Deep Adaptive Sampling for Low Sample Count Rendering

Alexandr Kuznetsov [*]
Nima Khademi Kalantari [*]
Ravi Ramamoorthi [*]

Eurographics Symposium on Rendering ‘18

Presenter : Jaehyun Ha

UC San Diego [*]

NetS&P Research Lab @ KAIST

Recap: Last week's presentation

- **<u>ReLight My NeRF: A Dataset for Novel View Synthesis and Relighting of Real World Objects</u>** (M. Toschi et al., SIGGRAPH 2023) presented by Team 5
- Presents a novel dataset for relightable NeRF – framing real world objects under challenging OLAT conditions and annotated with accurate camera and light poses
- Limitations: Does not support 360 degrees scans & varying light temperature/color

Background: Denoisers

- Interest about real-time ray tracing with ***extremely low sampling rates*** has been risen
- Chaitanya et al., "Interactive reconstruction of monte carlo image sequences using a recurrent denoising auto encoder", ACM TOG 2017

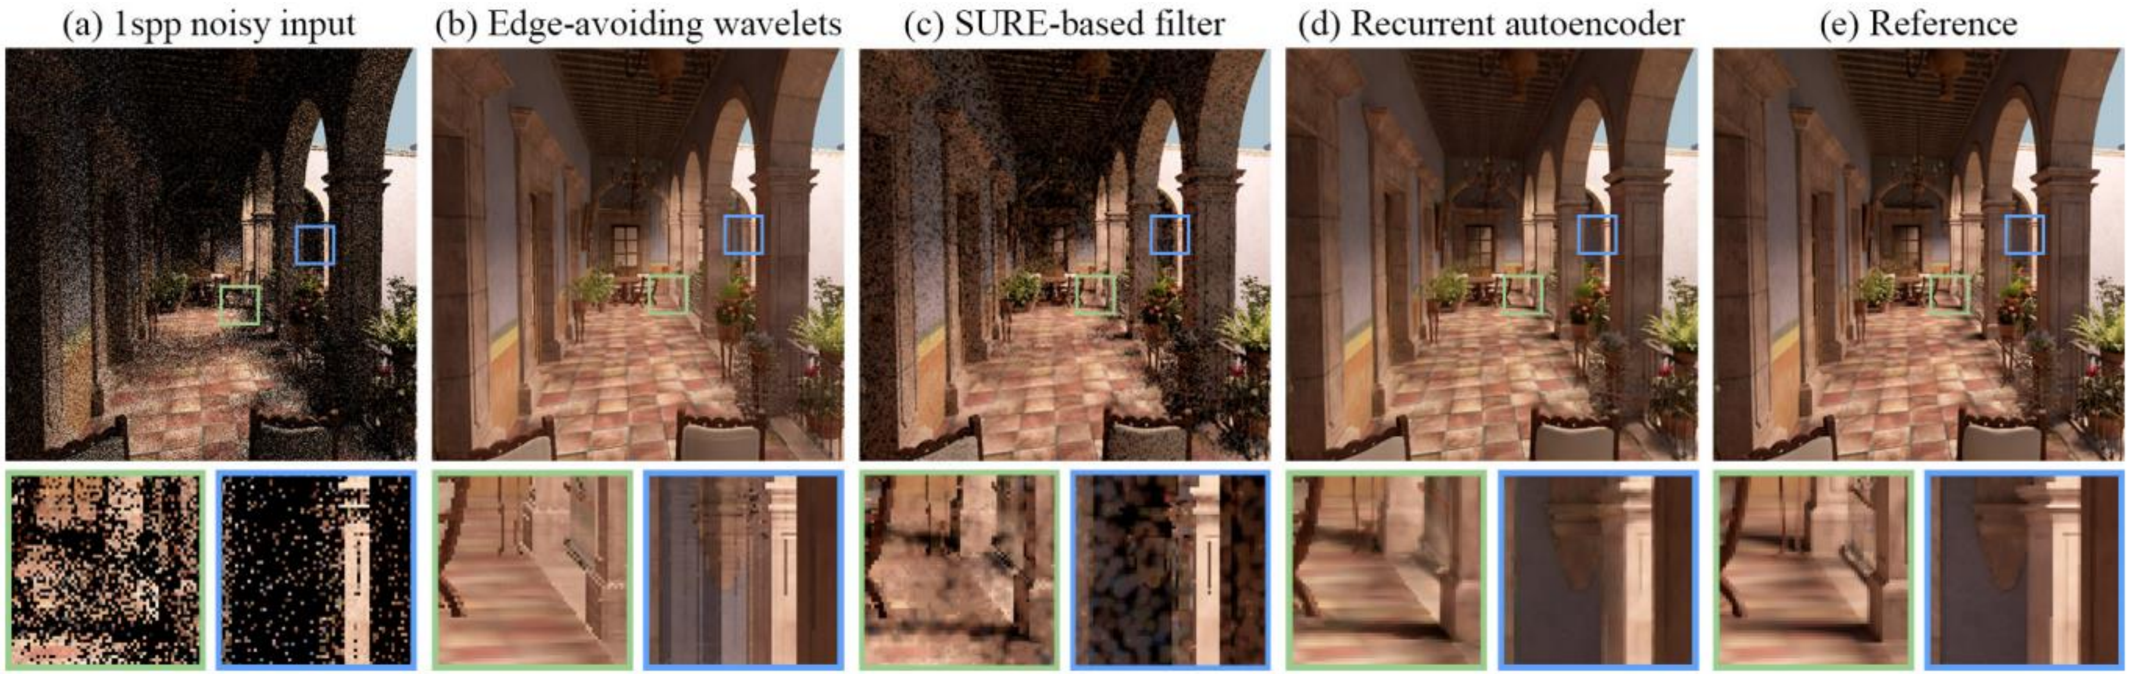

Background: Drawbacks of existing Denoisers for low SPP

- Existing approaches focus ***only*** on denoising problem
- Use ***uniformly sampled images*** as the input

Uniformly sampled

Ground truth

Background: Adaptive sampling approaches

- ***Adaptive sampling***: Optimizing techniques in MC rendering which allocates more samples to the areas of the suspicious part of the image
- Estimate a ***sampling map*** from variance, coherence, frequency data of the image

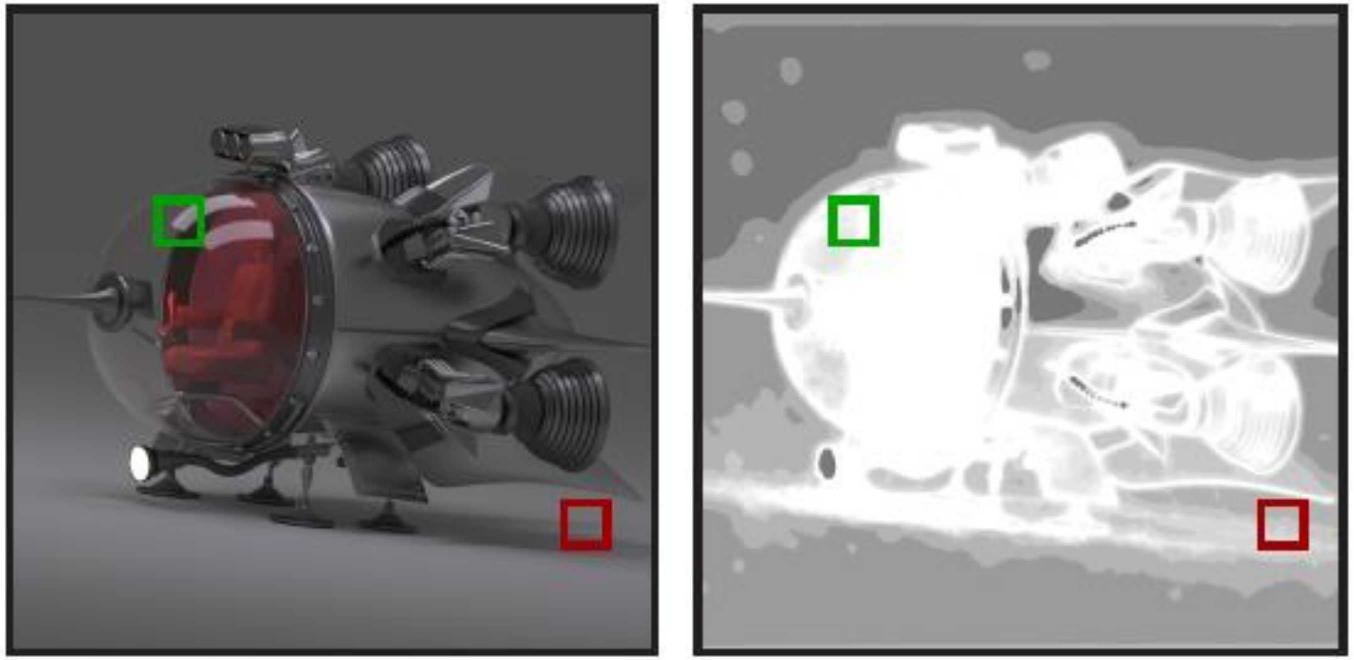

Sampling map

Background: Adaptive sampling approaches

- ***Sampling map*** from previous approaches requires uniformly sampled image rendered with high spps
- Thus, recent denoising techniques only use uniformly sampled images as the input

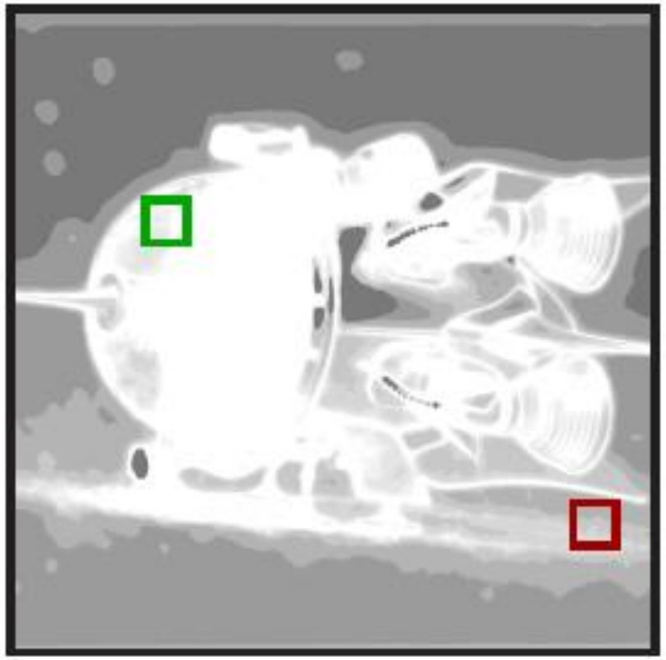

Sampling map

Pipeline of deep adaptive sampling for low SPPs

- Use deep learning for both adaptive sampling and denoising stage

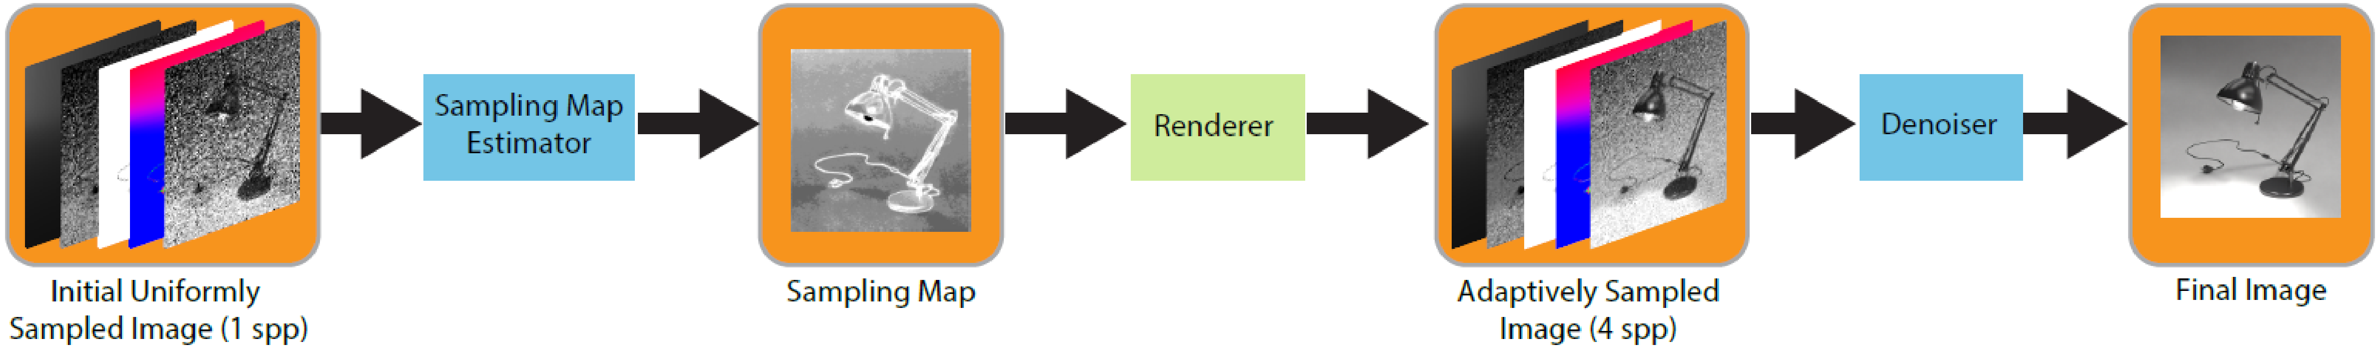

Figure 2: *In our system, we first render a scene with one spp and use it to calculate the sampling map. Then, we use the sampling map to render three additional samples. Finally, we denoise the resulting rendered image with an average of 4 spp to obtain the final denoised image.*

Step 1: Making a input for sampling map estimator

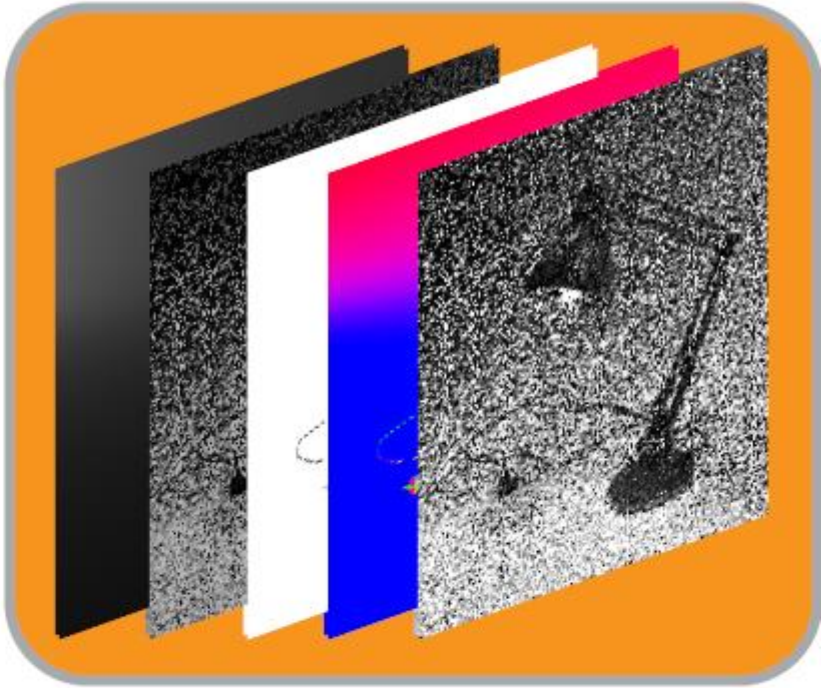

Initial Uniformly
Sampled Image (1 spp)

- ***Noisy image with 1 spp*** in RGB format (3 channels)
- ***Textures*** in RGB format (3 channels)
- ***Shading normal*** (3 channels)
- ***Depth*** (1 channel)
- ***Direct illumination visibility*** (1 channel)

Step 2: Sampling Map Estimator

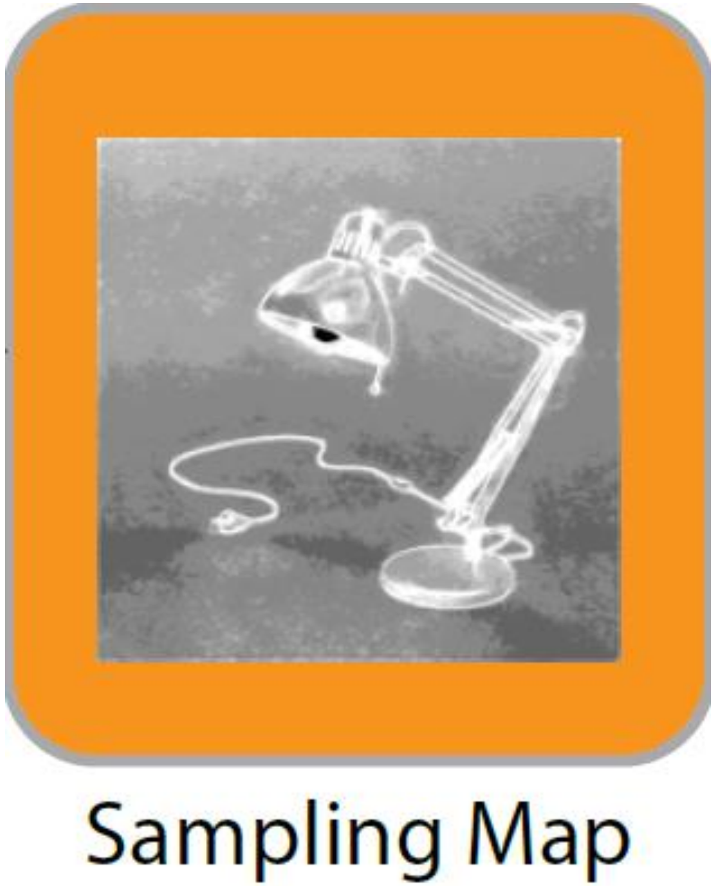

n=3

$$s(p) = \text{round}\left(\frac{M}{\sum_{j=1}^{M} e^{x(j)}} \times e^{x(p)} \times n\right)$$

- Calculates how much extra sample should be applied to each pixel
- CNN trained in ***end-to-end*** fashion: reducing error between final denoised image and the ground truth image

Step 3: Render adaptively sampled image

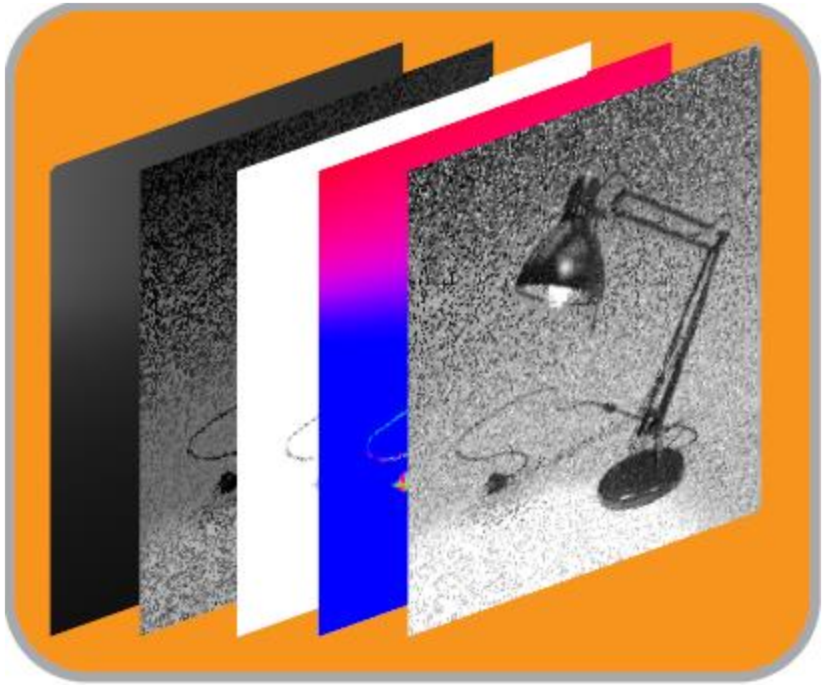

Adaptively Sampled Image (4 spp)

- Sampling map is used by the renderer to adaptively distribute 3 spp and produce an image with 4 spp (in avg.)

Step 4: Denoise the adaptively sampled image

Final Image

- Denoiser takes the adaptively rendered image and produce a high-quality image comparable to ground truth
- Train denoiser CNN with randomly generated 4 spp image as an input

Results

- ***Deep adaptive sampling (DA)*** outperforms ***equal quality uniform sampling (EQ)***

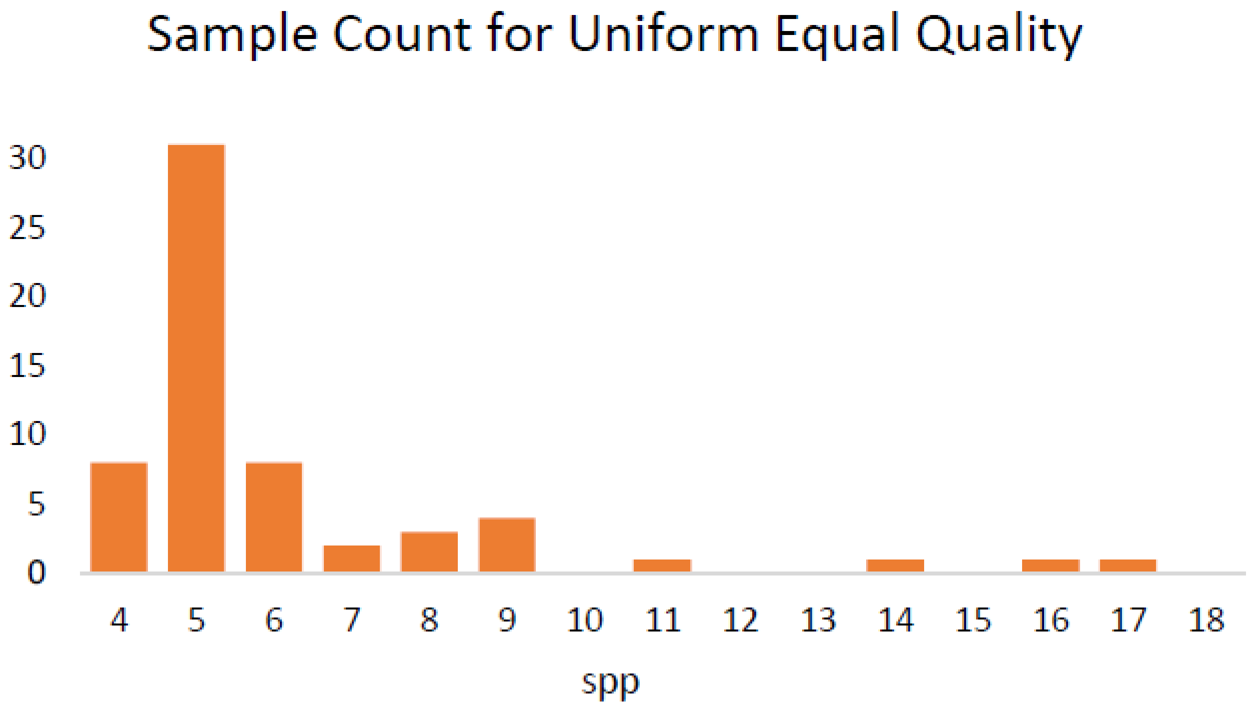

Figure 6: *We show the histogram of required sampling rate to achieve equal quality on all the 60 test scenes. The average of the required sampling rates are 6.12.*

DA		***EQ***
4	***<***	***6.12***

50% ⬆

Results

- ***Deep adaptive sampling (DA)*** outperforms ***other sampling approaches***

Results

- ***Deep adaptive sampling (DA)*** outperforms ***other sampling approaches***

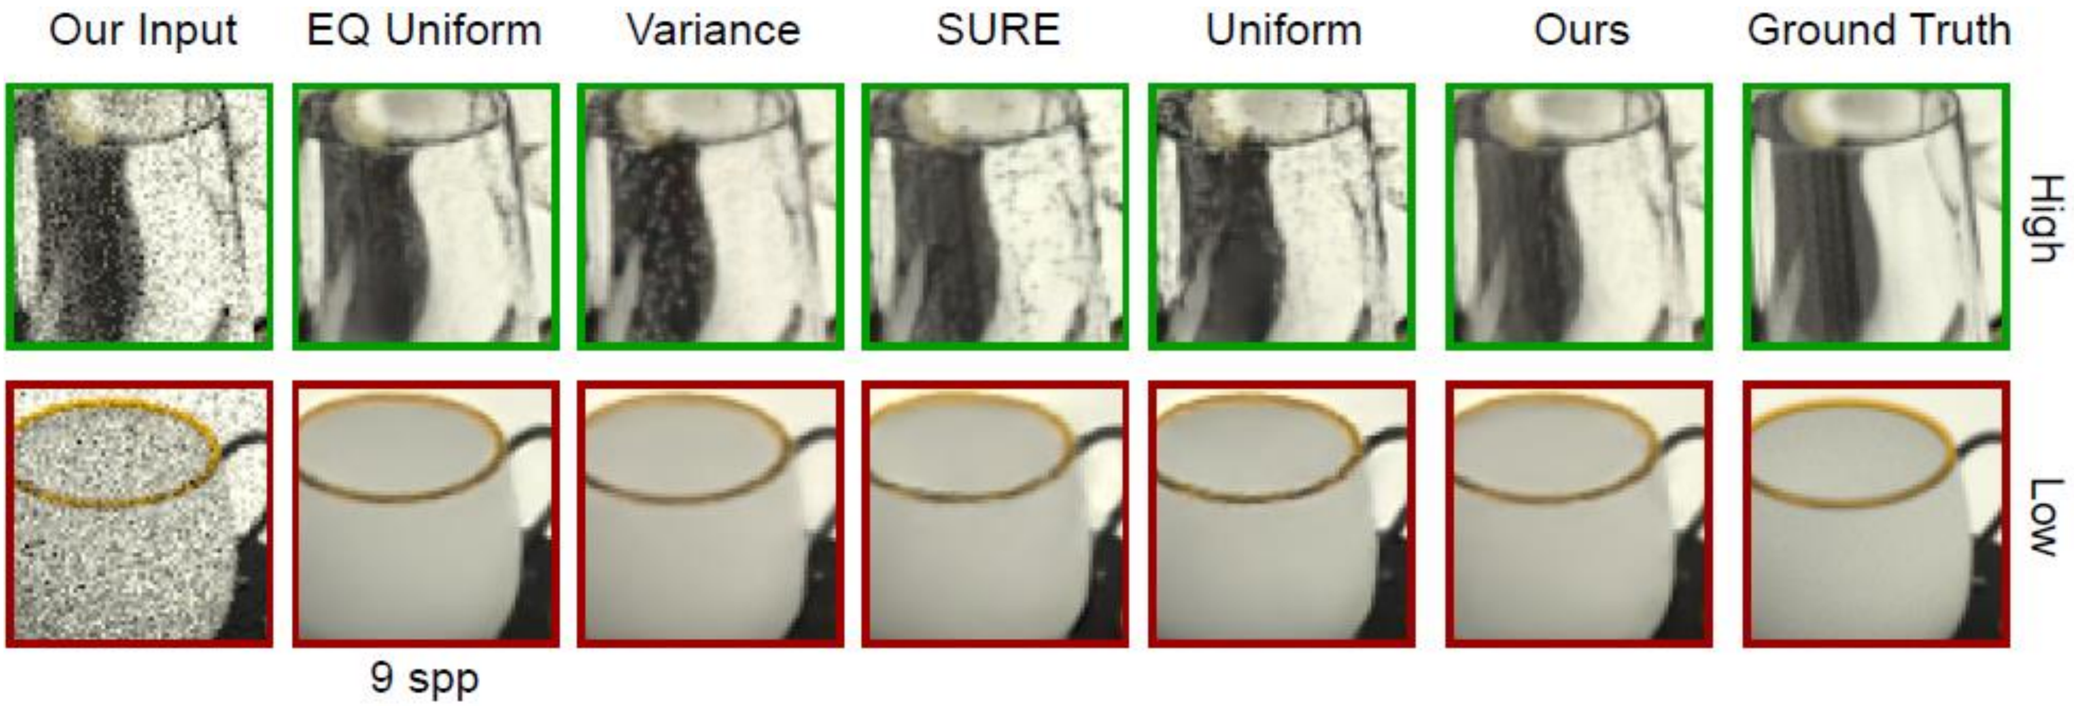

Limitations

- Not able to properly sample thin structures in noisy areas

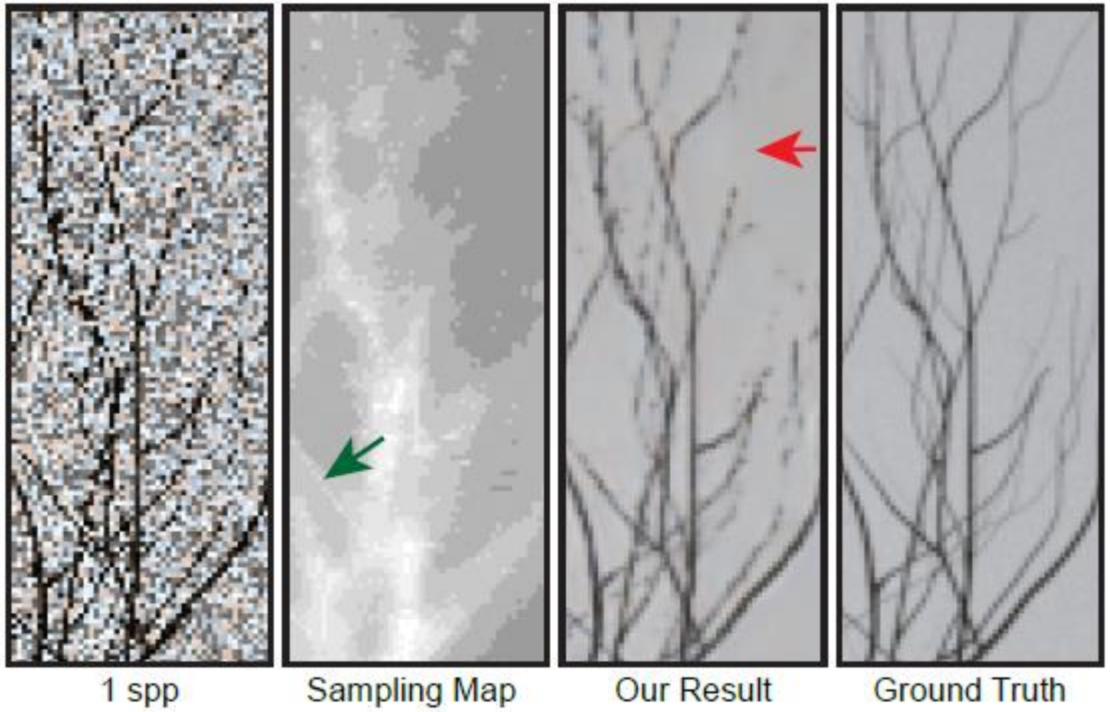

Figure 11: *Our sampling map cannot capture very thin structures such as some branches on this bush shown by the red arrow. However, it still able to capture a branch 2 pixel wide shown by the green arrow.*

Takeaways

Deep adaptive sampling for low sample count rendering is a first learning-based approach that enables adaptive sampling in extremely low sample count MC rendering.

- Use CNN to estimate a sampling map from a uniformly sampled 1 spp image.
- Use the map to render 3 additional samples (in avg.) adaptively
- Denoise the rendered image using another CNN to produce the final image

Future Work

- Using GAN to improve the perceptual quality of the results
- Investigate the performance at <1 sample per pixel

Thank You

Jaehyun Ha

jaehyunh1@kaist.ac.kr

NetS&P
Research Lab @ KAIST