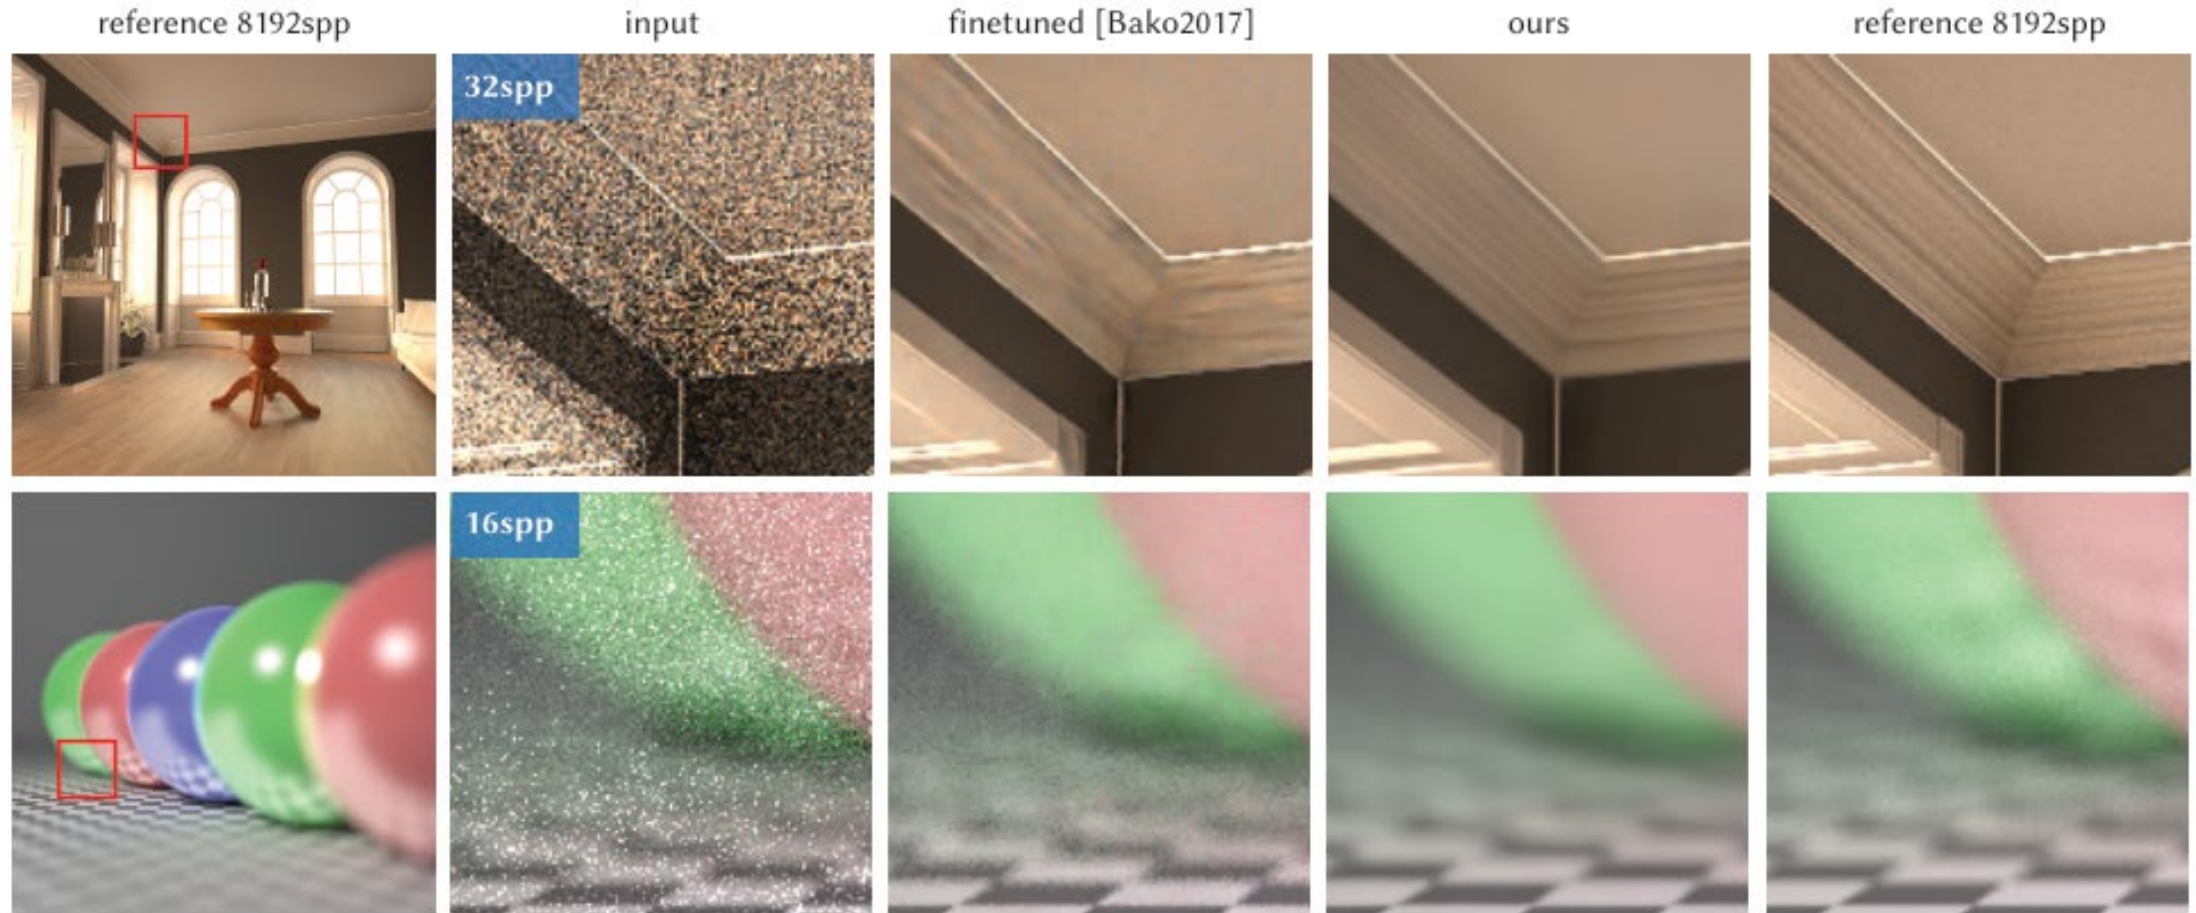
(g) ours

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



regression loss

convolutional network

concat



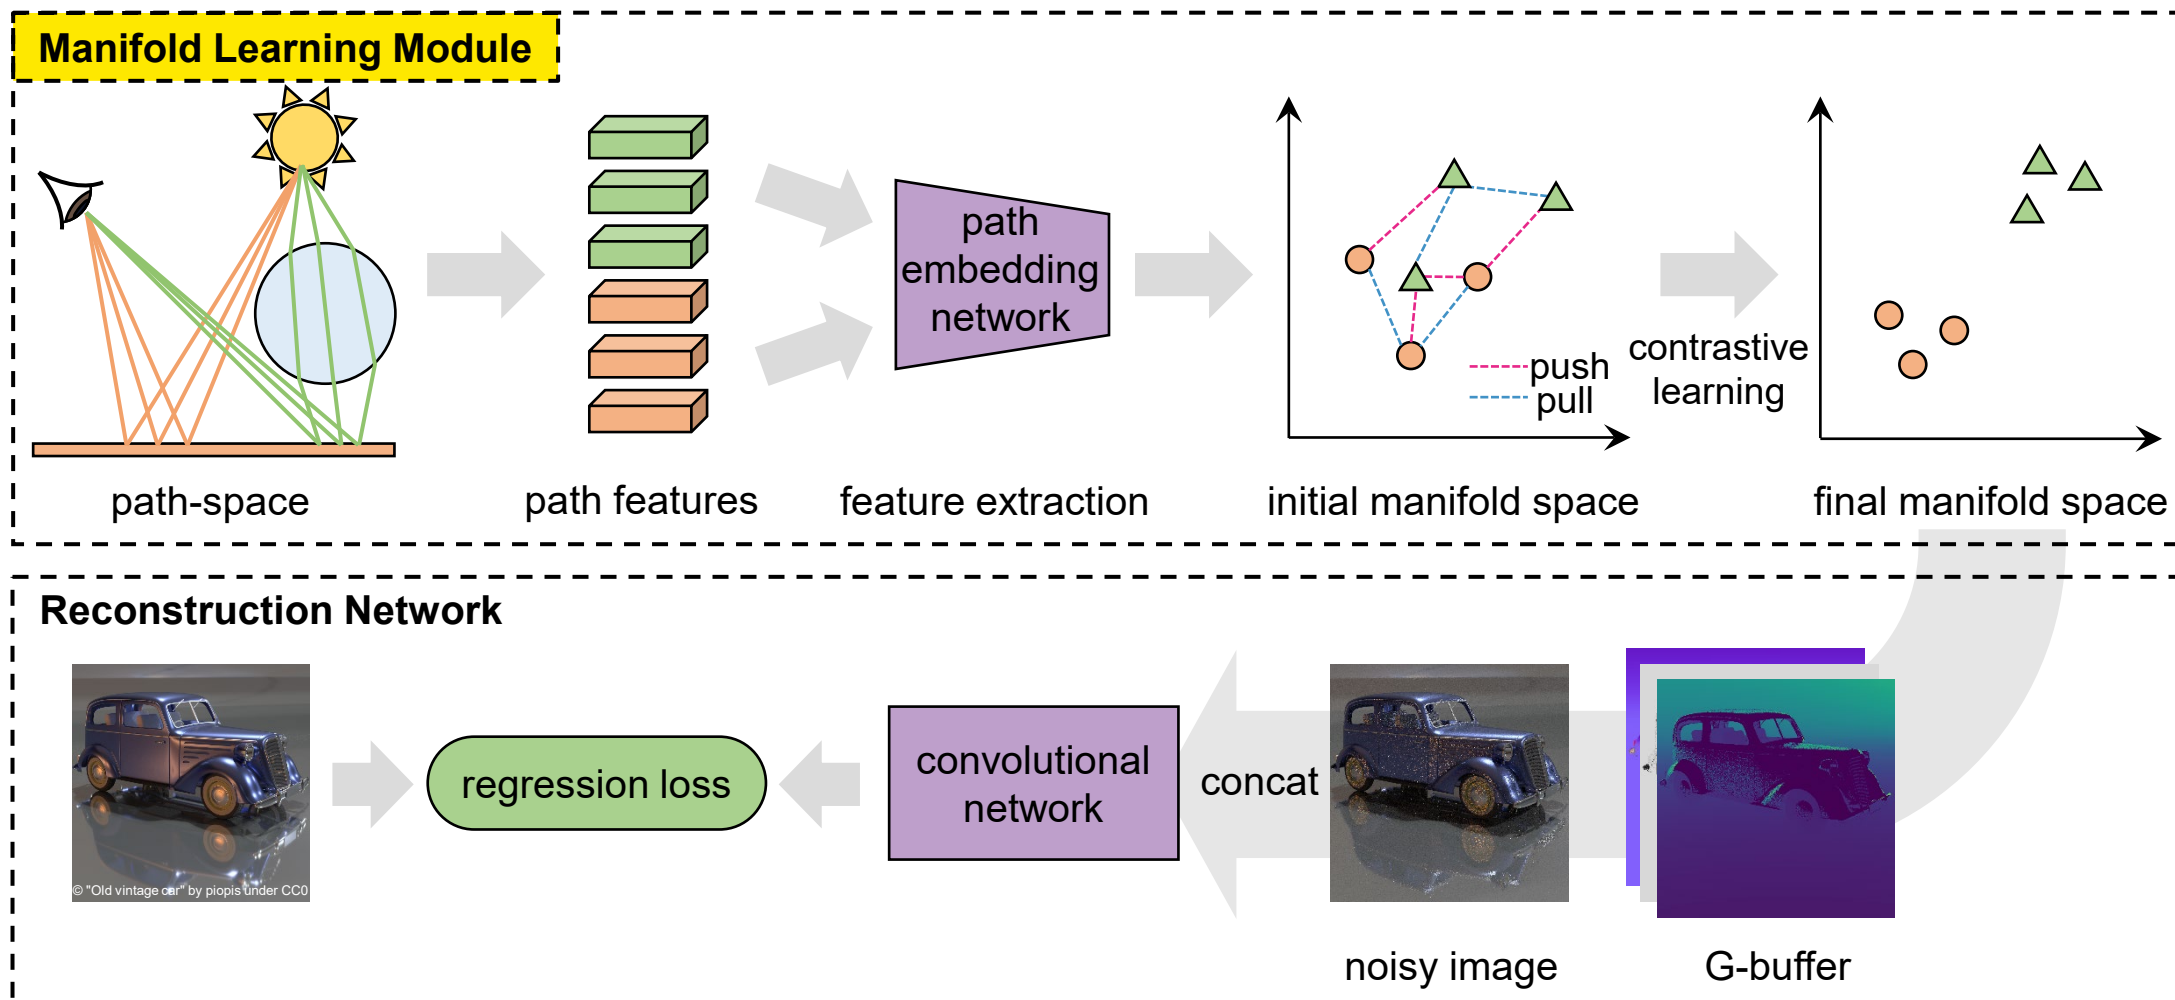
noisy image



G-buffer

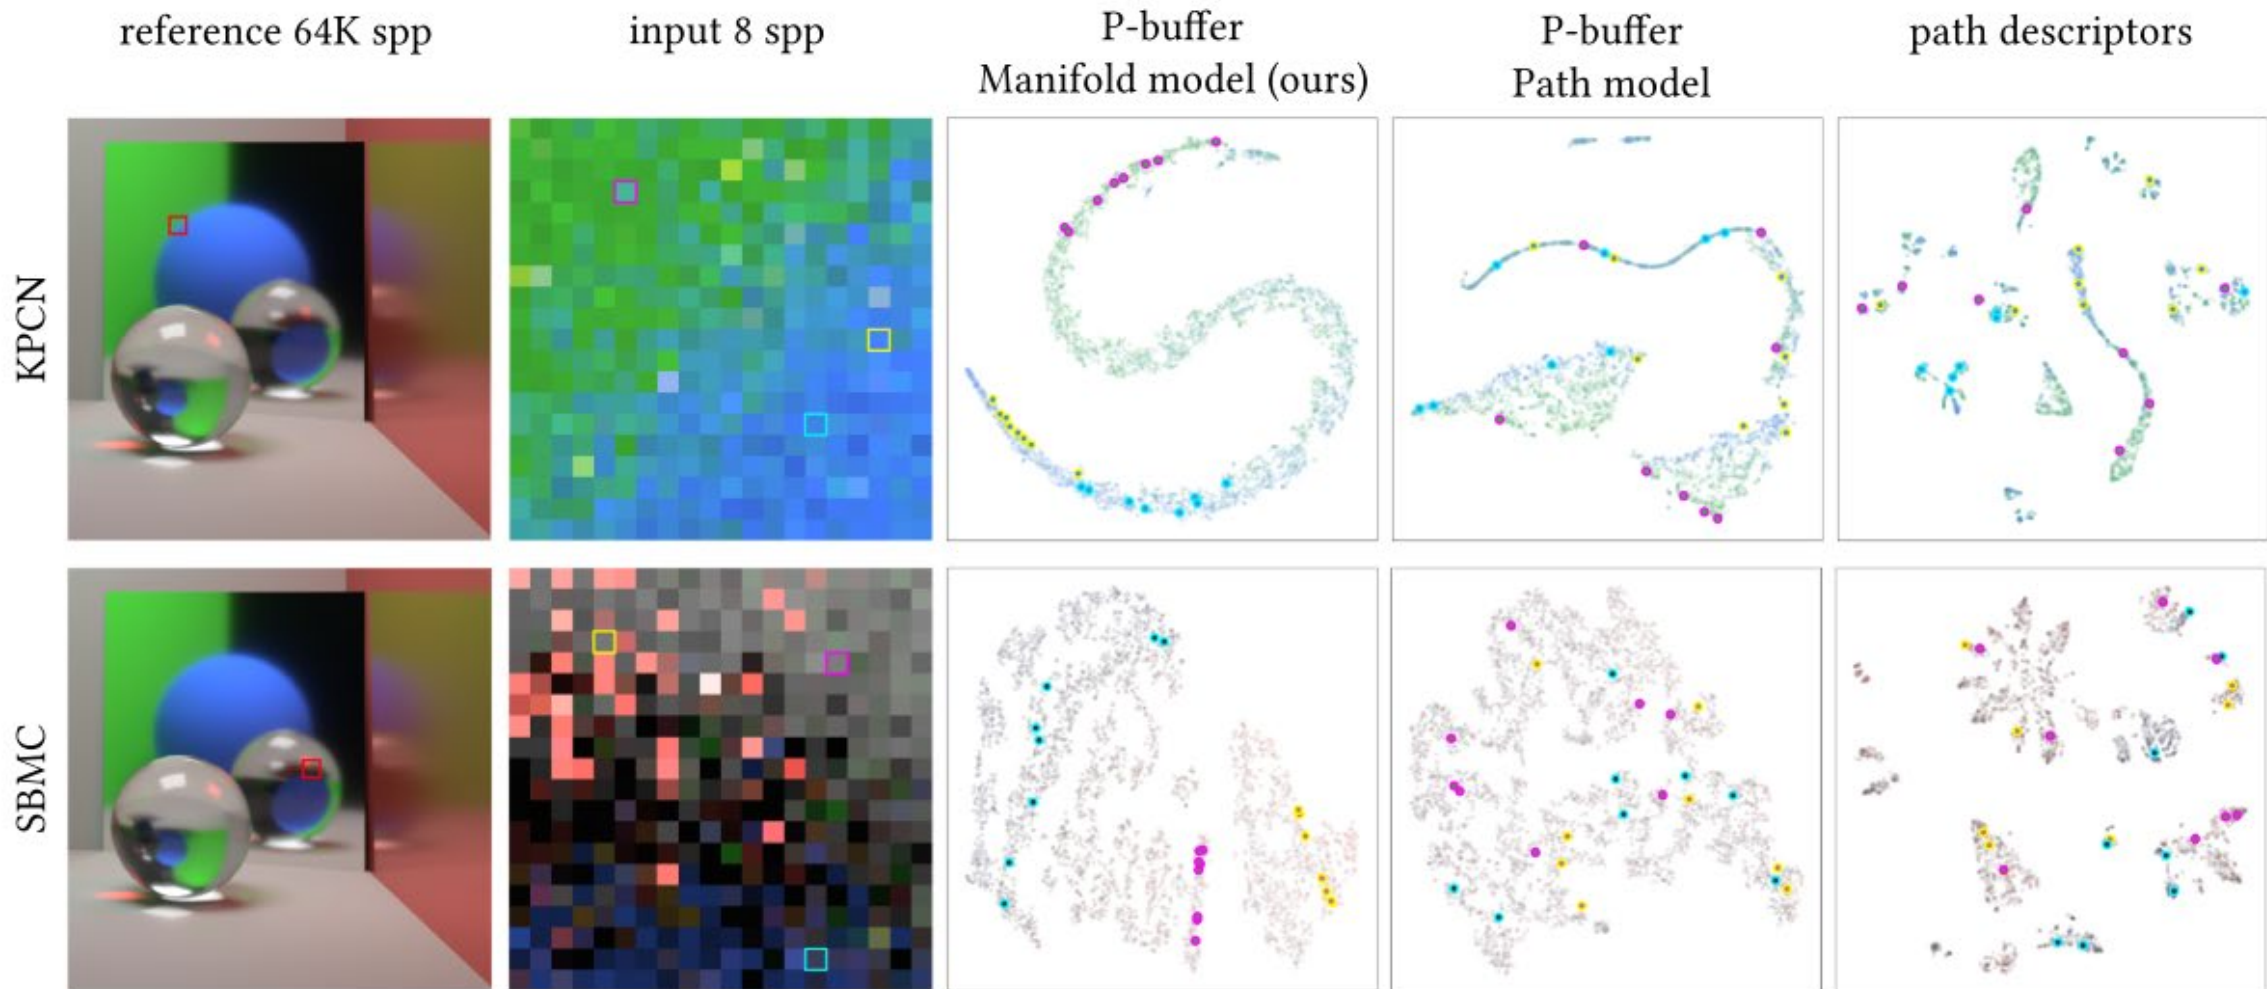
Manifold Learning for Path-space Features

- 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



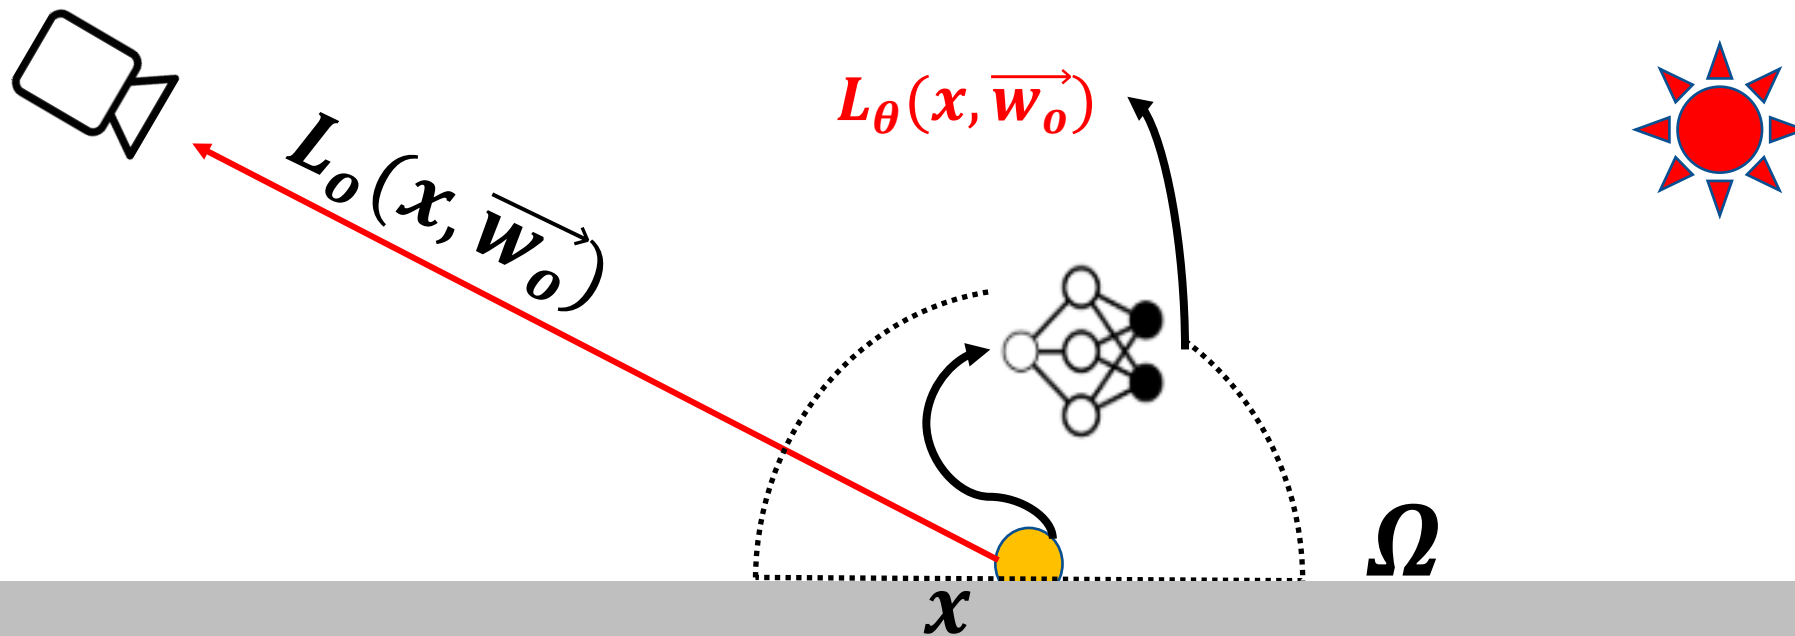
Manifold Learning for Path-space Features



Neural Radiance Caching

- Solving rendering equation via **Radiance-predicting Neural Network L_θ**


$$L_o(\mathbf{x}, \overrightarrow{\mathbf{w}}_o) = L_e(\mathbf{x}, \overrightarrow{\mathbf{w}}_o) + \int_{\Omega} f_r(\mathbf{x}, \overrightarrow{\mathbf{w}}_i, \overrightarrow{\mathbf{w}}_o) L_i(\mathbf{x}, \overrightarrow{\mathbf{w}}_i) (\overrightarrow{\mathbf{w}}_i \cdot \vec{\mathbf{n}}) d\overrightarrow{\mathbf{w}}_i$$
$$\sim L_e(\mathbf{x}, \overrightarrow{\mathbf{w}}_o) + \mathbf{L}_\theta(\mathbf{x}, \overrightarrow{\mathbf{w}}_o)$$



Neural Radiance Caching


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

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



Self-training for Neural Radiance Cache

- Minimize the loss between the **calculated radiance** and the **estimated radiance** 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 ($x_0 x_1 x_2$) where we used the cached(estimated) radiance in vertex x_2 for rendering

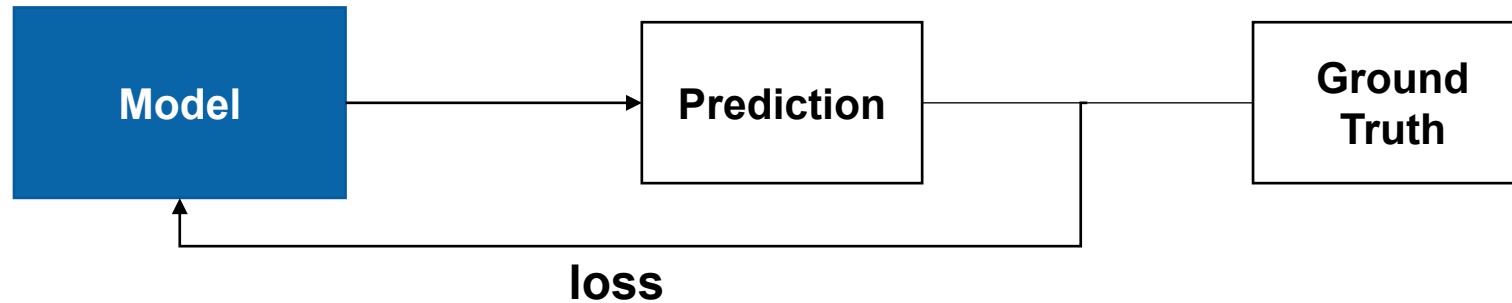


Result – 1spp Video


Result – 1spp Video w/ Image Denoising

Path Guiding using Reinforcement Learning

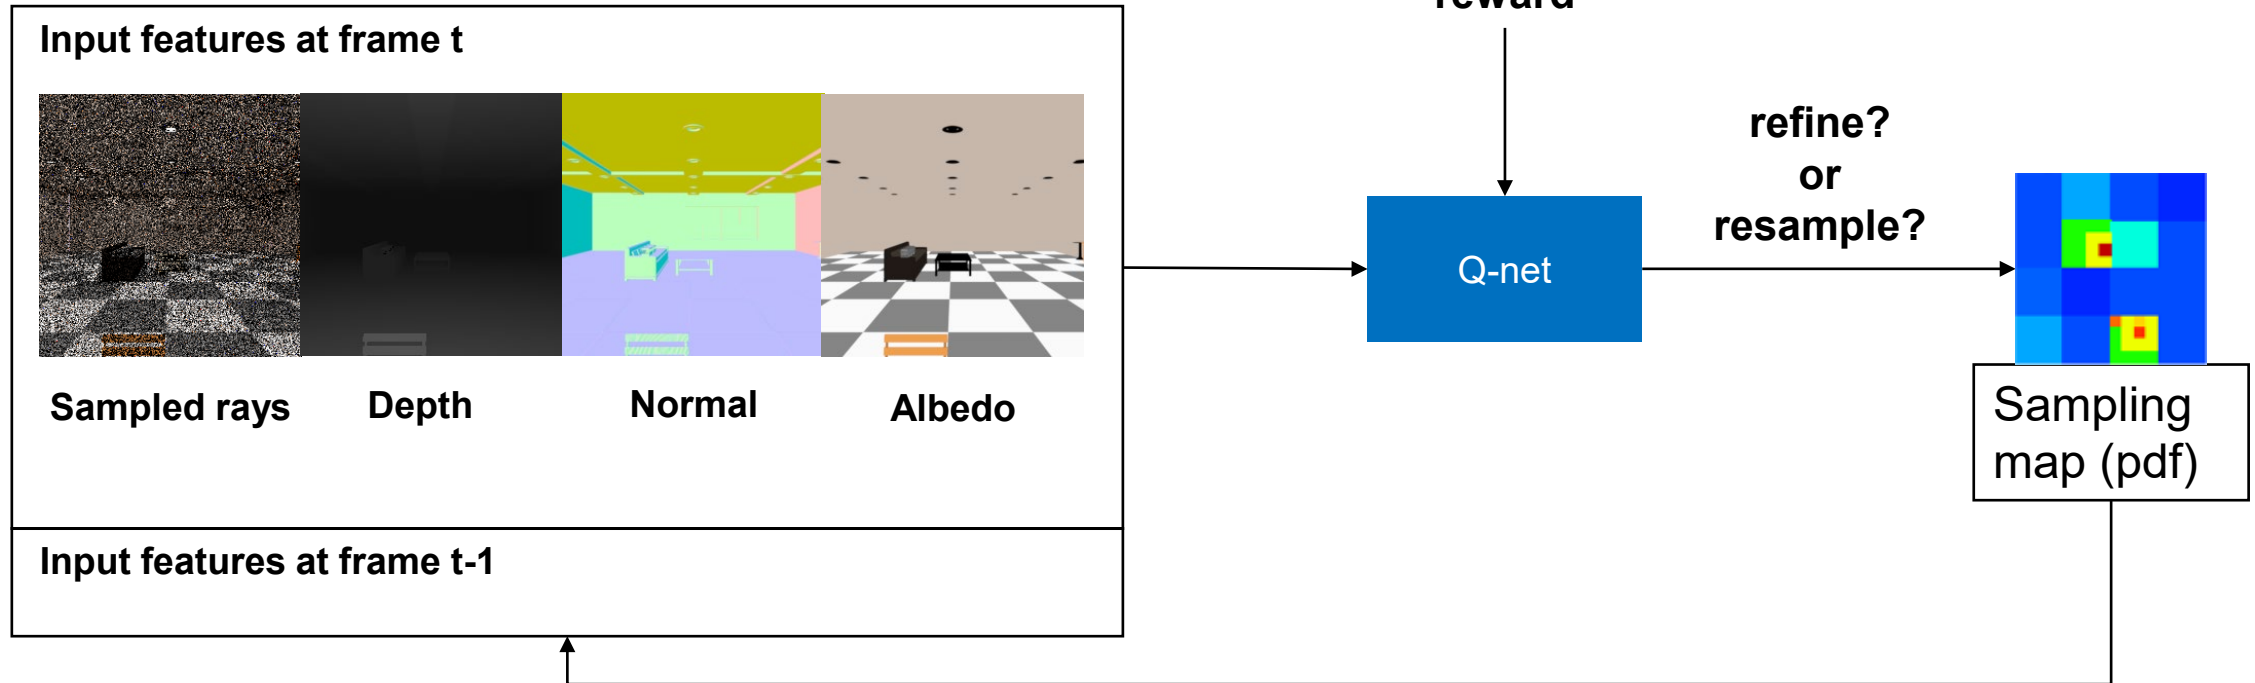
- Supervised Learning



- Reinforcement Learning
 - Find a policy $\pi: S \rightarrow A$ with maximum rewards




Sampling Map Generation (Q-Network)




generates new samples

1. Refine

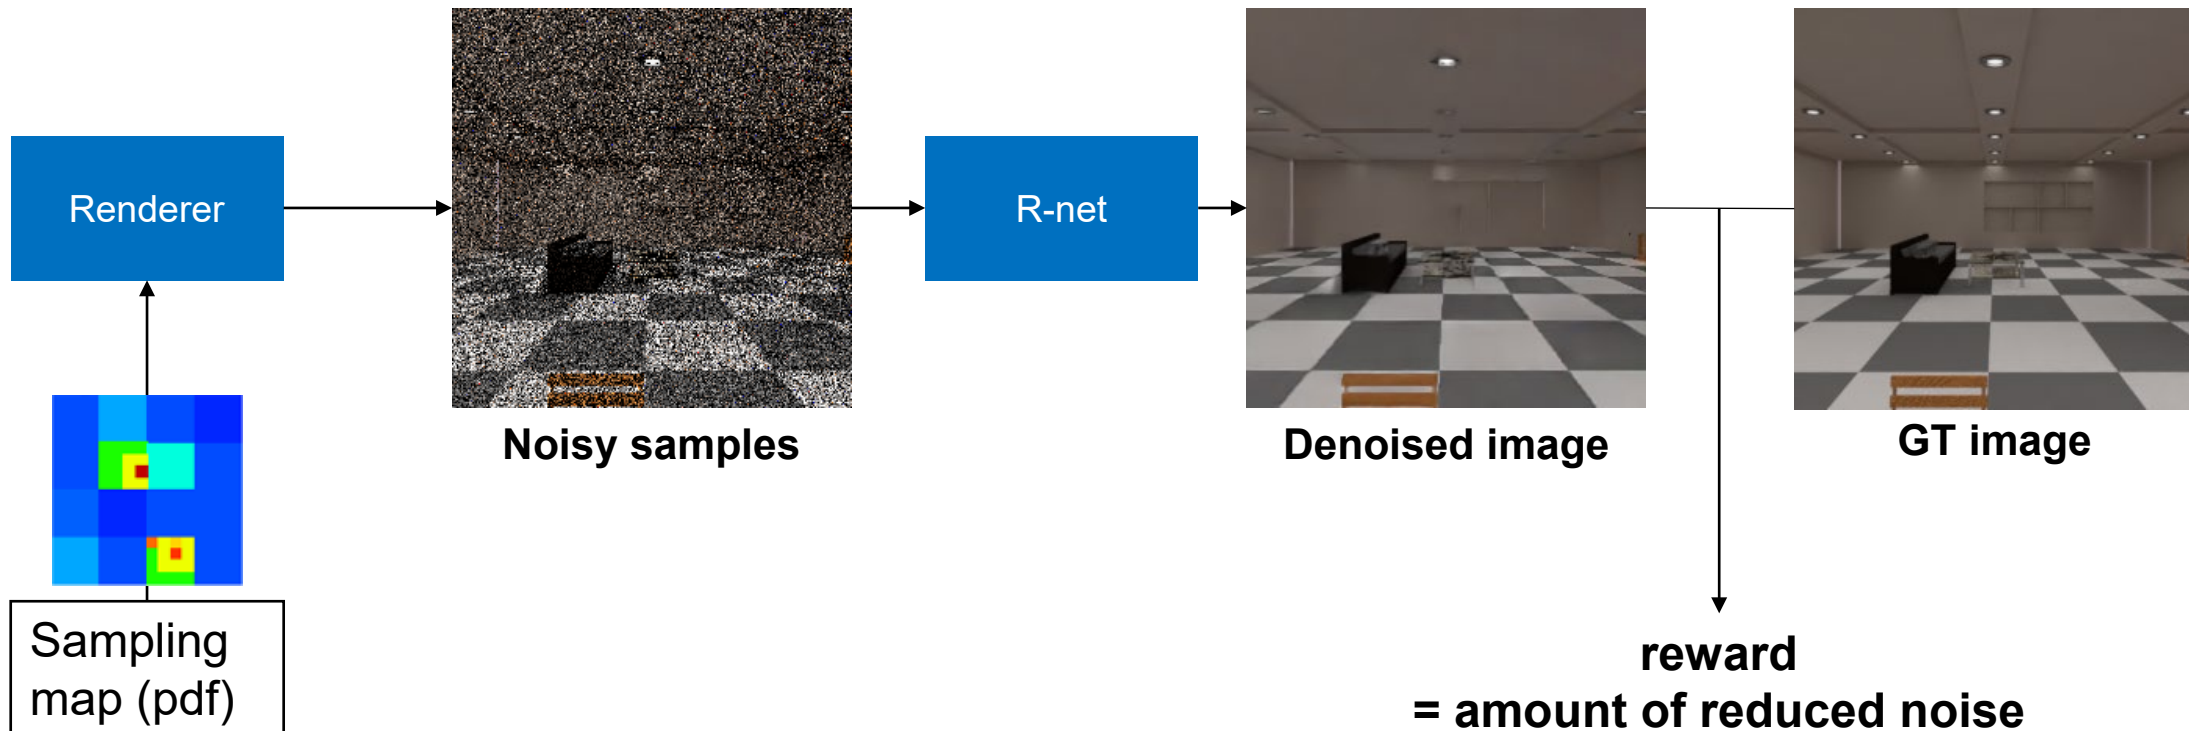


2. Resample



Rendering and Reconstruction (R-Network)

- Q-Net trained to estimate a sampling network that maximizes the reward



Path Guiding using Reinforcement Learning



(a) PT

(b) MEMLT


(c) GMM

(d) SD-tree

(e) Ours

(f) Reference

(g) Distribution



(a) PT

(b) MEMLT

(c) GMM

(d) SD-tree

(e) Ours

(f) Reference

(g) Distribution

rMSE 0.0011

rMSE 0.0015

rMSE 0.00072

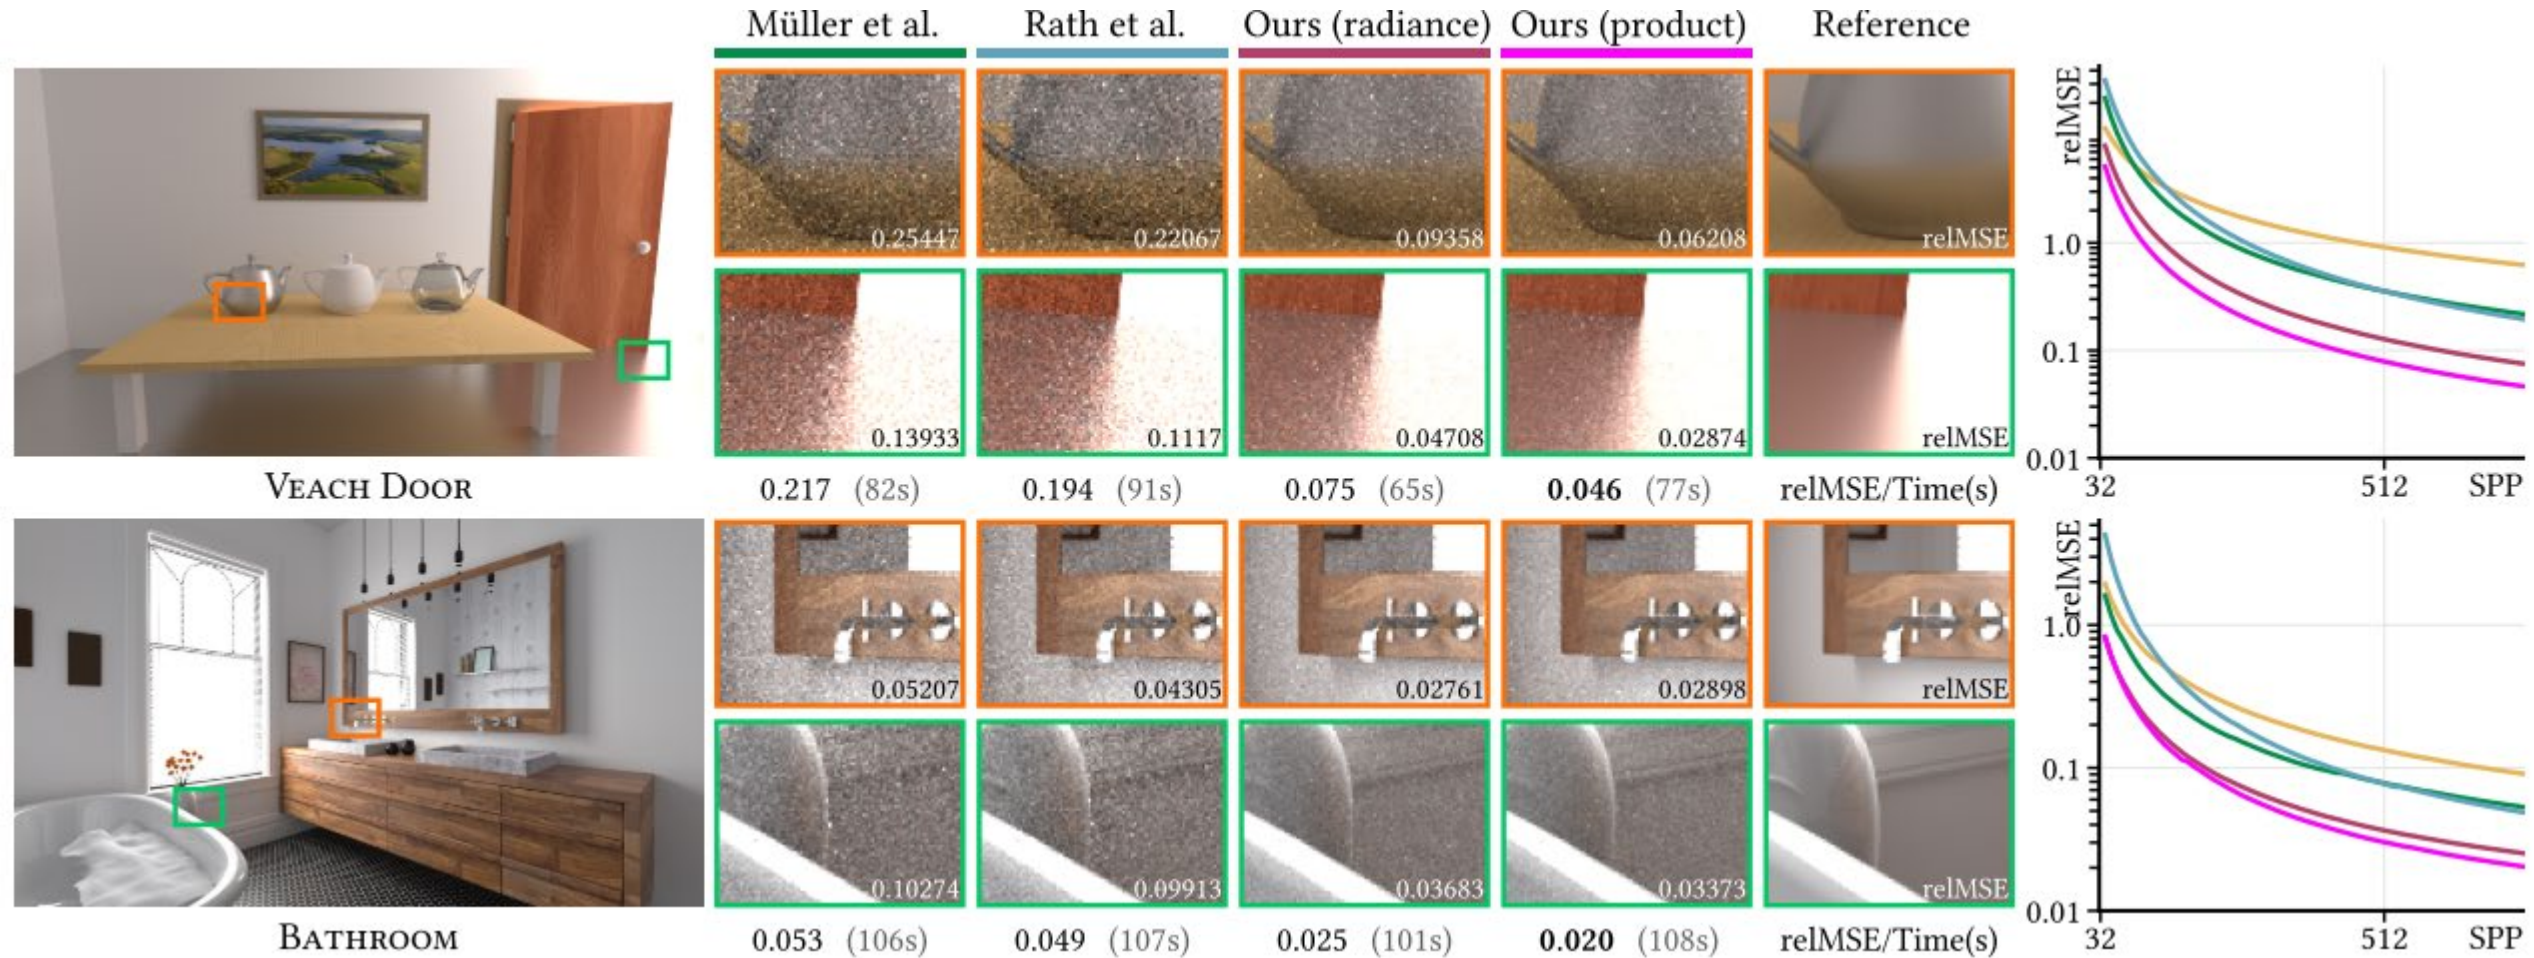
rMSE 0.020

rMSE 0.00035



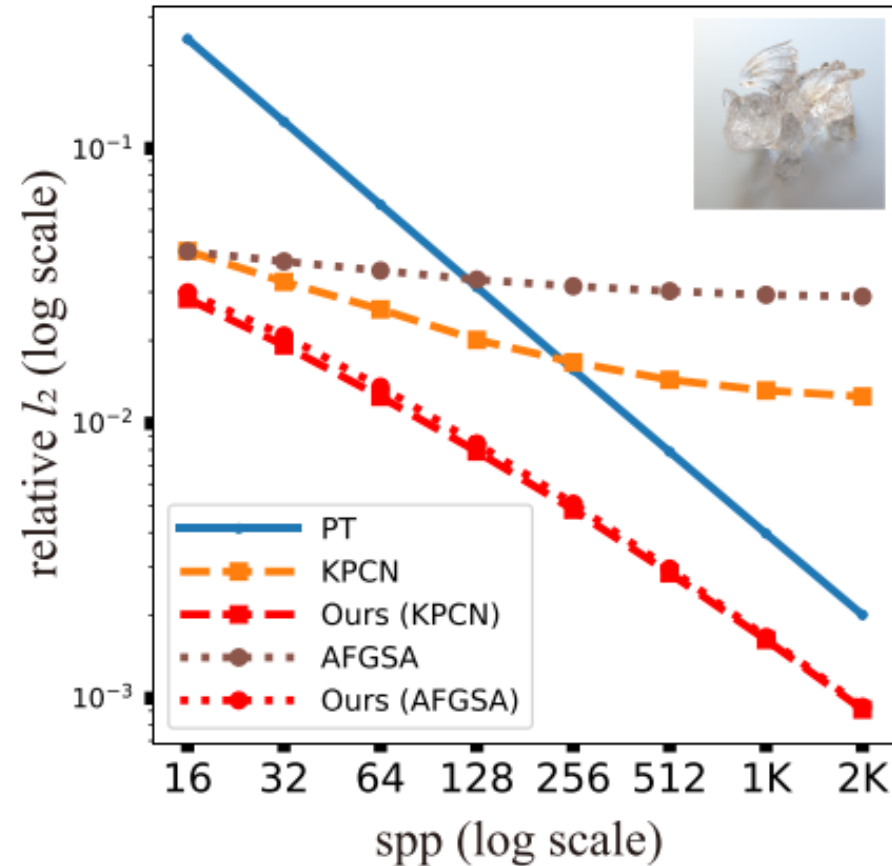
Implicit Neural Representation for Path Guiding

- Using implicit neural representation that encodes parameters for mixture models
- Using vMF (von Mises-Fisher) distribution $v(\omega | \mu, \kappa) = \frac{\kappa}{4\pi \sinh \kappa} \exp(\kappa \mu^T \omega)$



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), σ : 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), σ : 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

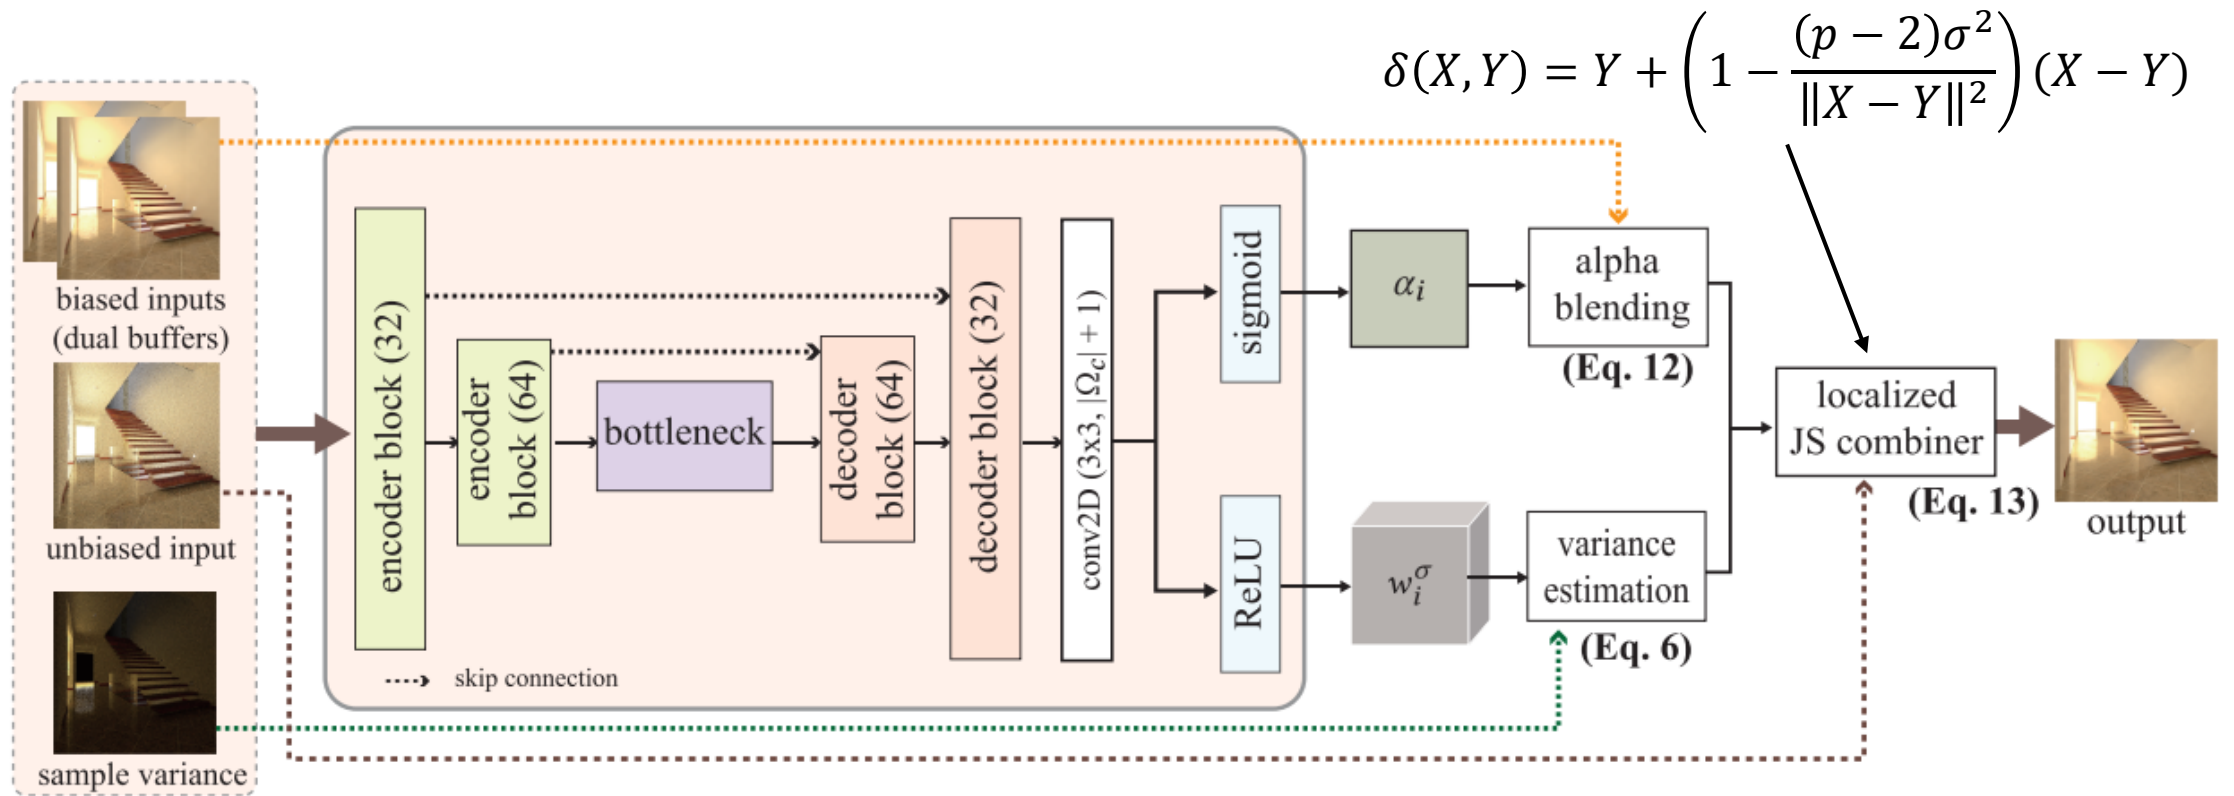
- Path Traced Result X : Noisy but Unbiased (Bias \downarrow , Variance \uparrow)
- Denoised Result Y : Smooth but Biased (Bias \uparrow , Variance \downarrow)
- James-Stein Estimator shrinks X towards Y as $\delta(X, Y) = Y + \left(1 - \frac{(p-2)\sigma^2}{\|X-Y\|^2}\right) (X - Y)$
 - p : Dimension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



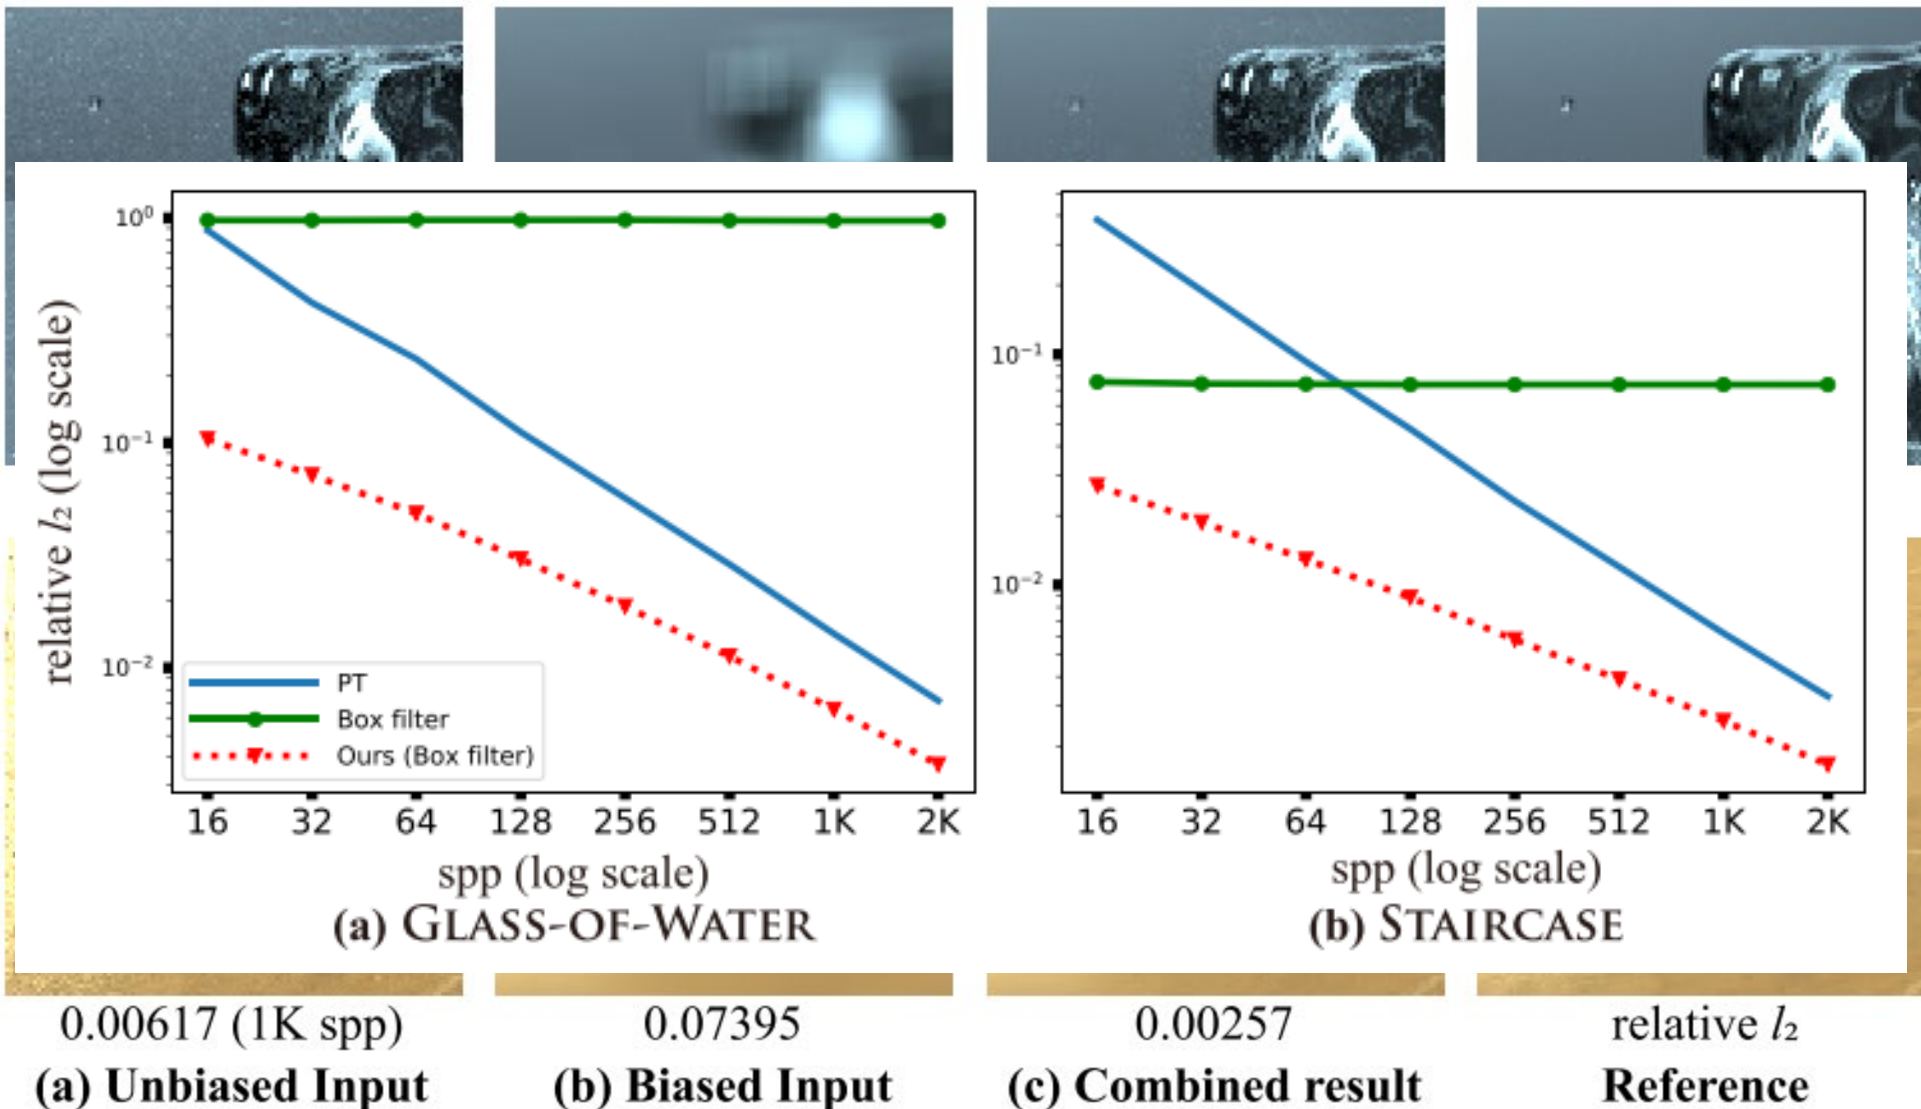
<Intuition>
Balancing the bias and variance of between
path traced result and denoised result

Neural James-Stein Combiner


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




Neural James-Stein Combiner



Neural James-Stein Combiner



time / relative l_2



36.7 m / 0.07092

36.7 m / 0.03529

36.7 m / 0.00727

1M spp



time / relative l_2



121.0 s / 0.00661

121.0 s / 0.08584

120.4 s / 0.00180

64K spp

(a) BDPT

(b) SPPM

(c) Ours

(d) Reference

What We Covered

- Path-space MC noise reduction
 - Path Guiding
 - Path Reuse (Path-space Filtering)
- Image-space MC noise reduction
 - Image Denoising
 - Adaptive Sampling
- Learning-based MC noise reduction
 - Image-space methods
 - Sample- & Path-space methods
 - Path-guiding
 - Post-post processing