


# ARC-NeRF: Area Ray Casting for Broader Unseen View Coverage in Few-shot Object Rendering

Seunghyeon Seo<sup>1</sup>   Yeonjin Chang<sup>1</sup>   Jayeon Yoo<sup>1</sup>   Seungwoo Lee<sup>1</sup>   Hojun Lee<sup>1,2</sup>   Nojun Kwak<sup>1</sup>

<sup>1</sup>Seoul National University   <sup>2</sup>Xperty Corp.




# Motivation




- NeRF suffers from a severe performance degradation with only a sparse set of inputs.




# Motivation




- Ray augmentation methods:
  - > More augmented rays  $\Rightarrow$  higher training cost
  - > Many potentially useful rays in unseen views remain unused.



(a) No aug.




(b) Patch rendering  
with aug. rays




(c) Flipped  
reflection rays


# Motivation




- Ray augmentation methods:
  - > More augmented rays  $\Rightarrow$  higher training cost
  - > Many potentially useful rays in unseen views remain unused.
  - >> **Area Ray** covers a broad area of unseen views.




(a) No aug.



(b) Patch rendering  
with aug. rays




(c) Flipped  
reflection rays

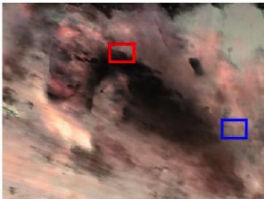

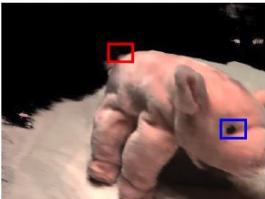
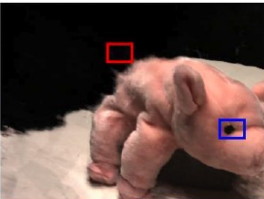
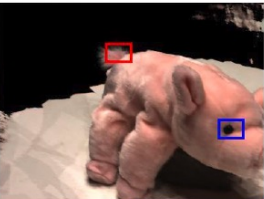
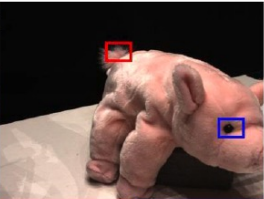




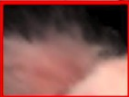


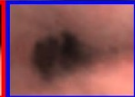

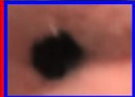




(d) Area Ray  
(Ours)



# Motivation




- **Area Ray:**
  - > Adaptive high-frequency regularization
  - > Achieves better training efficiency than brute-force ray augmentation

Mip-NeRF [1]	RegNeRF [25]	FlipNeRF [33]	FreeNeRF [49]	ARC-NeRF	Ground Truth
					
					
					
✗	✓ Patch rendering with aug. rays	✓ Flipped reflection rays	✗	✓ Area Ray	
✗	✗	✗	✓ Masking curriculum	✓ Adaptive reg. with IPE	

# Area Ray





(a) Metric distance reparameterization




(b) Area Ray featurization by IPE

- Reparameterize the metric distance  $t$  as  $\tilde{t}$  to derive the variance  $\tilde{\sigma}_\rho^2$ .

# Area Ray





(a) Metric distance reparameterization




(b) Area Ray featurization by IPE

- Reparameterize the metric distance  $t$  as  $\tilde{t}$  to derive the variance  $\tilde{\sigma}_\rho^2$ .
- Derive a base radius  $\tilde{\rho}$  of the Area Ray using the trigonometric function.

# Area Ray



(a) Metric distance reparameterization



(b) Area Ray featurization by IPE

- Area Ray:

$$\tilde{\mathbf{r}}(t) = \tilde{\mathbf{o}} + t\tilde{\mathbf{d}},$$


where  $\tilde{\mathbf{d}} = -\hat{\mathbf{n}}$  and  $\tilde{\mathbf{o}} = \mathbf{p}_s - t_s\tilde{\mathbf{d}}$ .

- High-frequency details are more tightly regularized in pixels with lower photo-consistency.  
=> Adaptive high-frequency reg.

[1]




Mean, variance of spatial region



Integrated positional encoding features




# Luminance Consistency Regularization



- Luminance maps offer a “free lunch” training signal that can be easily extracted from RGB images.

# Luminance Consistency Regularization




- Luminance maps offer a “free lunch” training signal that can be easily extracted from RGB images.

1) Compute the GT relative luminance of a target pixel:

$$y_{\text{GT}} = \sum_{\bar{c}}^{\{\bar{r}, \bar{g}, \bar{b}\}} \lambda_{\bar{c}} \bar{c}$$

where  $\bar{c} = c_{\text{GT}}^{2.2}$  is the linear RGB value obtained from gamma-compressed input via a simple power transformation.

# Luminance Consistency Regularization

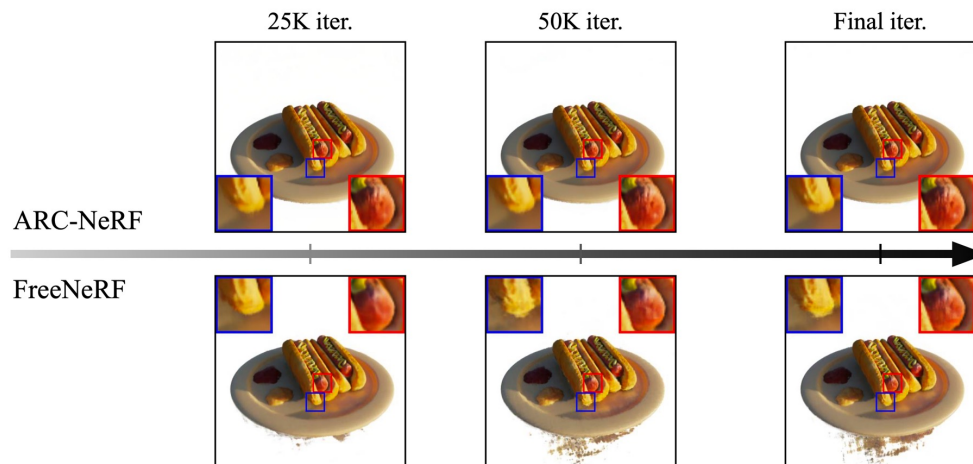



- Luminance maps offer a “free lunch” training signal that can be easily extracted from RGB images.

2) Predict the per-sample luminance  $y_i$  along a ray and aggregate the final luminance  $\hat{y}$  by volume rendering:



$$\hat{y}(\mathbf{r}) = \sum_{i=1}^N w_i y_i$$

# Analysis



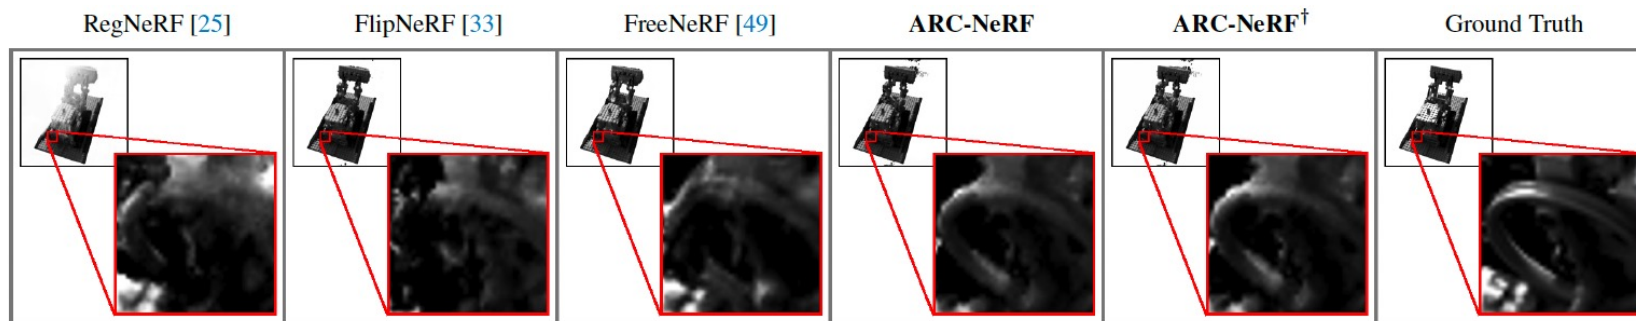

- FreeNeRF: hard masking of high-frequencies in early training
- Ours: adaptive frequency regularization based on ray-pixel consistency (angle-based)  
→ Sharper details at just 25K iters, even better than fully trained FreeNeRF

# Analysis



- Single Area Ray with adaptive regularization outperforms multicasting multiple augmented rays across the same region.
- ARC-NeRF improves rendering quality and training efficiency without the cost of processing extra rays.


# Analysis



- ARC-NeRF produces sharper relative luminance maps compared to other baselines.
- Luminance estimation provides additional supervision to blending weights via shared rendering pipeline.

	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Average $\downarrow$
RegNeRF [25]	5.71	0.780	0.291	0.368
FlipNeRF [33]	16.49	0.878	0.080	0.092
FreeNeRF [49]	15.63	0.869	0.091	0.113
<b>ARC-NeRF</b>	<u>16.87</u>	<u>0.886</u>	<u>0.074</u>	<u>0.088</u>
<b>ARC-NeRF<sup>†</sup></b>	<b>17.18</b>	<b>0.891</b>	<b>0.062</b>	<b>0.079</b>


# Analysis



	Area Ray	$\mathcal{L}_{\text{lum.}}$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Avg. $\downarrow$
FlipNeRF [33]			<u>19.55</u>	0.767	0.180	0.101
(1)	✓		19.51	<b>0.774</b>	<u>0.147</u>	<u>0.097</u>
(2)		✓	18.44	0.747	0.201	0.119
(3)	✓	✓	<b>19.85</b>	<u>0.773</u>	<b>0.146</b>	<b>0.096</b>

- Replacing FlipNeRF's reflection rays with Area Ray improves SSIM and LPIPS.
- Using only  $L_{\text{lum.}}$  without Area Ray leads to degraded performance.
  - > View-dependent nature of reflection rays conflicts with  $L_{\text{lum.}}$ , causing mismatch.
- Combining Area Ray with  $L_{\text{lum.}}$  yields best results, especially in PSNR with sharper fine details.

# Results



## Realistic Synthetic 360

	PSNR $\uparrow$		SSIM $\uparrow$		LPIPS $\downarrow$		Avg. $\downarrow$	
	4-view	8-view	4-view	8-view	4-view	8-view	4-view	8-view
Mip-NeRF [1]	8.70	13.31	0.792	0.848	0.250	0.176	0.285	0.188
DietNeRF [9]	10.86	16.08	0.814	0.870	0.194	0.113	0.223	0.123
InfoNeRF [15]	13.65	16.74	0.834	0.865	0.134	0.094	0.139	0.095
RegNeRF [25]	7.24	13.47	0.795	0.856	0.292	0.158	0.318	0.177
MixNeRF [34]	16.13	19.31	0.863	0.902	0.099	0.058	0.101	0.065
FreeNeRF [49]	15.71	18.99	0.857	0.894	0.103	0.064	0.114	0.072
FlipNeRF [33]	16.47	19.54	0.866	0.903	0.091	0.057	0.095	0.062
<b>ARC-NeRF</b>	<b>16.86</b>	<b>20.29</b>	<b>0.873</b>	<b>0.910</b>	<b>0.084</b>	<b>0.052</b>	<b>0.091</b>	<b>0.057</b>

## Shiny Blender 4-view

	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Average $\downarrow$
Ref-NeRF [41]	17.10	0.821	0.190	0.142
FreeNeRF [49]	16.99	0.828	0.157	0.131
FlipNeRF [33]	18.14	0.847	0.141	0.109
<b>ARC-NeRF</b>	<b>18.68</b>	<b>0.851</b>	<b>0.141</b>	<b>0.107</b>

## DTU 3-view

	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Avg. $\downarrow$
Mip-NeRF [1]	8.68	0.571	0.353	0.323
3DGS [14]	14.18	0.628	0.301	0.191

### Pre-training.


PixelNeRF [50]	16.82	0.695	0.270	0.147
PixelNeRF <sup>†</sup> [50]	18.95	0.710	0.269	0.125
SRF [5]	15.32	0.671	0.304	0.171
SRF <sup>†</sup> [5]	15.68	0.698	0.281	0.162
MVSNerF [4]	18.63	0.769	0.197	0.113
MVSNerF <sup>†</sup> [4]	18.54	0.769	0.197	0.113

### Regularization.

DietNeRF [9]	11.85	0.633	0.314	0.243
RegNeRF [25]	18.89	0.745	0.190	0.112
MixNeRF [34]	18.95	0.744	0.203	0.113
SimpleNeRF [35]	16.25	0.751	0.249	0.143
DiffusioNeRF [46]	16.20	0.698	0.160	0.128
SparseNeRF [42]	19.55	0.769	0.201	0.102
FreeNeRF <sup>†</sup> [49]	19.92	0.781	0.125	0.086
FreeNeRF [49]	19.23	0.769	0.149	0.103
FlipNeRF [33]	19.55	0.767	0.180	0.101
SparseGS [47]	18.89	0.702	0.229	0.117
<b>ARC-NeRF</b>	<b>19.85</b>	<b>0.773</b>	<b>0.146</b>	<b>0.096</b>



# Results




Realistic Synthetic 360 4-view




Realistic Synthetic 360 4-view

# Results




3-view; Notable improvement in the detail of the tail.



6-view; Apple surface textures are more stably reconstructed across changing views.



9-view; Brick textures are also more consistently reproduced.



Thank you!