



# CharacterGen: Efficient 3D character generation from single images with multi-view pose calibration



A-poses



# EUROGRAPHICS 2022

43<sup>RD</sup> ANNUAL CONFERENCE OF  
THE EUROPEAN ASSOCIATION FOR COMPUTER GRAPHICS



**REIMS · FRANCE**  
**APRIL 25-29 / 2022**

## PROGRESSIVE DENOISING OF MONTE CARLO RENDERED IMAGES

ARTHUR FIRMINO, JEPPE REVAL FRISVAD, AND HENRIK WANN JENSEN

April 26<sup>th</sup>, 2022  
Centre des Congrès de Reims



**LUXION**

DTU Compute  
Department of Applied Mathematics and Computer Science



# Backgrounds - Denoisers


More samples

Secondary estimator

Generate  $N$  random samples  $x_i$

Estimator: 
$$\langle I \rangle = I_{\text{sec}} = \frac{1}{N} \sum_{i=1}^N f(\bar{x}_i)$$

Variance 
$$\sigma_{\text{sec}}^2 = \sigma_{\text{prim}}^2 / N$$



© Kavita Bala, Computer Science, Cornell University

Lecture 04. Montecarlo Integration



Special Lecture. Monte Carlo Noise Reduction

# Background - Image Filters

Source layer

5	2	6	8	2	0	1	2
4	3	4	5	1	9	6	3
3	9	2	4	7	7	6	9
1	3	4	6	8	2	2	1
8	4	6	2	3	1	8	8
5	8	9	0	1	0	2	3
9	2	6	6	3	6	2	1
9	8	8	2	6	3	4	5

Convolutional kernel


-1	0	1
2	1	2
1	-2	0

Destination layer


	5						

$$\begin{aligned} &(-1 \times 5) + (0 \times 2) + (1 \times 6) + \\ &(2 \times 4) + (1 \times 3) + (2 \times 4) + \\ &(1 \times 3) + (-2 \times 9) + (0 \times 2) = 5 \end{aligned}$$

# Background - Image Filters



# Background - Image Filters



(a) MC Input (8 spp) (b) Our approach (RPF)



Cross-Bilateral Filter

High-order filter

*SEN P., DARABI S.: On filtering the noise from the random parameters in Monte Carlo rendering.  
BITTERLI B et al. Nonlinearly weighted first-order regression for denoising Monte Carlo renderings.*

# Problem - Loss of Detail



Rendered Image



Denoised Image



Reference

# Problem - Loss of Detail



Rendered Image



Denoised Image



Reference




# Problem - Non-converging



Rendered Image



Denoised Image



# Method

**Goal:** To fix two major problems by mixing parameter  $\alpha$

Problems:

1. Loss of detail
2. Non - converging

# Method

$(1-\alpha)$



Rendered Image

+  $\alpha \cdot$




Denoised Image

= Good Image!

# Method

$(1-\alpha)$




Rendered Image

+  $\alpha \cdot$

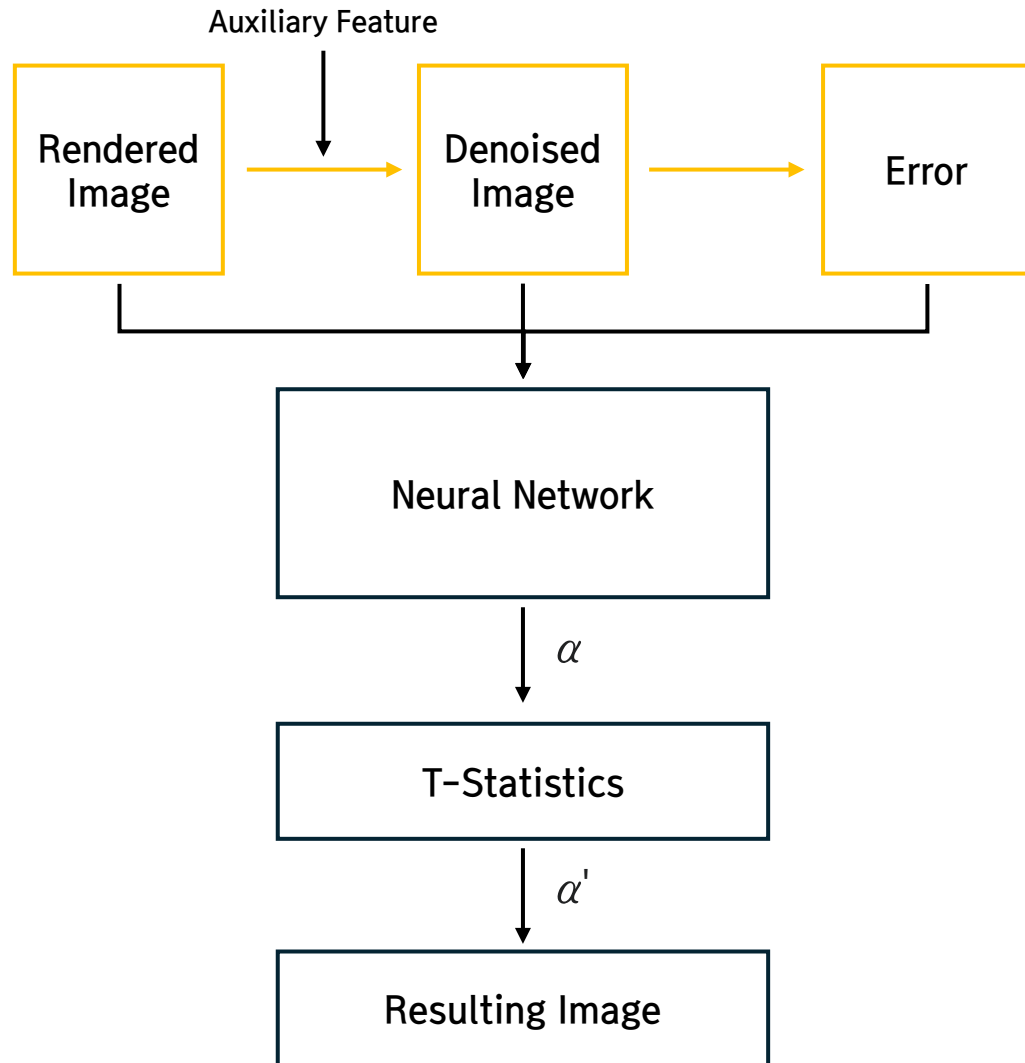


Denoised Image



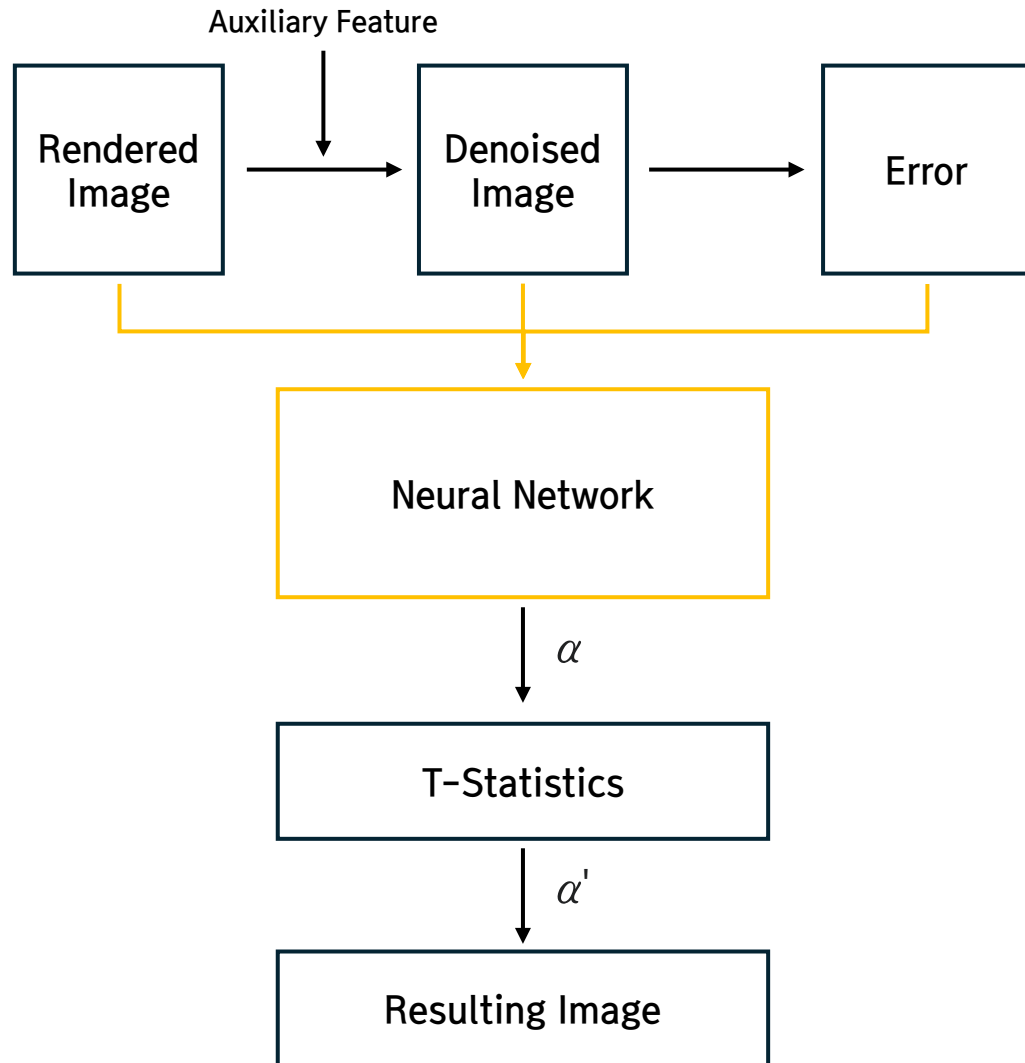
Resulting  $\alpha$

# Pipeline



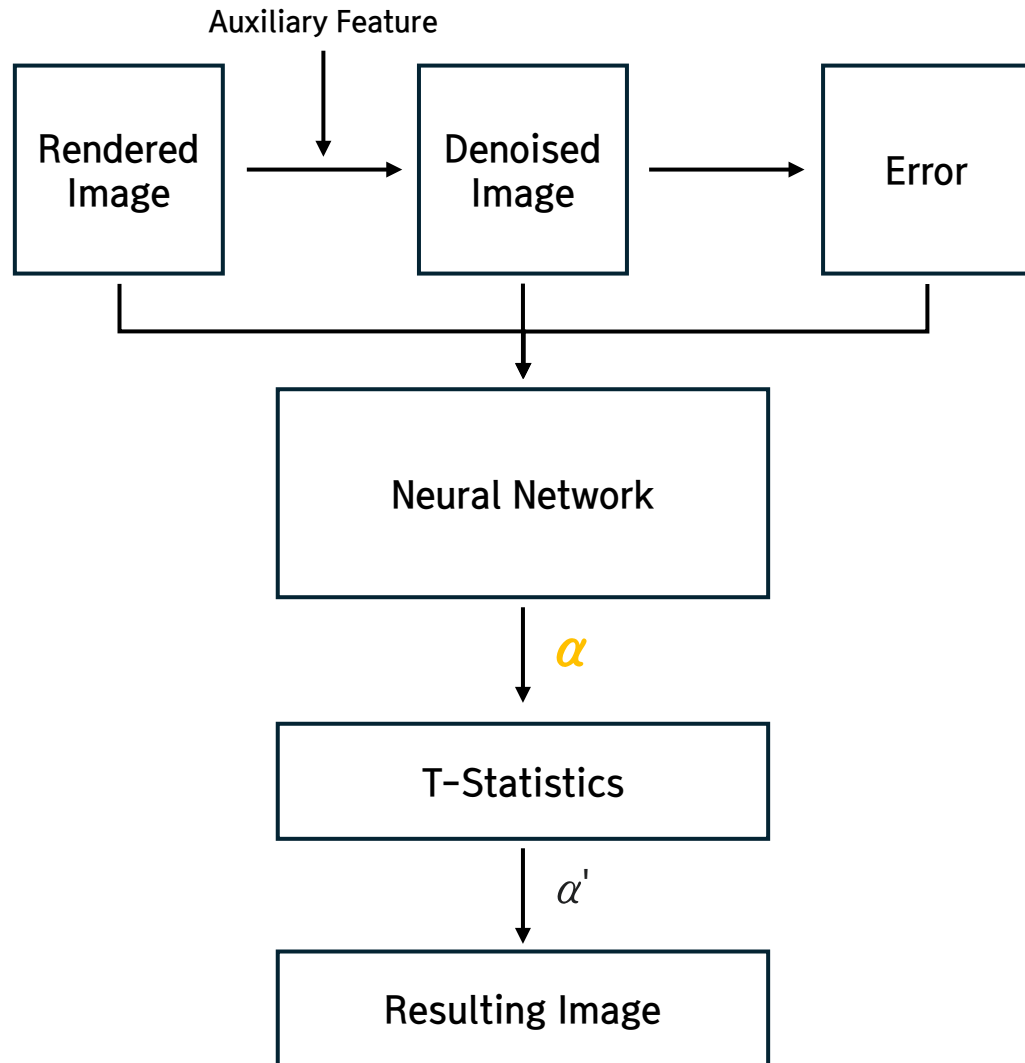
1. Generate denoised image from rendered image. Calculate error from denoised image.
2. Feed rendered image, denoised image and error to neural network.
3. Receive  $\alpha$  as output.
4. Rescale  $\alpha$  with t-statistics.
5. Generate resulting image.

# Pipeline



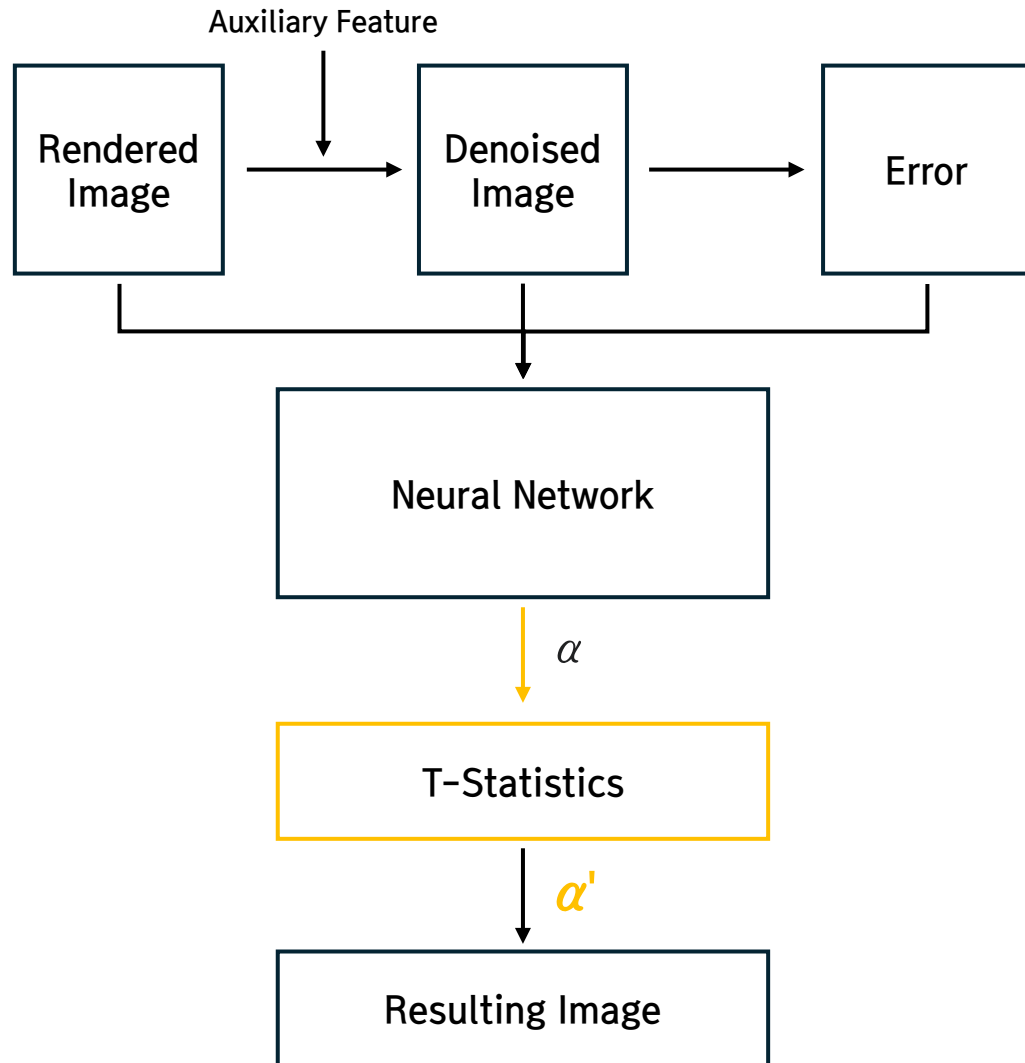
1. Generate denoised image from rendered image. Calculate error from denoised image.
2. Feed rendered image, denoised image and error to neural network.
3. Receive  $\alpha$  as output.
4. Rescale  $\alpha$  with t-statistics.
5. Generate resulting image.

# Pipeline



1. Generate denoised image from rendered image. Calculate error from denoised image.
2. Feed rendered image, denoised image and error to neural network.
3. Receive  $\alpha$  as output.
4. Rescale  $\alpha$  with t-statistics.
5. Generate resulting image.

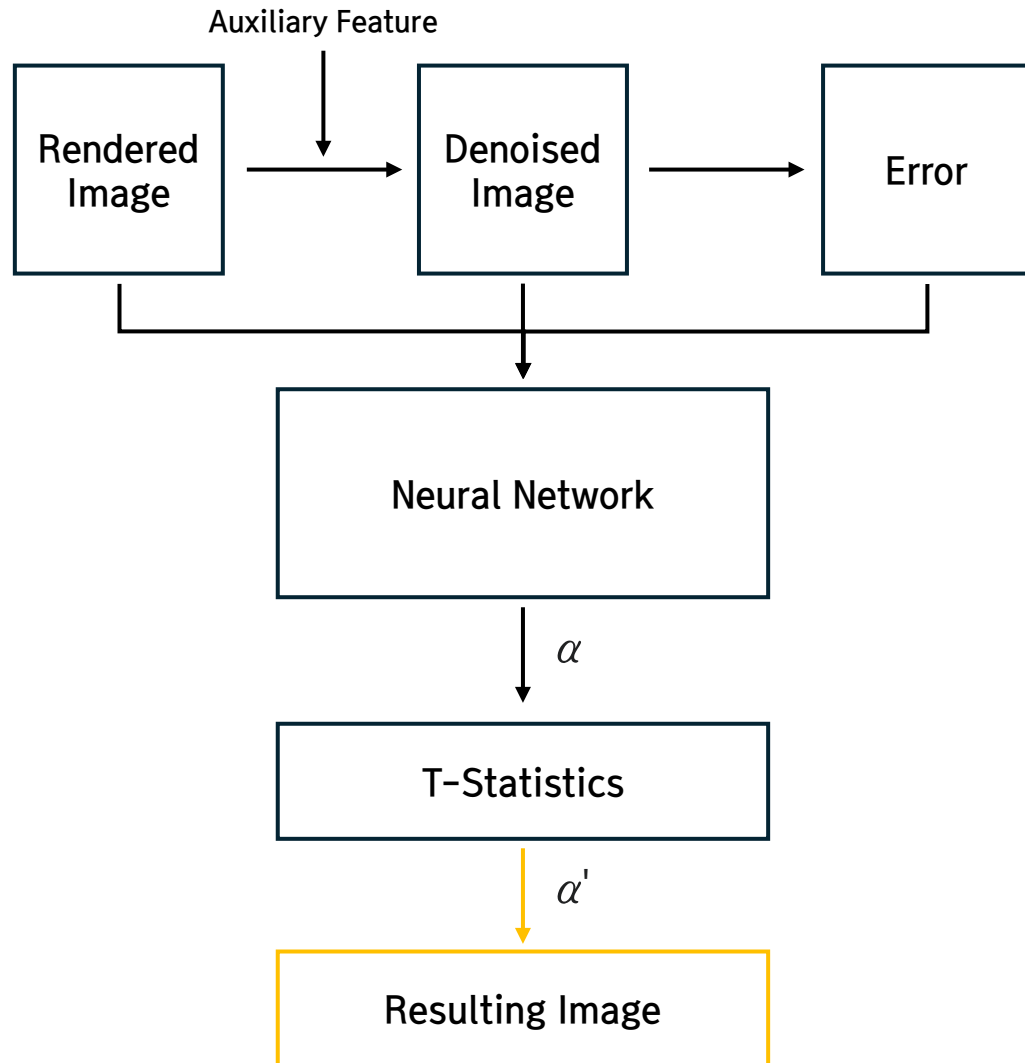
# Pipeline



1. Generate denoised image from rendered image. Calculate error from denoised image.
2. Feed rendered image, denoised image and error to neural network.
3. Receive  $\alpha$  as output.
4. Rescale  $\alpha$  with t-statistics.
5. Generate resulting image.



# Pipeline




1. Generate denoised image from rendered image. Calculate error from denoised image.
2. Feed rendered image, denoised image and error to neural network.
3. Receive  $\alpha$  as output.
4. Rescale  $\alpha$  with t-statistics.
5. Generate resulting image.


# Error Estimation - SURE

1. Generate denoised image from rendered image. Calculate **error** from denoised image.

$$\text{SURE}(F, x) = \frac{1}{d} \left( \|F(x) - x\|^2 + 2\text{tr}(J_F(x) \cdot \Sigma) - \text{tr}(\Sigma) \right)$$




# Error Estimation - SURE



**Figure 3:** *Per-pixel squared error estimate of a denoised image using SURE (left), and its actual squared error (right).*

# Input swapping

2. Feed rendered image, denoised image and error to neural network.
3. Receive  $\alpha$  as output.



**Figure 6:** Overview of our network architecture. The initial  $5 \times 5$  convolutional layer follows all but the final convolutional layer. The network also takes as input the rendered and denoised images, along with estimates of their error (see Fig. 5).

The network also takes as input the rendered and denoised images, along with estimates of their error (see Fig. 5) of Vogels et al. [VRM\*18].

# T-statistics


4. Rescale  $\alpha$  with t-statistics.

$$t_p = \frac{\bar{z}_p - \bar{x}_p}{\sqrt{\text{Var}[\bar{x}_p]} + \epsilon}$$

$\bar{x}_p$ : averages around pixel p in rendered image

$\bar{z}_p$ : averages around pixel p in mixed image

Large  $t_p \rightarrow$  Decrease alpha



# Results




Denoised Image




Mixed Image

# Results

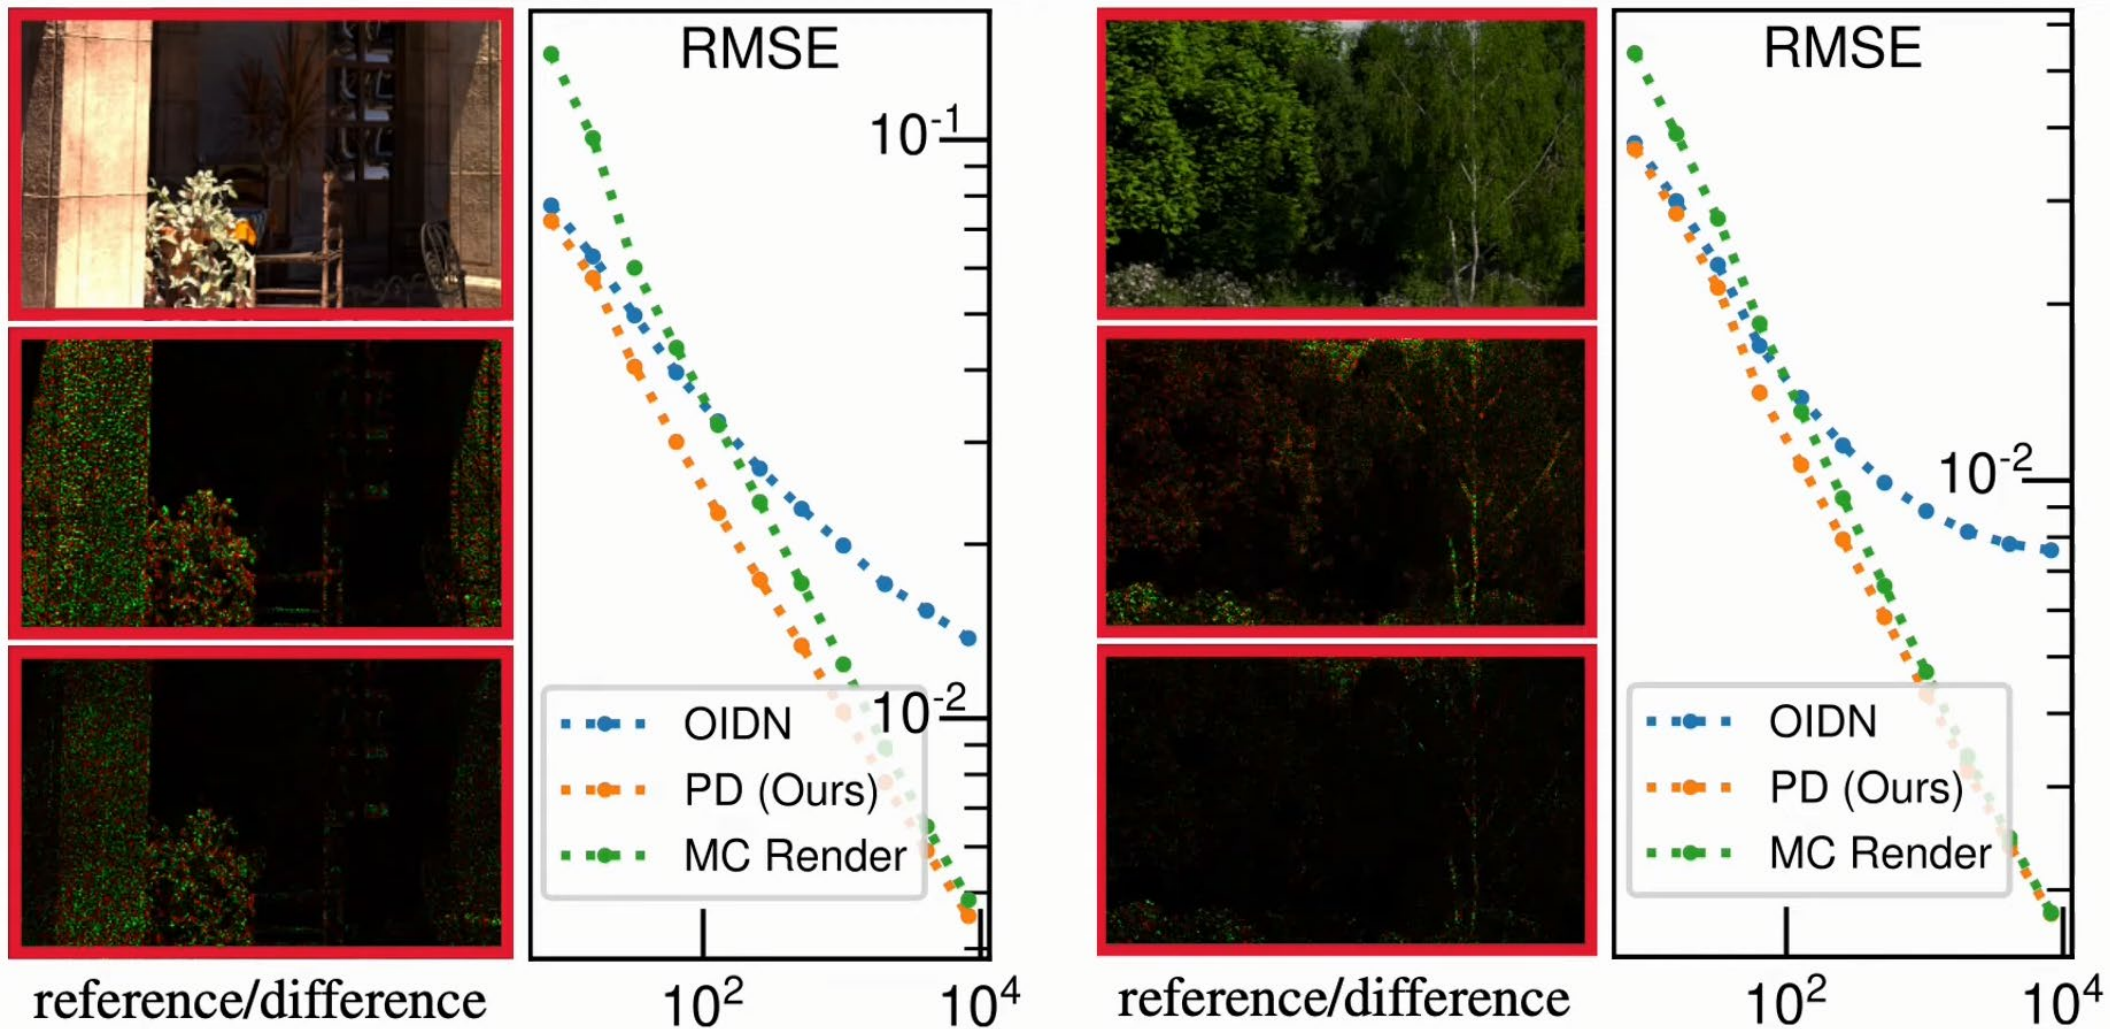


Denoised Image



Mixed Image

# Results

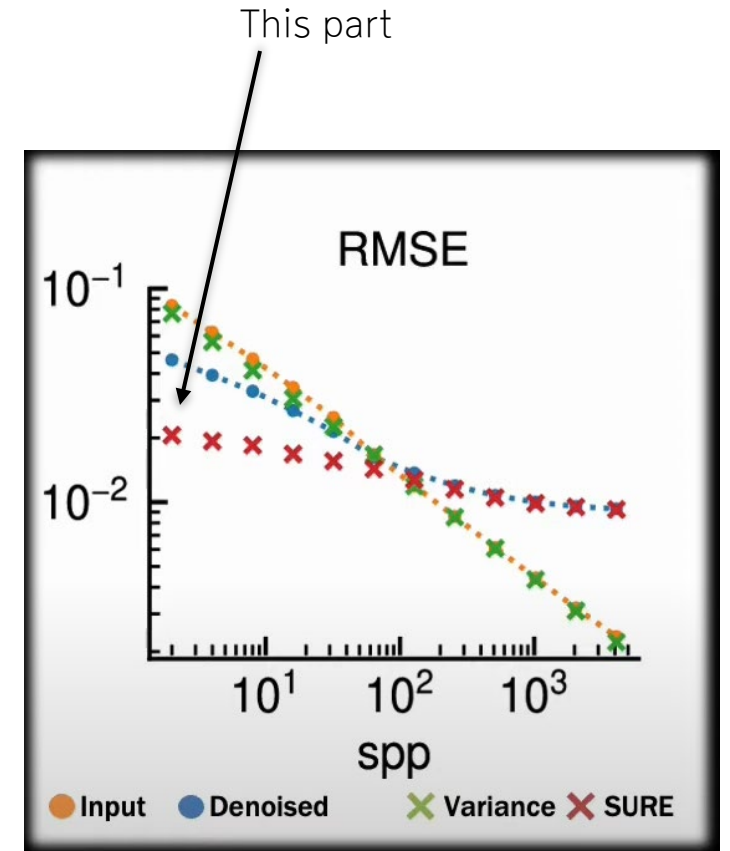




# Limitations



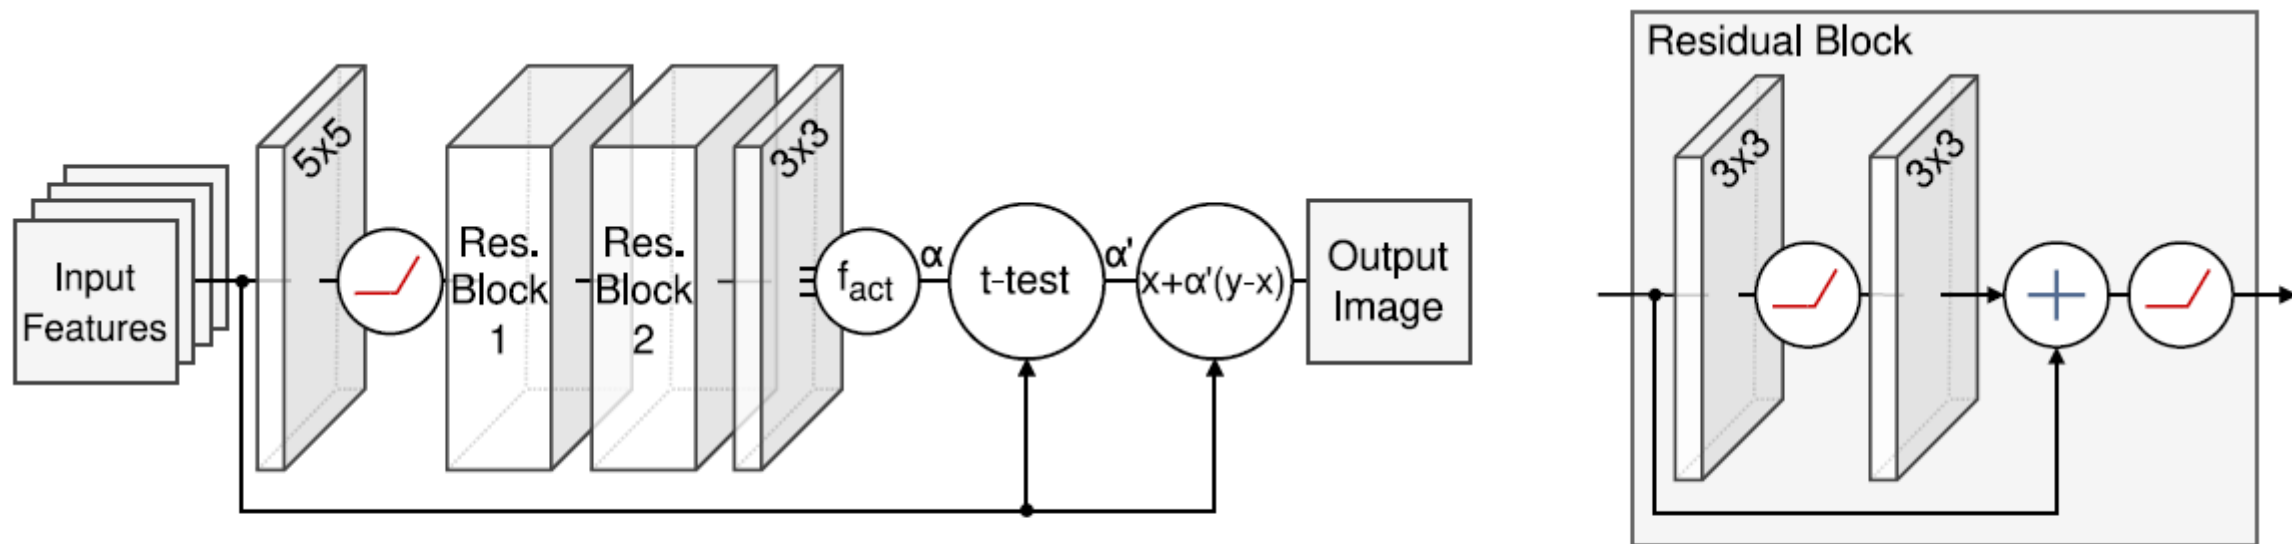
**Figure 13:** *Limitation of our method at very low sample counts, 2spp in this example, arising from insufficiently accurate sample variance estimates.*



Quiz! :D







**Figure 6:** Overview of our network  $h_{NN}(x, y, \dots; \Theta_h)$  for mixing an unbiased MC rendering  $x$  with its denoised counterpart  $y$  using per-pixel mixing parameters  $\alpha^l$ . The initial  $5 \times 5$  convolutional layer is followed by two residual blocks and a final convolutional layer. ReLU activation follows all but the final convolution, which uses  $f_{\text{act}}$  (see Sections 4.1 and 4.3). Input features include the rendered and denoised images, along with estimates of their error (Section 4.3). The use of residual blocks is inspired by the networks of Vogels et al. [VRM\*18].

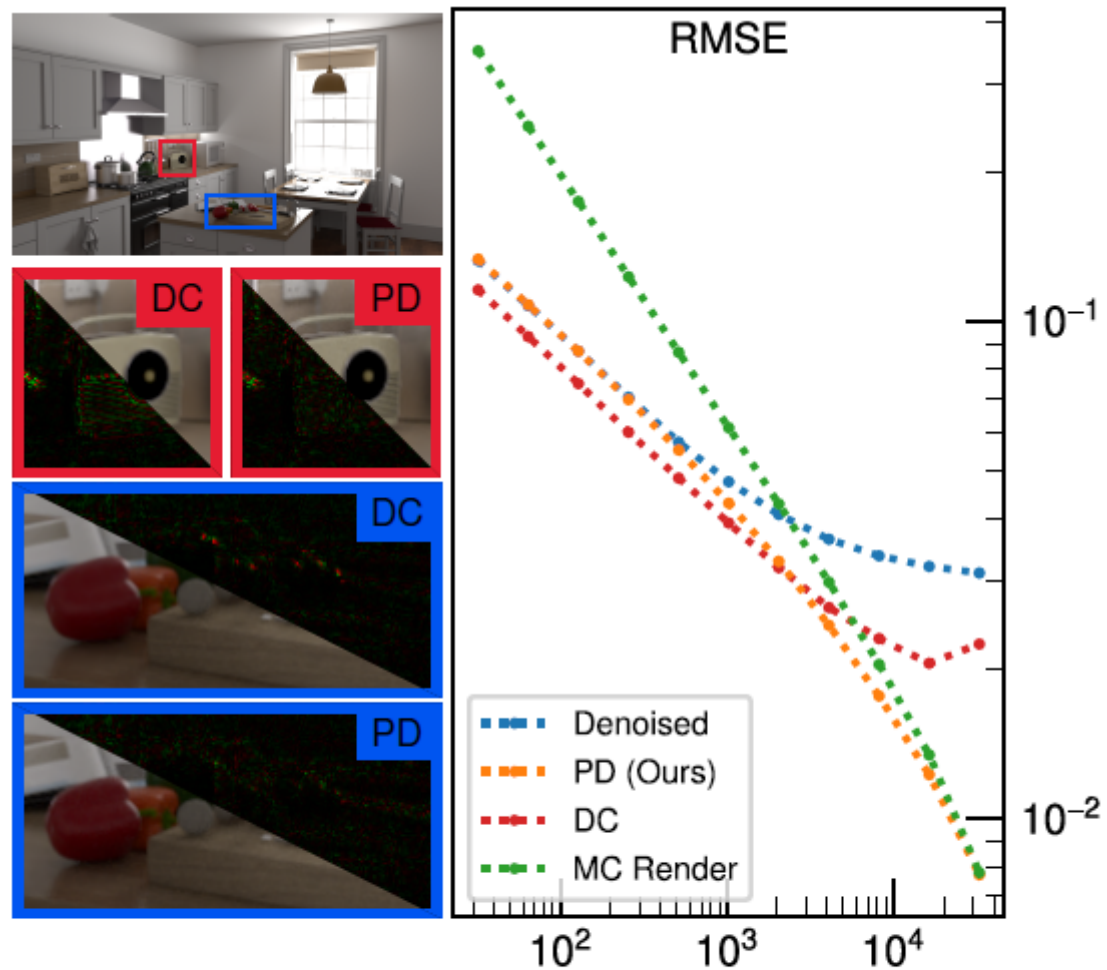
$$\bar{x}_p = x_{M_p} + \sum_{q \in \mathcal{N}_p} \kappa_{p,q} x_q, \quad (12)$$

$$\text{Var}[\bar{x}_p] = \text{Var}[x_{M_p}] + \sum_{q \in \mathcal{N}_p} \kappa_{p,q}^2 \text{Var}[x_q] \quad (13)$$

$$\bar{z}_p = z_{M_p} + \sum_{q \in \mathcal{N}_p} \kappa_{p,q} (x_q + \alpha_q (y_q - x_q)) \quad (14)$$

$$\kappa_{p,q} = \frac{1}{\sigma_s \sqrt{2\pi}} \exp\left(-\frac{\|p - q\|^2}{2\sigma_s^2}\right) \quad (15)$$

$$M_p = \arg \max_{q \in \mathcal{N}_p} \text{Var}[x_q]. \quad (16)$$



**Figure 11:** Comparison of our method, Progressive Denoising (PD), with Deep Combiner (DC), kitchen test scene. While DC continually improves upon the denoised image, our method performs the best as the quality of the MC rendered input increases.

inputs	model	error		
		RMSE	SMAPE	FLIP
HDR	den	0.1964	0.0344	0.0816
	pro-den	<b>0.0784</b>	<b>0.0284</b>	<b>0.0726</b>
HDR, VAR	den	0.1022	0.0292	0.0736
	pro-den	<b>0.0587</b>	<b>0.0277</b>	<b>0.0713</b>
HDR, ALB, NRM	den	0.1247	0.0393	0.0947
	pro-den	<b>0.0523</b>	<b>0.0298</b>	<b>0.0789</b>
HDR, ALB, NRM, VAR	den	0.0971	0.0326	0.0838
	pro-den	<b>0.0536</b>	<b>0.0278</b>	<b>0.0742</b>
	mc-render	0.0841	0.0669	0.1049

**Table 1:** Comparison of our approach (pro-den) with simple denoising (den) for different input feature combinations.

inputs	model	error		
		RMSE	SMAPE	FLIP
HDR, ALB, NRM, VAR	den-kpcn	0.1342	0.0362	0.0990
	pro-den	<b>0.0640</b>	<b>0.0285</b>	<b>0.0788</b>
	mc-render	0.0841	0.0669	0.1049

**Table 2:** Our approach (pro-den) applied to a denoiser based on a kernel predicting network (den-kpcn). The improvement is similar as to when applied to the U-Net based denoisers of Table 1.

inputs	model	error		
		RMSE	SMAPE	FLIP
HDR	oidn	0.0499	0.0278	0.0735
	pro-den	<b>0.0394</b>	<b>0.0251</b>	<b>0.0685</b>
HDR, ALB, NRM	oidn	0.0586	0.0265	0.0743
	pro-den	<b>0.0489</b>	<b>0.0248</b>	<b>0.0705</b>
	mc-render	0.0841	0.0669	0.1049

**Table 3:** Applying our method (pro-den) to a pre-trained denoiser, Intel Open Image Denoise (oidn). Despite the already high-quality of OIDN, our method is still able to lower the overall error.




Fig. 13. The ED, which takes a pair of unbiased and biased images, i.e., ED (PT, KPCN) and ED (PT, AFGSA), shows much higher errors than the other biased results, including ours. This MSE-based method can be more robust when it takes only biased inputs, i.e., ED (NFOR, KPCN) and ED (NFOR, AFGSA), but produces higher errors than its input NFOR for the DRAGON (from 256 to 2K spp) and CURLY-HAIR (from 16 to 128 spp and 2K spp) scenes. Our technique, however, robustly improves our input denoisers and shows the best errors over the tested ranges.