


# NeRF Revisited: Fixing Quadrature Instability in Volume Rendering

Mikaela Angelina Uy<sup>1</sup>, George Nakayama<sup>1</sup>, Guandao Yang<sup>1,2</sup>,  
Rahul Thomas<sup>1</sup>, Leonidas Guibas<sup>1</sup>, Ke Li<sup>3,4</sup>



Stanford University<sup>1</sup>



Cornell University<sup>2</sup>



Simon Fraser University<sup>3</sup>

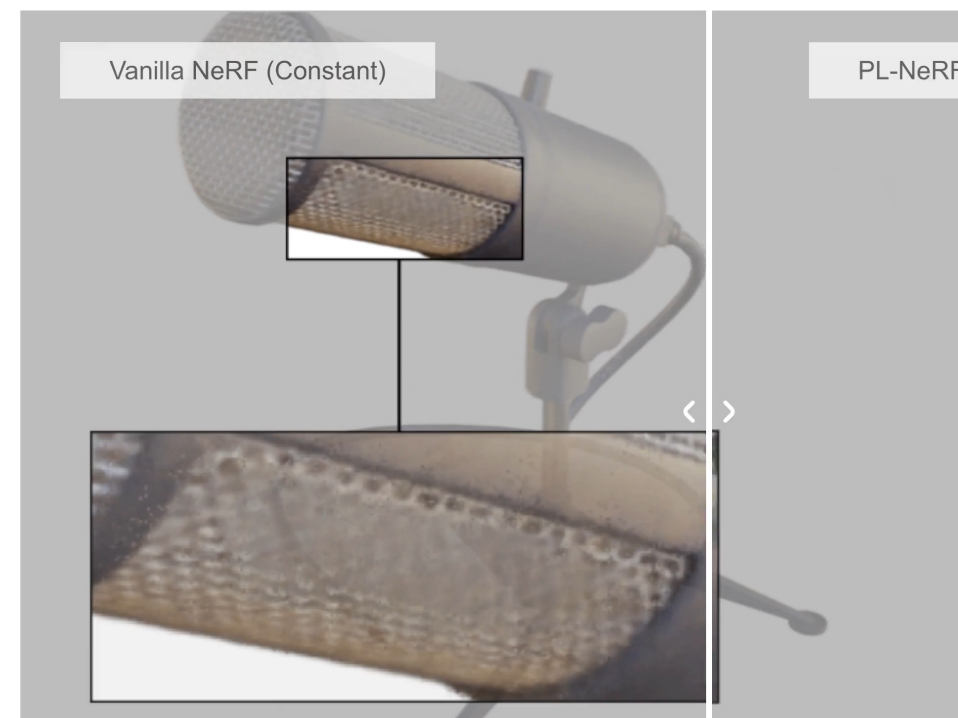


Google<sup>4</sup>

Vanilla NeRF




PL-NeRF (Ours)



Observe that the structure inside the mic is lost in piecewise constant opacity approximation.

# NeRF: Neural Radiance Fields


- Novel view synthesis given only input images with known poses



- Input: 5D continuous coordinate
- Output: volume density (1D), color (rgb)

# NeRF: Neural Radiance Fields

Our focus!



- We dive into the actual **volume rendering equation**.

# What We ~~Are~~ Were Used to

- The formula that we are all used takes the following form:

$$\hat{C}(\mathbf{r}) = \sum_{i=1}^N T_i (1 - \exp(-\sigma_i \delta_i)) \mathbf{c}_i, \text{ where } T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j \delta_j\right),$$

- This comes from the continuous integral:


$$\hat{y} = \mathbb{E}_{s \sim p(s)} [c(s)] = \int_0^{\infty} \underline{p(s)} c(s) ds$$

$$C(\mathbf{r}) = \int_{t_n}^{t_f} T(t) \sigma(\mathbf{r}(t)) \mathbf{c}(\mathbf{r}(t), \mathbf{d}) dt, \text{ where } T(t) = \exp\left(-\int_{t_n}^t \sigma(\mathbf{r}(s)) ds\right).$$

# Where did it come from?

- Continuous integral:

$$C(\mathbf{r}) = \int_{t_n}^{t_f} T(t)\sigma(\mathbf{r}(t))\mathbf{c}(\mathbf{r}(t), \mathbf{d})dt, \text{ where } T(t) = \exp\left(-\int_{t_n}^t \sigma(\mathbf{r}(s))ds\right).$$




- In practice, it is approximated with **quadrature**, resulting in the expression we are used to.

$s_1, s_2, \dots, s_N$  be  $N$  (ordered) samples

$$\forall s \in [s_j, s_{j+1}], \tau(s) = \tau(s_j),$$



$$\hat{C}(\mathbf{r}) = \sum_{i=1}^N T_i(1 - \exp(-\sigma_i\delta_i))\mathbf{c}_i, \text{ where } T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j\delta_j\right),$$




This is the **exact integral** under the **piecewise constant** opacity approximation!



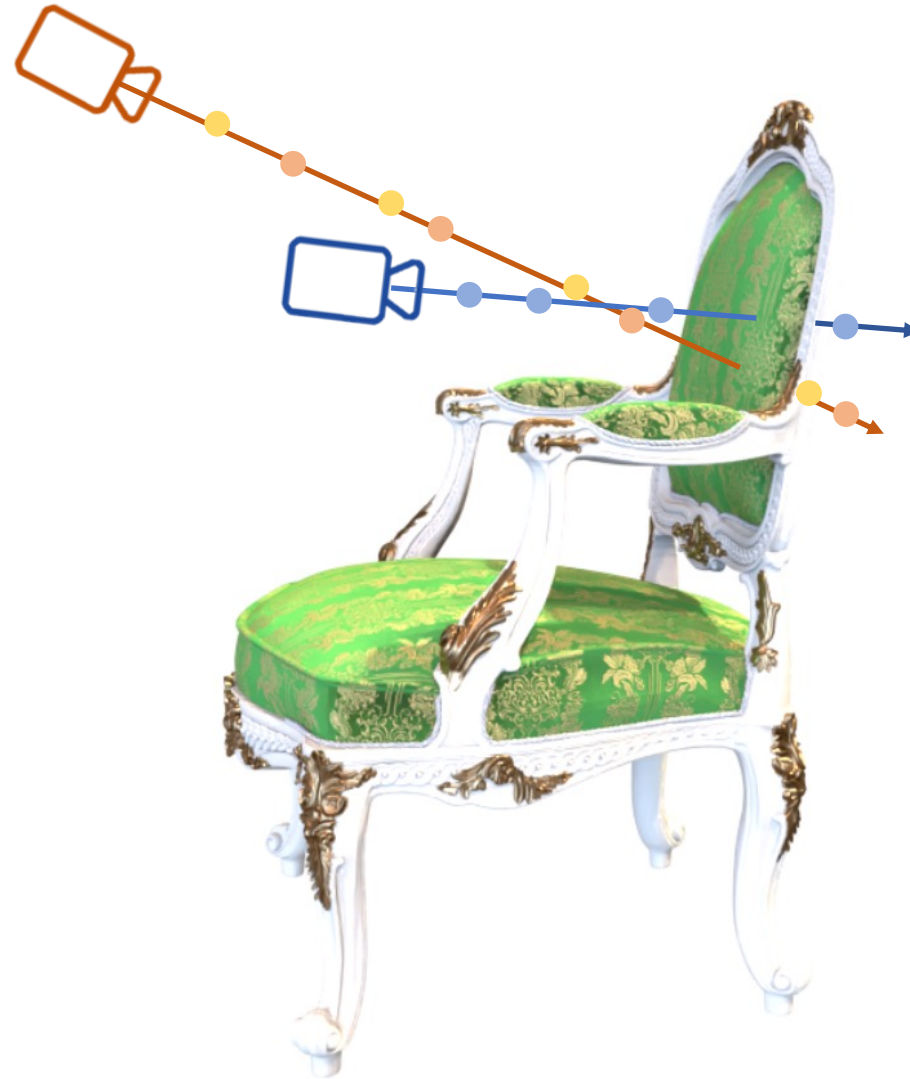
# NeRFs with Piecewise Constant Opacity

- Different from the classical works, the difference with the emergence of NeRFs is **optimization**.
  - Opacity and color are **learned and trained** using the volume rendering equation.
  - Instead of just being **evaluated** as in classical works.
- Note: for each set of samples along the ray, the opacity at the **left bin** is assigned for each interval.

$$\forall s \in [s_j, s_{j+1}], \tau(s) = \tau(s_j),$$



# Quadrature Instability




$$\forall s \in [s_j, s_{j+1}], \tau(s) = \tau(s_j),$$

# Problems with Quadrature Instability


- Ray Conflicts

$$\forall s \in [s_j, s_{j+1}], \tau(s) = \tau(s_j),$$

Piecewise Constant Opacity



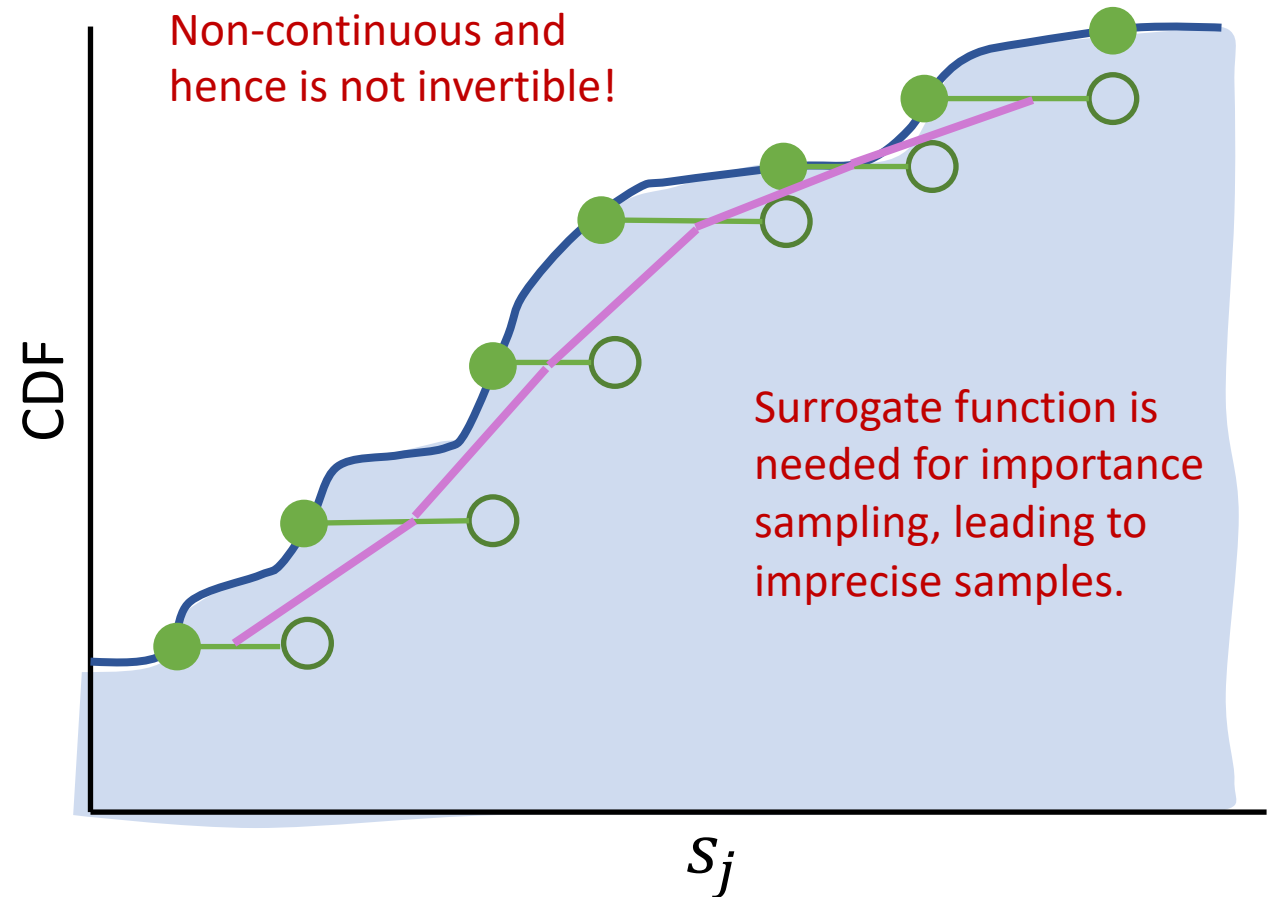

Rendered Test View



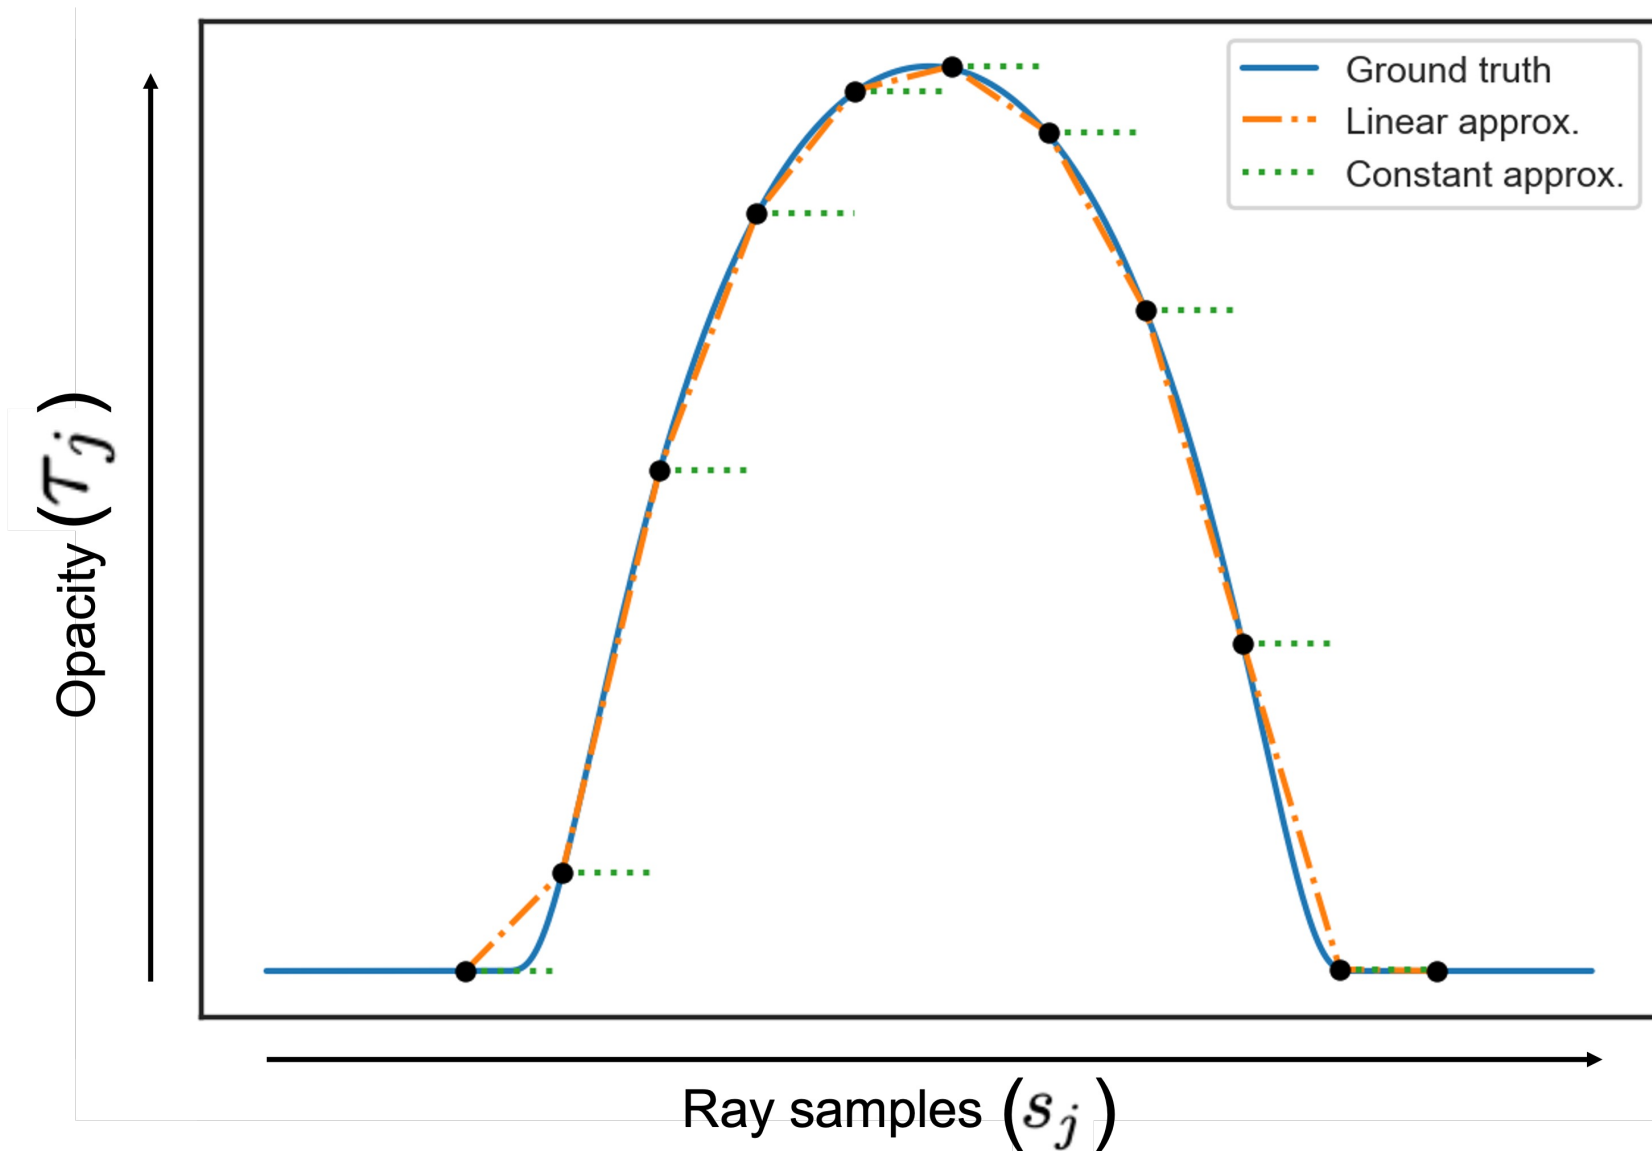


# Problems with Quadrature Instability

- Non-invertibility of the CDF



# Let's do better!



# General Form for $P_j$

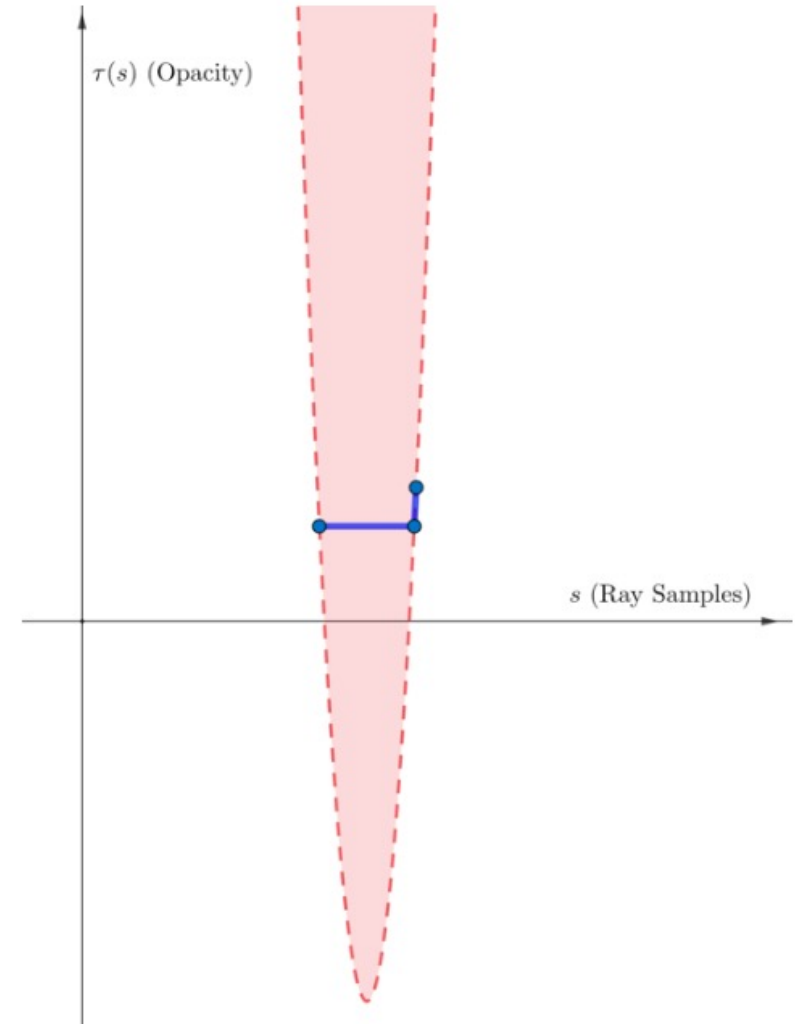
- We first showed that the probability of an interval can be exactly evaluated iff transmittance ( $T$ ) is in closed-form:

$$P_j = \int_{s_j}^{s_{j+1}} \tau(s)T(s) ds = - \int_{s_j}^{s_{j+1}} T'(s) ds = T(s_j) - T(s_{j+1}).$$

- We know that this holds for any **piecewise polynomial** approximation for opacity.

# Quadratic and Higher Order Polynomials


- Degree = 0 results in the original piecewise constant assumption with quadrature instability.
- We further show that for degree  $\geq 2$ , it would lead to **poor numerical conditioning**.
  - Interpolating a higher degree polynomial can give negative opacity values when samples are close.



# Our Piecewise Linear Derivation for Opacity

- We use a quadrature under a **piecewise linear** assumption for opacity.
  - For  $s \in [s_j, s_{j+1}]$ , where  $\tau_j = \tau(s_j)$ ,  $\tau(s_{j+1}) = \tau(s_{j+1})$

$$\tau(s) = \left( \frac{s_{j+1} - s}{s_{j+1} - s_j} \right) \tau_j + \left( \frac{s - s_j}{s_{j+1} - s_j} \right) \tau_{j+1}.$$



# Our Piecewise Linear Derivation for Opacity

- Under this assumption, we get the following **simple and closed-form** expressions:


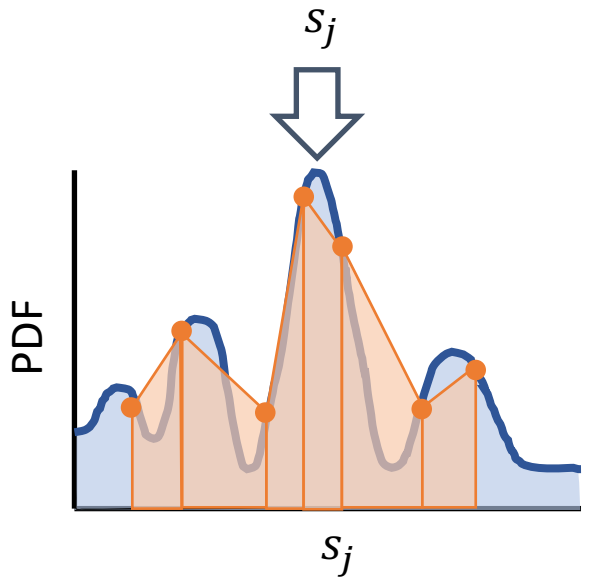

$$P_j = T(s_j) \cdot \left( 1 - \exp \left[ -\frac{(\tau_{j+1} + \tau_j)(s_{j+1} - s_j)}{2} \right] \right).$$

$$T(s_j) = \prod_{k=1}^i \exp \left[ -\frac{(\tau_k + \tau_{k-1})(s_k - s_{k-1})}{2} \right].$$



# Our Precise Importance Sampling

- The CDF is increasing and continuous, thus it is **invertible**!




# Invertibility

- The CDF is increasing and continuous, thus it is **invertible!**

$$F(t) = \int_0^t p(s) ds = \sum_{s_j < t} P_j + \int_{s_j}^t p(s) ds = \sum_{s_j < t} P_j + \int_{s_j}^t \tau(s) T(s) ds$$



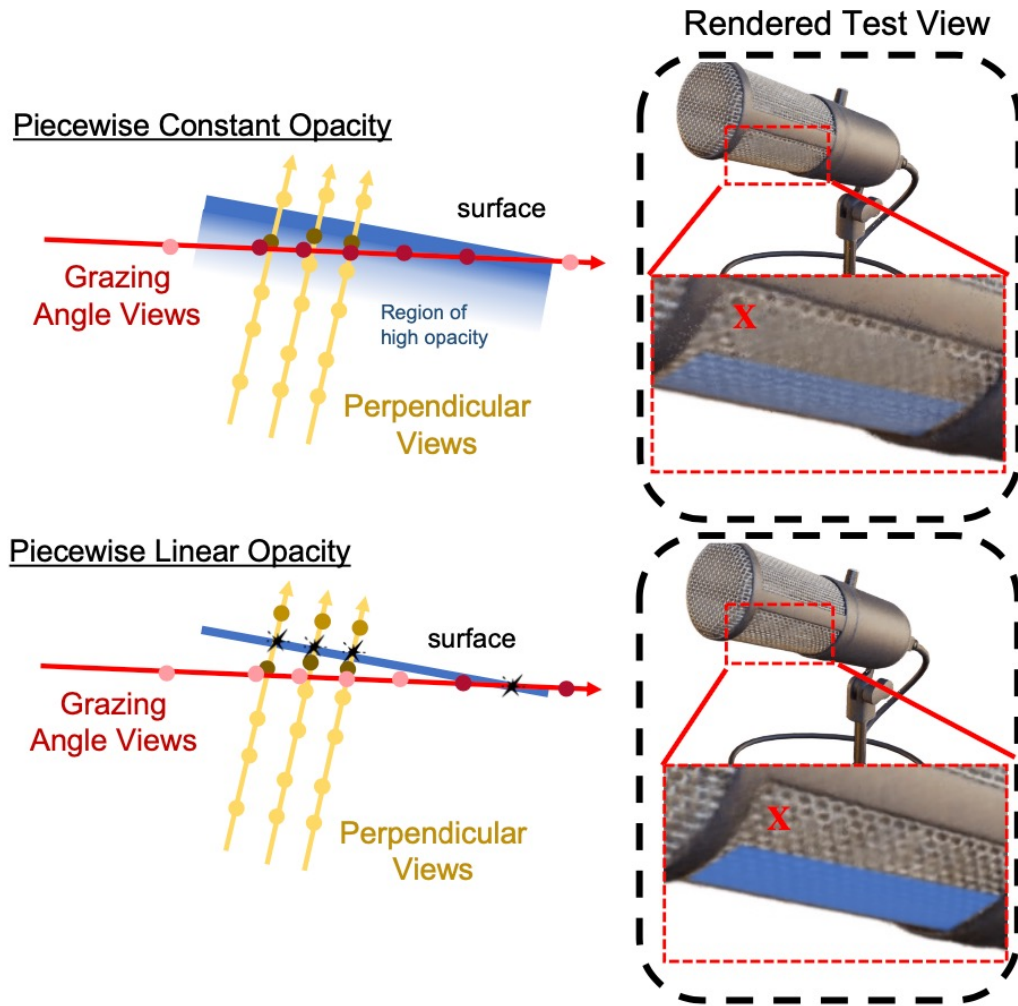
- We are also able to derive a closed-form solution for inverse transform sampling -- our **precise importance sampling.**

$$t = \frac{s_{k+1} - s_k}{\tau_{k+1} - \tau_k} \left[ -\tau_k + \sqrt{\tau_k^2 + \frac{2(\tau_{k+1} - \tau_k) \left( -\ln \frac{1-u}{T(s_k)} \right)}{(s_{k+1} - s_k)}} \right].$$



Which also leads to more effective supervision on samples, e.g. depth.

# Our Precise Importance Sampling




# Results: Alleviating Ray Conflicts

## **Ray Conflicts**

# Results: Less Fuzzy Surfaces

Vanilla NeRF (Constant)



PL-NeRF

# Results: Crisper Textures

Vanilla NeRF (Constant)



PL





# Results: Quantitative

<b>Blender</b>		Avg.	Chair	Drums	Ficus	Hotdog	Lego	Mat.	Mic	Ship
PSNR $\uparrow$	Const. (Vanilla)	30.61	32.54	24.79	29.63	36.08	32.01	29.31	32.55	27.95
	Linear (Ours)	<b>31.10</b>	<b>32.92</b>	<b>25.07</b>	<b>30.18</b>	<b>36.46</b>	<b>32.90</b>	<b>29.52</b>	<b>33.08</b>	<b>28.71</b>
SSIM $\uparrow$	Const. (Vanilla)	0.943	0.966	0.918	0.960	0.975	0.959	0.943	0.978	0.846
	Linear (Ours)	<b>0.948</b>	<b>0.969</b>	<b>0.923</b>	<b>0.965</b>	<b>0.977</b>	<b>0.966</b>	<b>0.948</b>	<b>0.981</b>	<b>0.857</b>
LPIPS $\downarrow$	Const. (Vanilla)	5.17	3.19	7.97	4.14	2.48	2.33	4.32	2.16	14.8
	Linear (Ours)	<b>4.39</b>	<b>2.85</b>	<b>7.10</b>	<b>3.03</b>	<b>2.28</b>	<b>1.81</b>	<b>3.21</b>	<b>1.73</b>	<b>13.1</b>
<b>RFF</b>		Avg.	Fern	Flower	Fortress	Horns	Leaves	Orchid	Room	Trex
PSNR $\uparrow$	Const. (Vanilla)	27.53	26.79	28.23	32.53	28.54	22.35	21.20	33.03	27.58
	Linear (Ours)	<b>28.05</b>	<b>26.85</b>	<b>28.71</b>	<b>32.95</b>	<b>29.38</b>	<b>22.51</b>	<b>21.25</b>	<b>33.99</b>	<b>28.79</b>
SSIM $\uparrow$	Const. (Vanilla)	0.874	0.746	0.886	0.925	0.893	0.816	0.746	0.956	0.916
	Linear (Ours)	<b>0.885</b>	<b>0.863</b>	<b>0.902</b>	<b>0.932</b>	<b>0.911</b>	<b>0.826</b>	<b>0.754</b>	<b>0.961</b>	<b>0.933</b>
LPIPS $\downarrow$	Const. (Vanilla)	7.37	9.67	6.34	2.92	7.26	11.0	11.8	4.33	5.66
	Linear (Ours)	<b>6.06</b>	<b>7.92</b>	<b>4.93</b>	<b>2.46</b>	<b>5.51</b>	<b>9.59</b>	<b>10.2</b>	<b>3.54</b>	<b>4.38</b>

Table 1: **Quantitative Results on Blender and Real Forward Facing Datasets.** Reported LPIPS scores are multiplied by  $10^2$

# Results: Qualitative





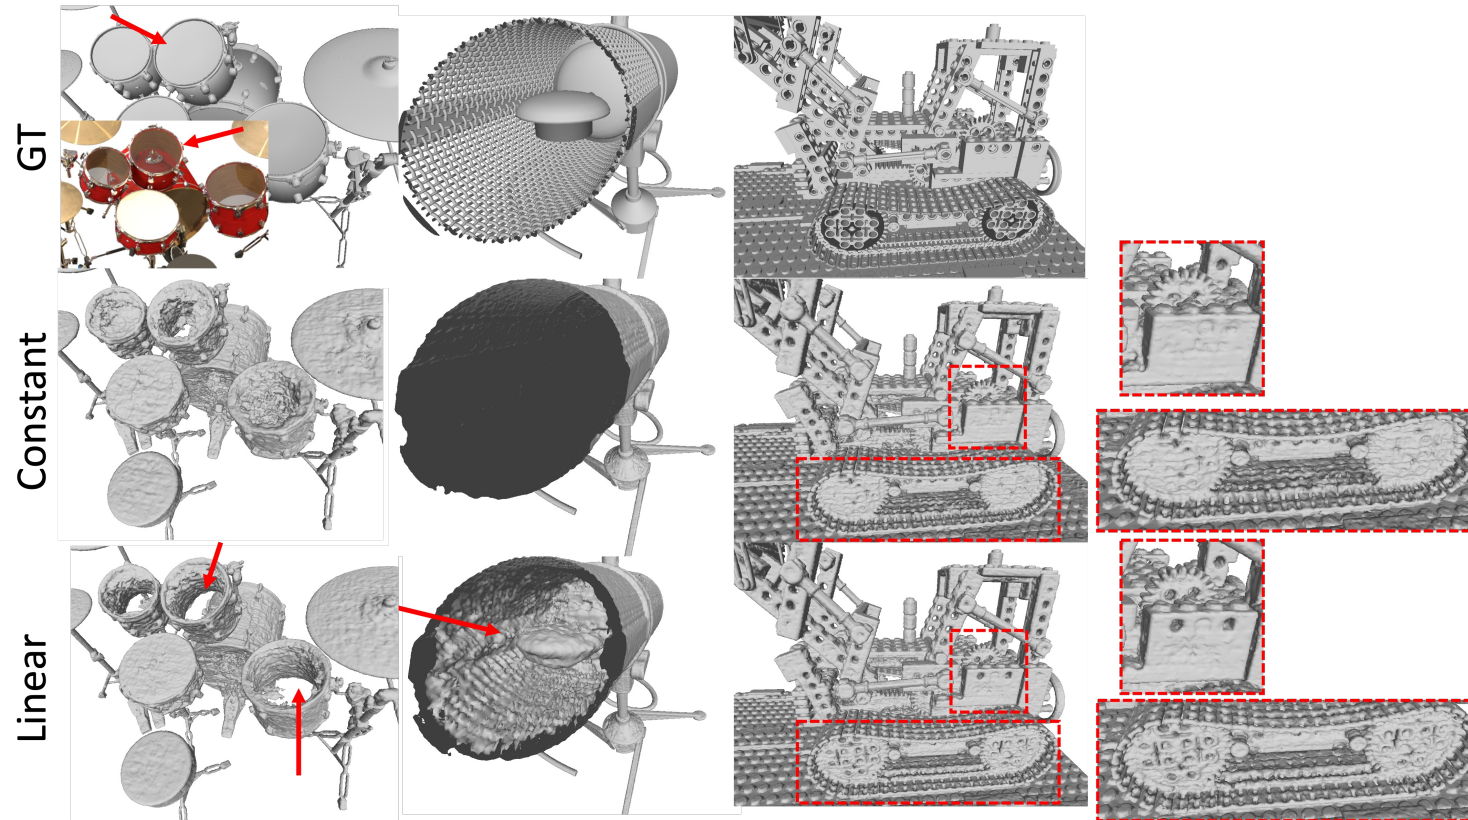
Figure 4: Qualitative Results for Blender and Real Forward Facing.

# Results: Better Geometry Extraction







# Results: Better Geometry Extraction



	<b>Blender</b>	Avg.	Chair	Drums	Ficus	Hotdog	Lego	Mat.	Mic	Ship
CD↓	Vanilla NeRF	10.43	5.162	<b>6.842</b>	29.94	7.555	7.474	6.833	5.214	11.44
	PL-NeRF	<b>10.10</b>	<b>4.676</b>	7.754	<b>29.58</b>	<b>7.004</b>	<b>6.825</b>	<b>6.061</b>	<b>5.213</b>	<b>10.44</b>

Table 3: **Geometry Extraction Error** Distance between the surface of the GT to the predicted meshes. Scores are  $\times 10^3$

# Drop-in Replacement in Existing Methods



GT

Mip-NeRF

PL-MipNeRF

# Drop-in Replacement in Existing Methods

<b>Blender</b>		Avg.	Chair	Drums	Ficus	Hotdog	Lego	Mat.	Mic	Ship
PSNR $\uparrow$	Mip-NeRF	31.76	33.95	24.39	31.20	36.12	33.84	30.55	34.63	29.41
	PL-MipNeRF	<b>32.48</b>	<b>35.11</b>	<b>24.92</b>	<b>32.25</b>	<b>36.51</b>	<b>35.15</b>	<b>30.69</b>	<b>35.22</b>	<b>30.00</b>
SSIM $\uparrow$	Mip-NeRF	0.955	0.975	0.921	0.971	0.978	0.971	0.957	0.987	0.876
	PL-MipNeRF	<b>0.959</b>	<b>0.981</b>	<b>0.928</b>	<b>0.977</b>	<b>0.980</b>	<b>0.976</b>	<b>0.959</b>	<b>0.989</b>	<b>0.882</b>
LPIPS $\downarrow$	Mip-NeRF	3.64	1.80	6.82	2.35	1.97	1.44	2.39	0.973	11.4
	PL-MipNeRF	<b>3.09</b>	<b>1.32</b>	<b>5.78</b>	<b>1.66</b>	<b>1.67</b>	<b>1.07</b>	<b>2.09</b>	<b>0.788</b>	<b>10.3</b>

Table 1: **Quantitative Results of Mip-NeRF v.s. PL-MipNeRF** Reported LPIPS scores are multiplied by  $10^2$

<b>Blender</b>		Avg.	Chair	Drums	Ficus	Hotdog	Lego	Mat.	Mic	Ship
PSNR $\uparrow$	DIVeR	30.78	32.01	<b>24.72</b>	30.1	<b>35.94</b>	29.03	29.31	32.10	<b>29.08</b>
	PL-DIVeR	<b>30.88</b>	<b>32.92</b>	24.7	<b>30.23</b>	<b>35.94</b>	<b>33.42</b>	<b>32.06</b>	<b>33.08</b>	28.99
SSIM $\uparrow$	DIVeR	0.956	0.959	<b>0.917</b>	<b>0.963</b>	0.974	0.965	<b>0.977</b>	0.978	0.870
	PL-DIVeR	<b>0.947</b>	<b>0.969</b>	0.916	<b>0.963</b>	<b>0.966</b>	<b>0.966</b>	<b>0.977</b>	<b>0.981</b>	<b>0.871</b>
LPIPS $\downarrow$	DIVeR	3.39	2.79	6.13	2.34	1.92	<b>1.46</b>	<b>1.77</b>	2.16	<b>7.77</b>
	PL-DIVeR	<b>3.28</b>	<b>2.85</b>	<b>6.01</b>	<b>2.12</b>	<b>1.83</b>	1.49	<b>1.77</b>	<b>1.73</b>	7.82

Table 2: **Quantitative Results of DIVeR v.s. PL-DIVeR** Reported LPIPS scores are multiplied by  $10^2$



# Thank you!



Visit our  
Project Page!

