Anti-Aliasing in Neural Rendering

Wenbo Hu

Tencent AI Lab

What is aliasing?

- Aliasing occurs when the sampling rate is below the Nyquist frequency

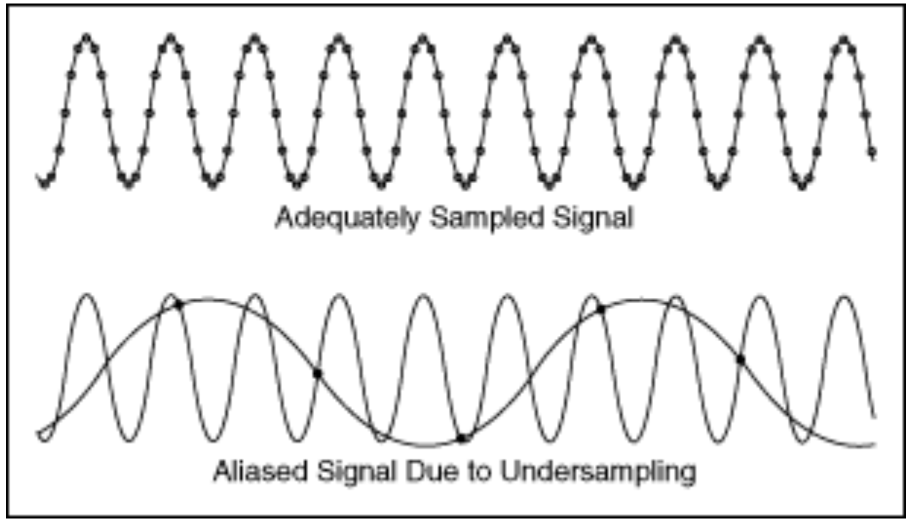

Factors for aliasing:

- Discrete sampling
- Sampling rate is low
- Repetitive pattern

Principles for anti-aliasing

- Area-sampling (Pre-filtering)
- Multi-sampling

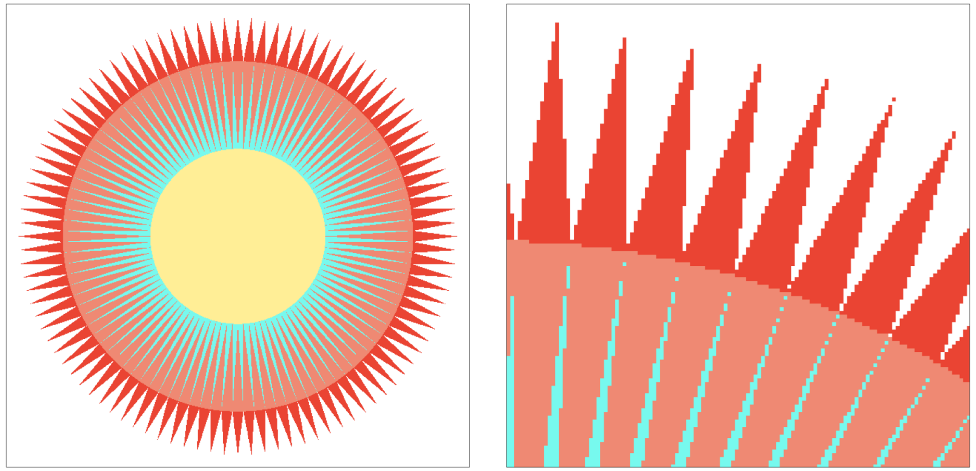

Jaggies (锯齿)

Moiré pattern (摩尔纹)

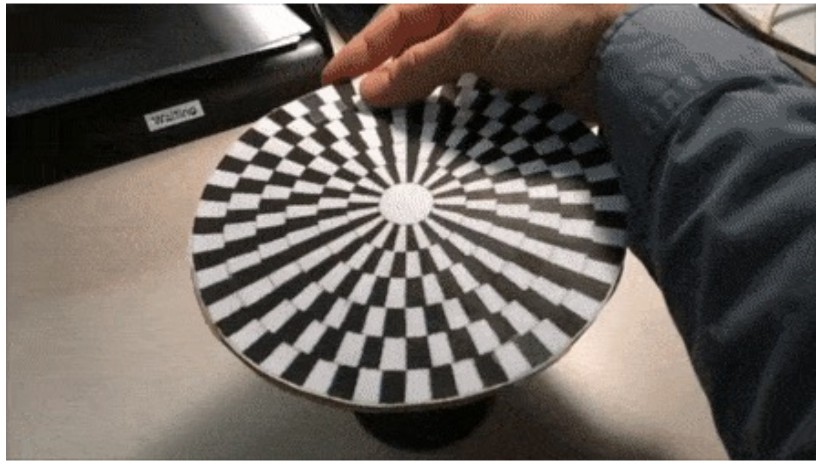

Wagon wheel effect (车轮效应)

Aliasing in NeRF

- The aliasing in NeRF is due to the point sampling in the image plane and 3D space

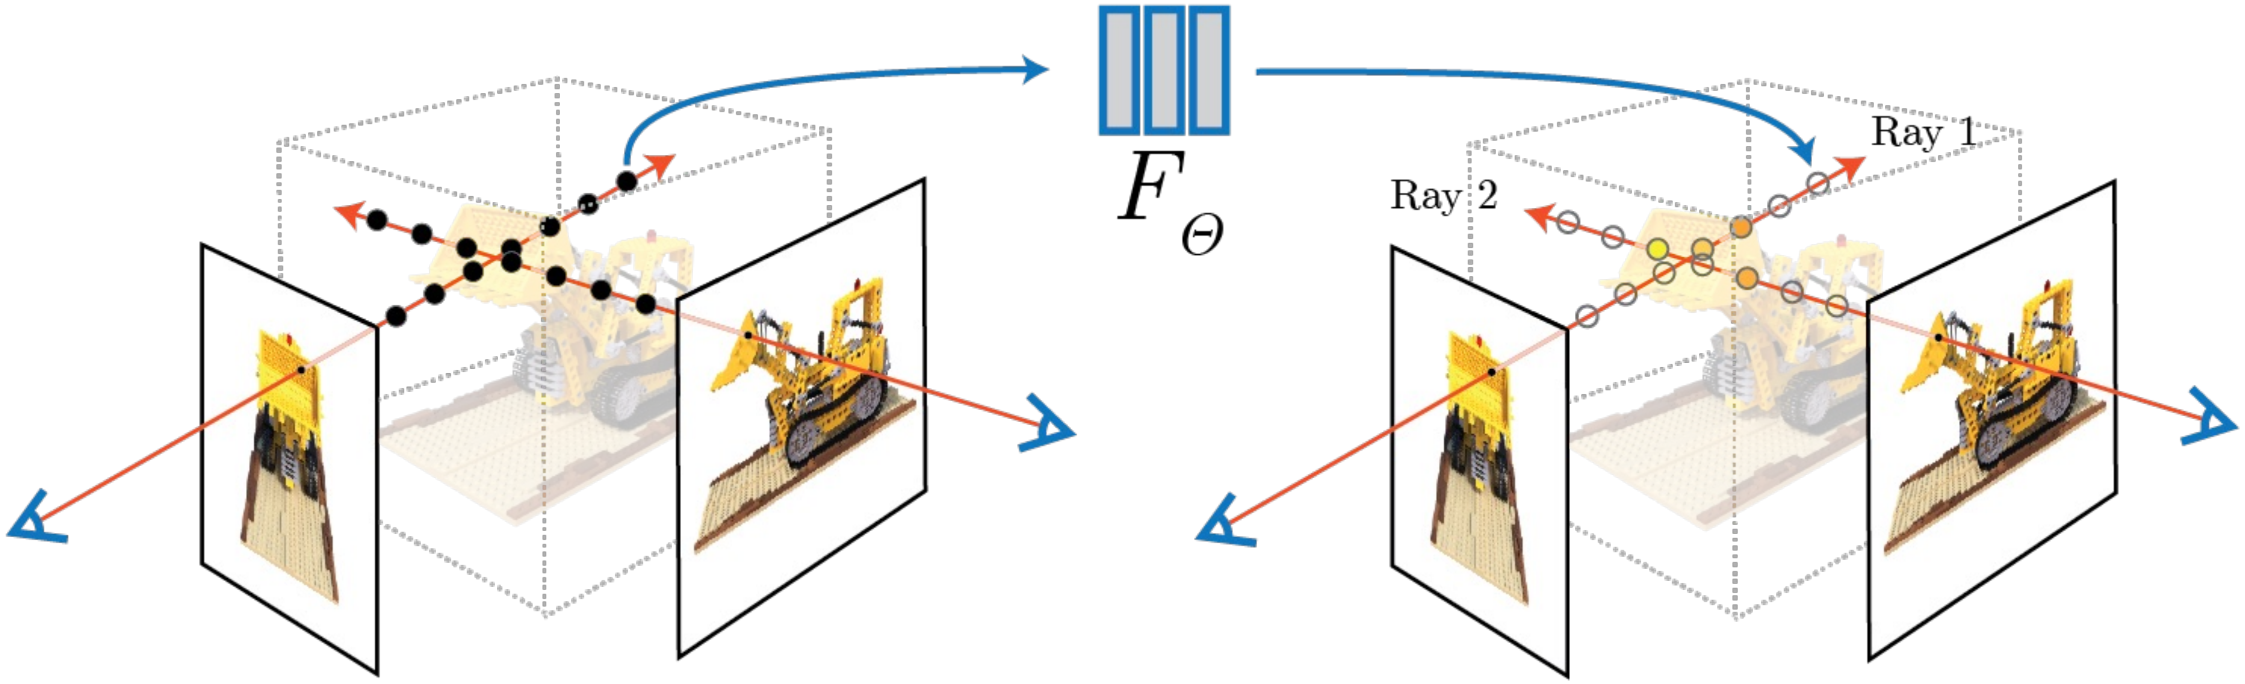

NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis, ECCV'20

Pioneer of Anti-aliased NeRF

- Mip-NeRF: A Multiscale Representation for Anti-Aliasing Neural Radiance Fields, *Barron et al. ICCV'21* (Oral, Best Paper Honorable Mention)

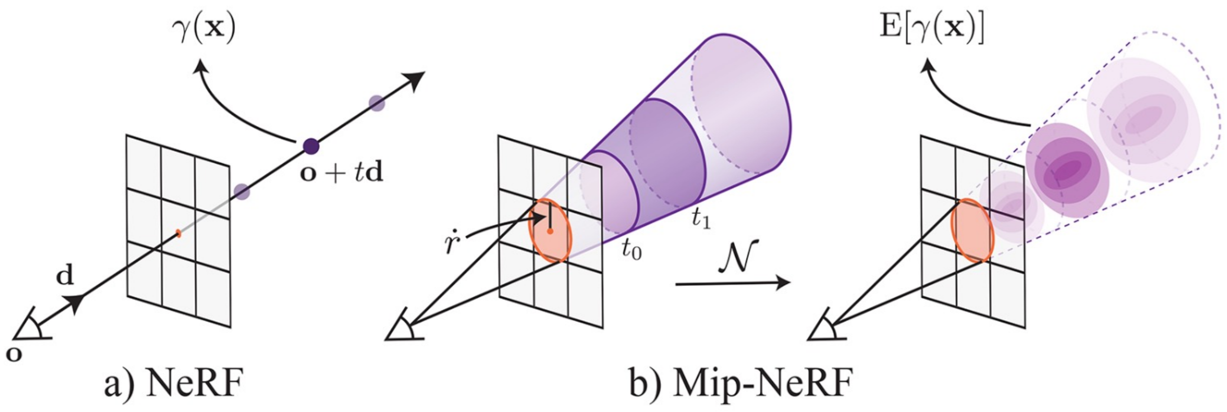

Cone Casting

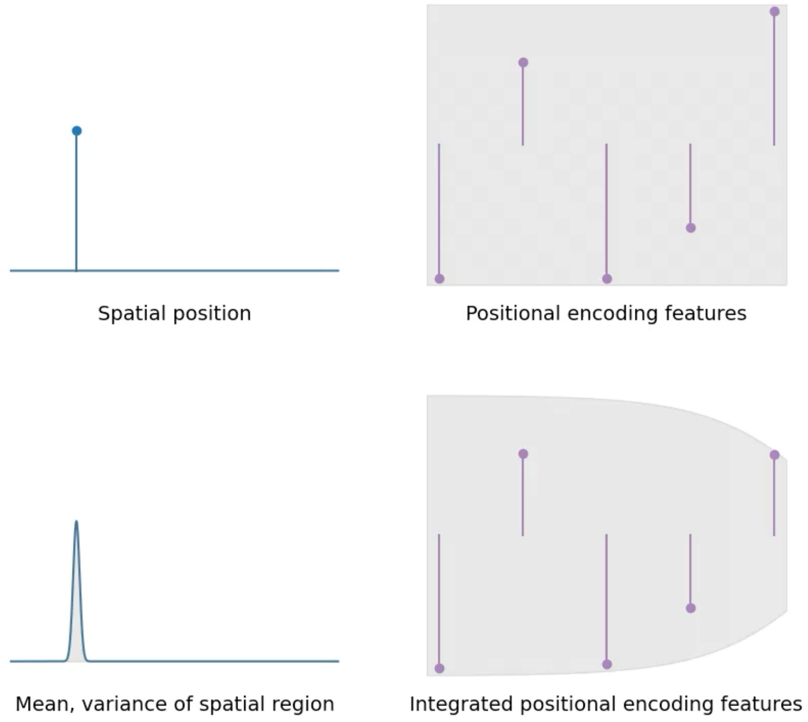

Integrated Positional Encoding

Anti-Aliasing (Quality) *vs.* Speed

Best Paper Finalist
ICCV'23

- **Tri-MipRF**: Tri-Mip Representation for Efficient Anti-Aliasing Neural Radiance Fields

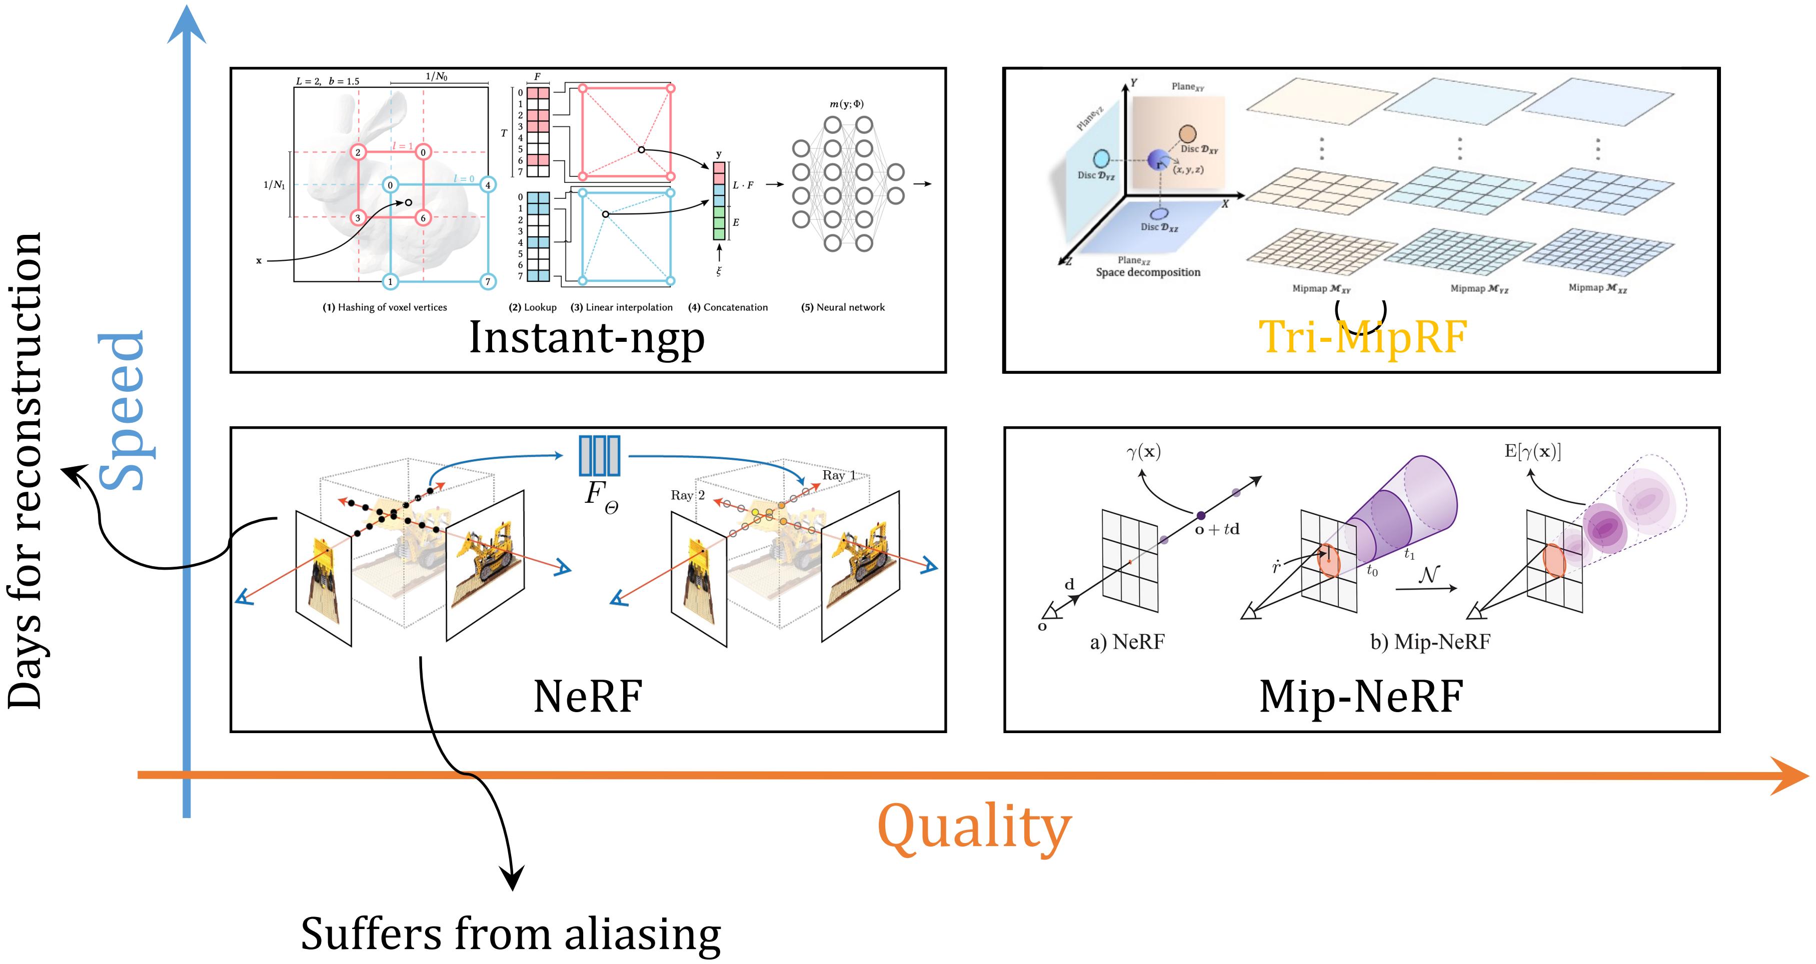

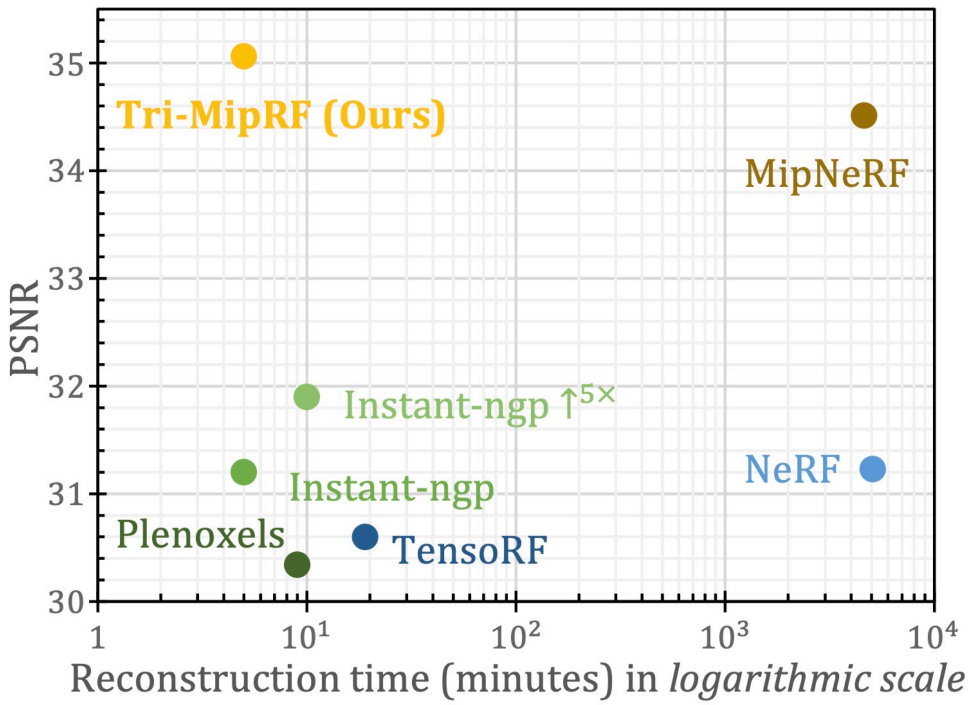

Tri-MipRF

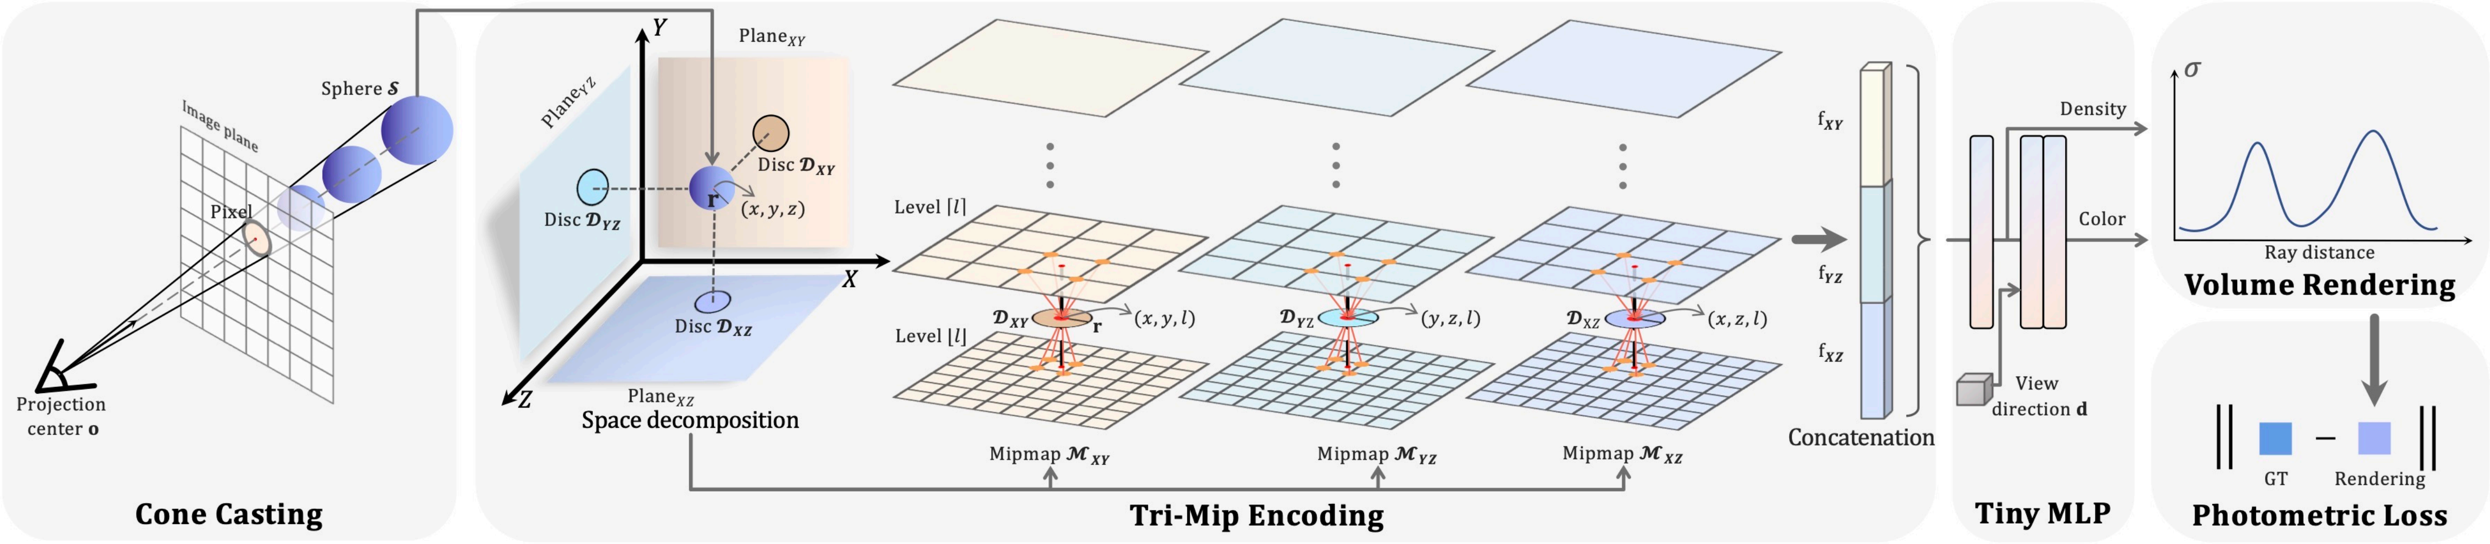

Cone Casting
Projection center o
Image plane
Pixel
Sphere 𝓢
Space decomposition
Plane$_{XY}$
Plane$_{YZ}$
Plane$_{XZ}$
Disc $\mathcal{D}_{XY}$
Disc $\mathcal{D}_{YZ}$
Disc $\mathcal{D}_{XZ}$
(x, y, z)
Level [l]
(x, y, l)
(y, z, l)
(x, z, l)
Mipmap $\mathcal{M}_{XY}$
Mipmap $\mathcal{M}_{YZ}$
Mipmap $\mathcal{M}_{XZ}$
Tri-Mip Encoding
f_{XY}
f_{YZ}
f_{XZ}
Concatenation
View direction d
Tiny MLP
Density
Color
σ
Ray distance
Volume Rendering
GT
Rendering
Photometric Loss

Tri-MipRF

Anti-Aliasing by Multi-sampling

- Zip-NeRF: Anti-Aliased Grid-Based Neural Radiance Fields, *Barron et al. ICCV'23* (Oral, Best Paper Finalist)

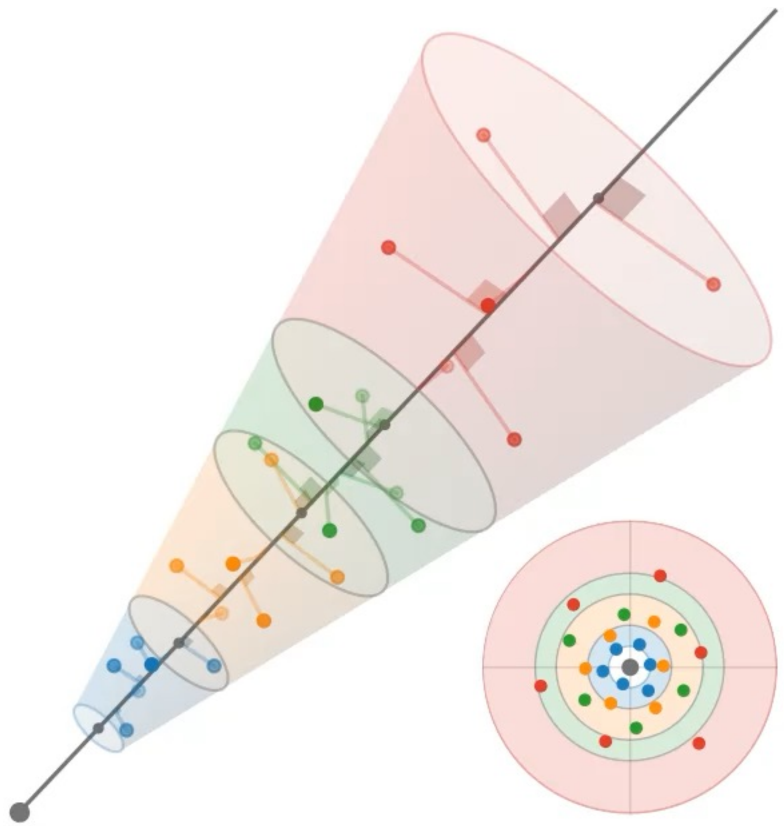

Training

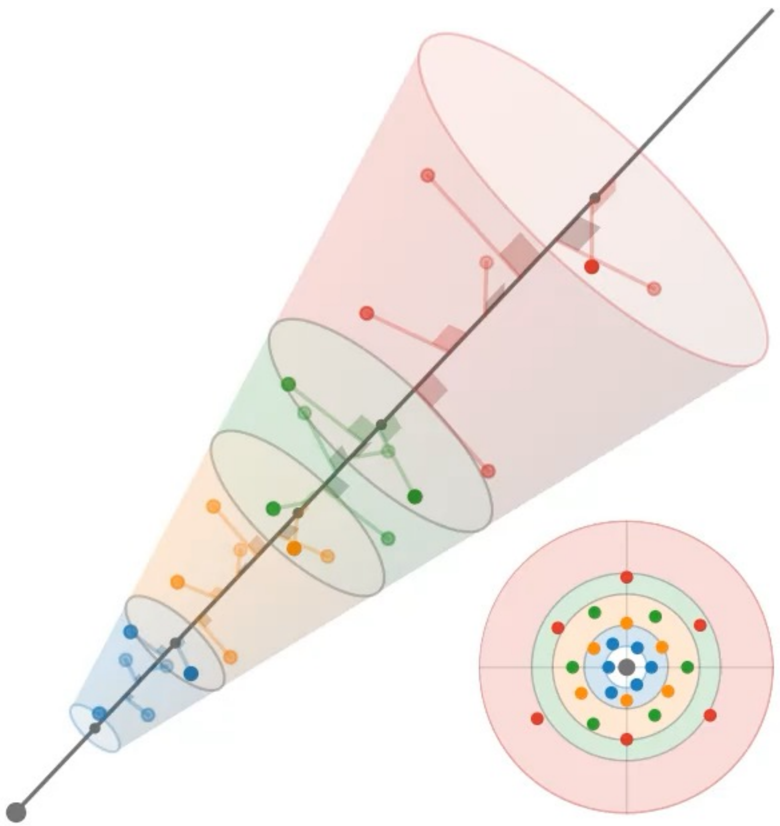

Rendering

Tri-MipRF *vs.* Zip-NeRF

Tri-MipRF (pre-filtering)

- **Pros:**
  - Quality-Efficiency balance
    - high-quality anti-aliasing
    - fast training (~5min) and rendering (real-time support)
  - Compact plane representation
    - friendly with NN
- **Cons:**
  - Isotropic approximation
  - Weak for unbounded scenes

Zip-NeRF (multi-sampling)

- **Pros:**
  - High-quality
    - SOTA quality for unbounded scenes
  - Anisotropic anti-aliasing
    - Unify isotropic and anisotropic anti-aliasing
- **Cons:**
  - Slow training and rendering

Anisotropic Anti-Aliasing

- Isotropic area-sampling leads to ambiguity, we need anisotropic area-sampling

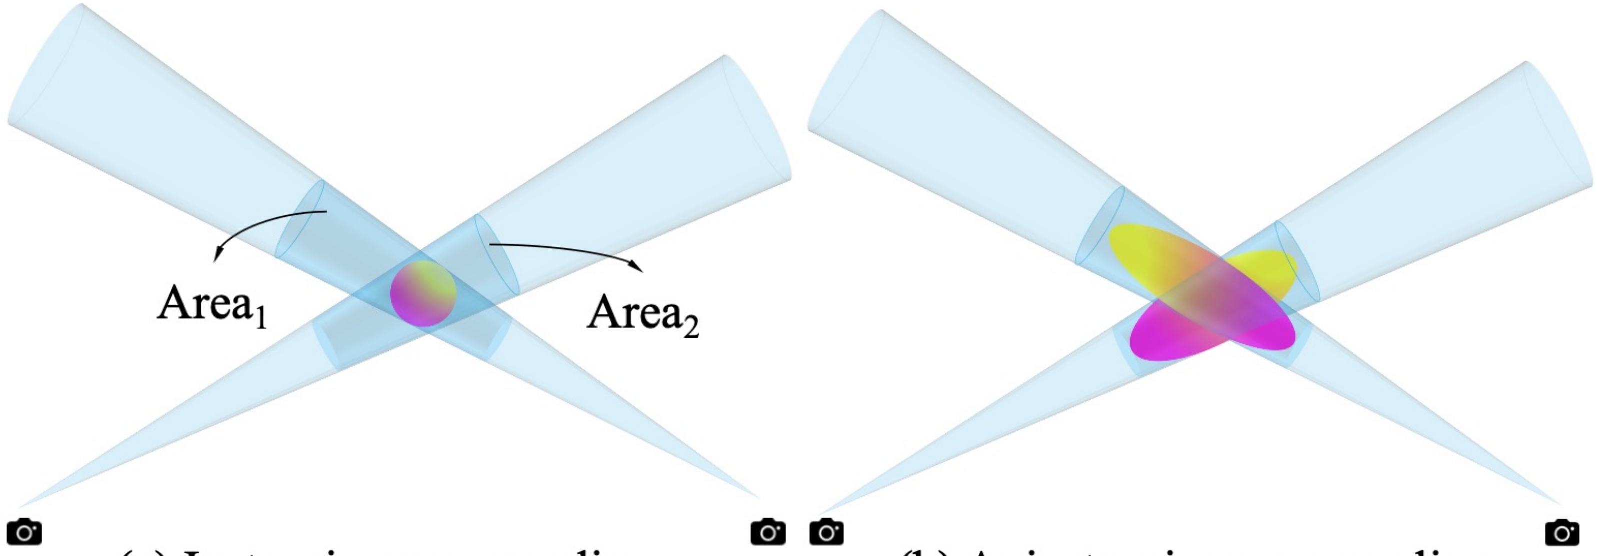

(a) Isotropic area-sampling (b) Anisotropic area-sampling

Anisotropic Anti-Aliasing

- Mipmap encoding

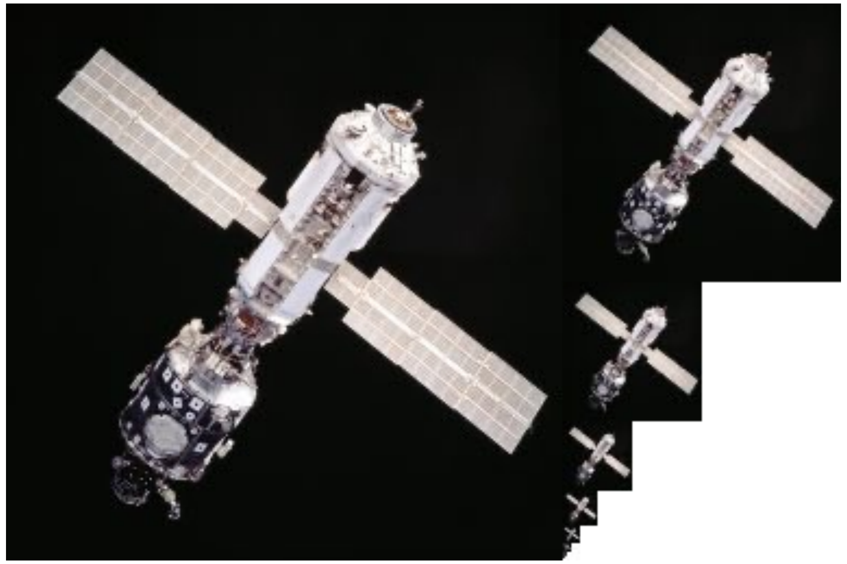

Memory consumption:

$1 + 1/4 + 1/16 + 1/64 + 1/256 + \ldots = ?$ 4/3

- Ripmap encoding

Memory consumption:

4

Anisotropic Anti-Aliasing

- Mipmap *vs.* Ripmap encoding

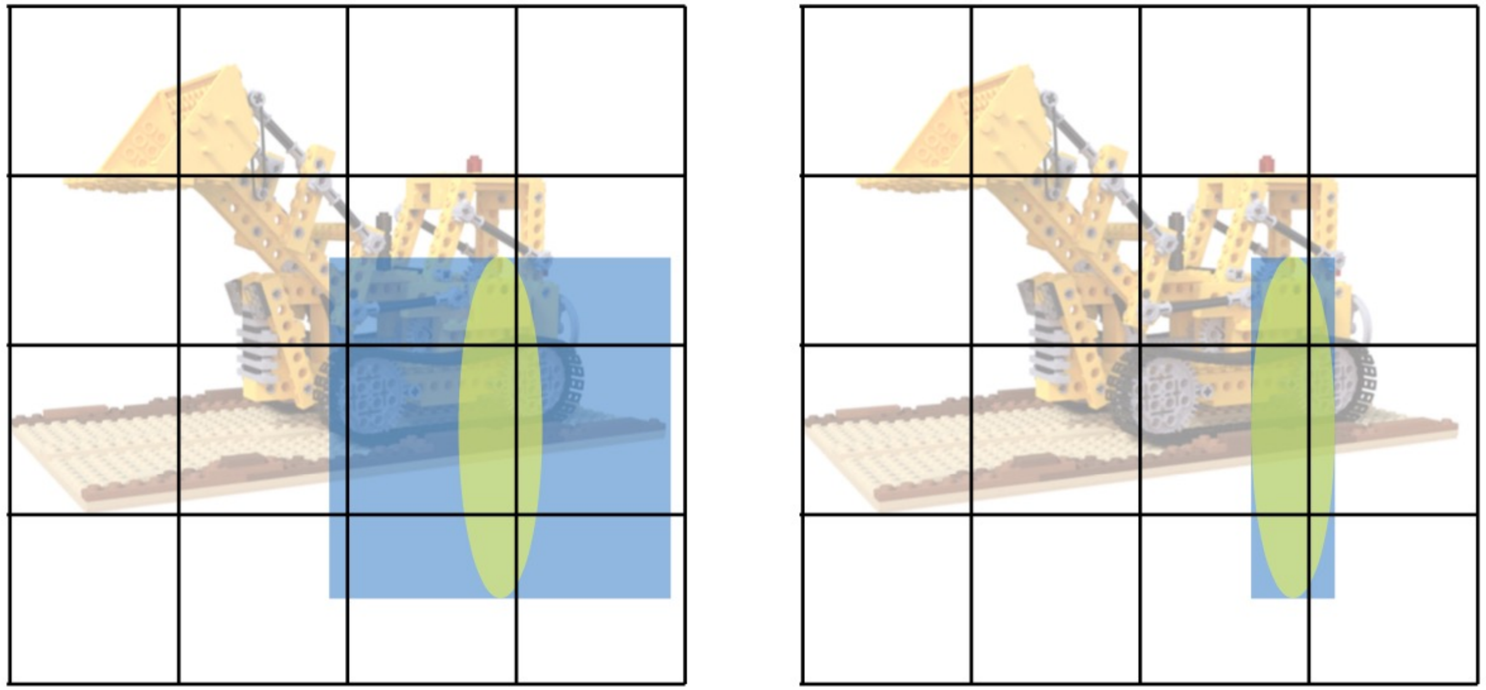

(a) Isotropic area sampling (b) Anisotropic area sampling

Anisotropic Anti-Aliasing

- Tri-plane *vs.* Platonic projection

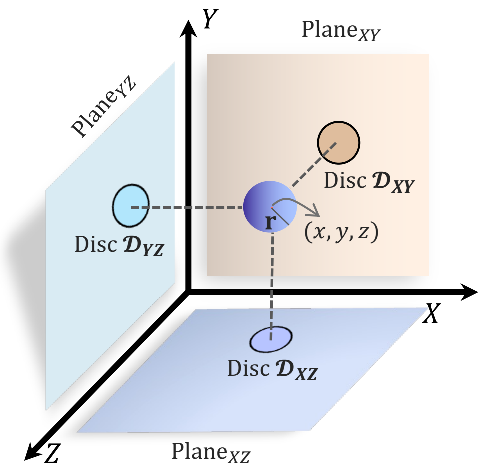

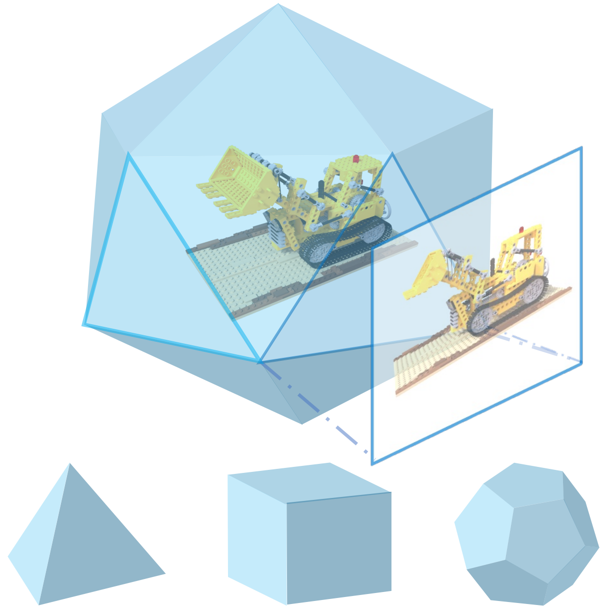

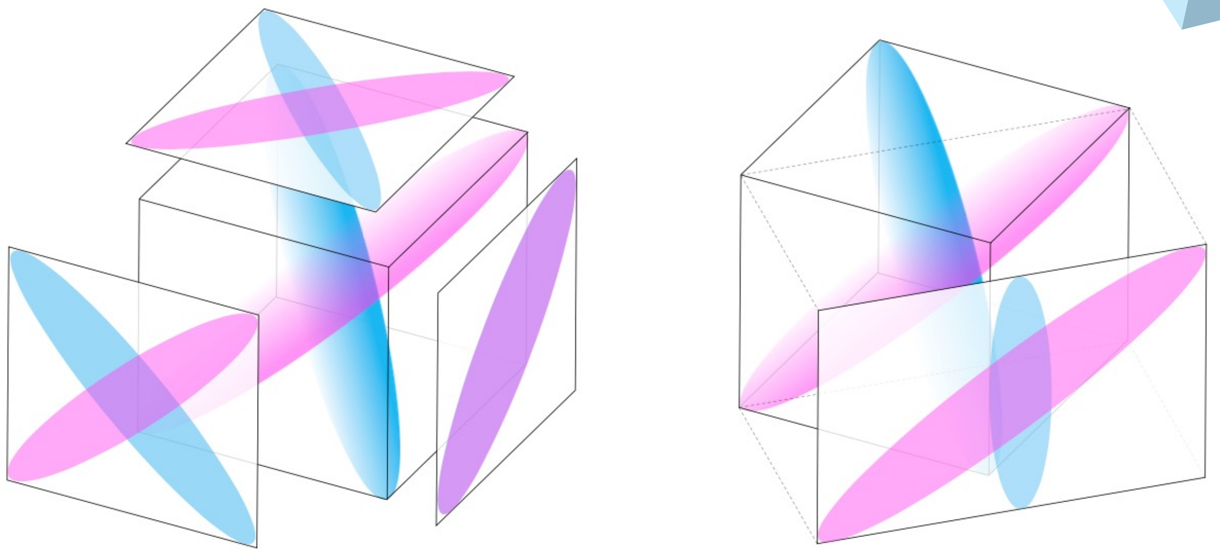

(a) Projection on orthogonal tri-plane (b) Projection on an additional plane

Rip-NeRF

- Rip-NeRF: Anti-aliasing Radiance Fields with Ripmap-Encoded Platonic Solids, SIGGRAPH'24

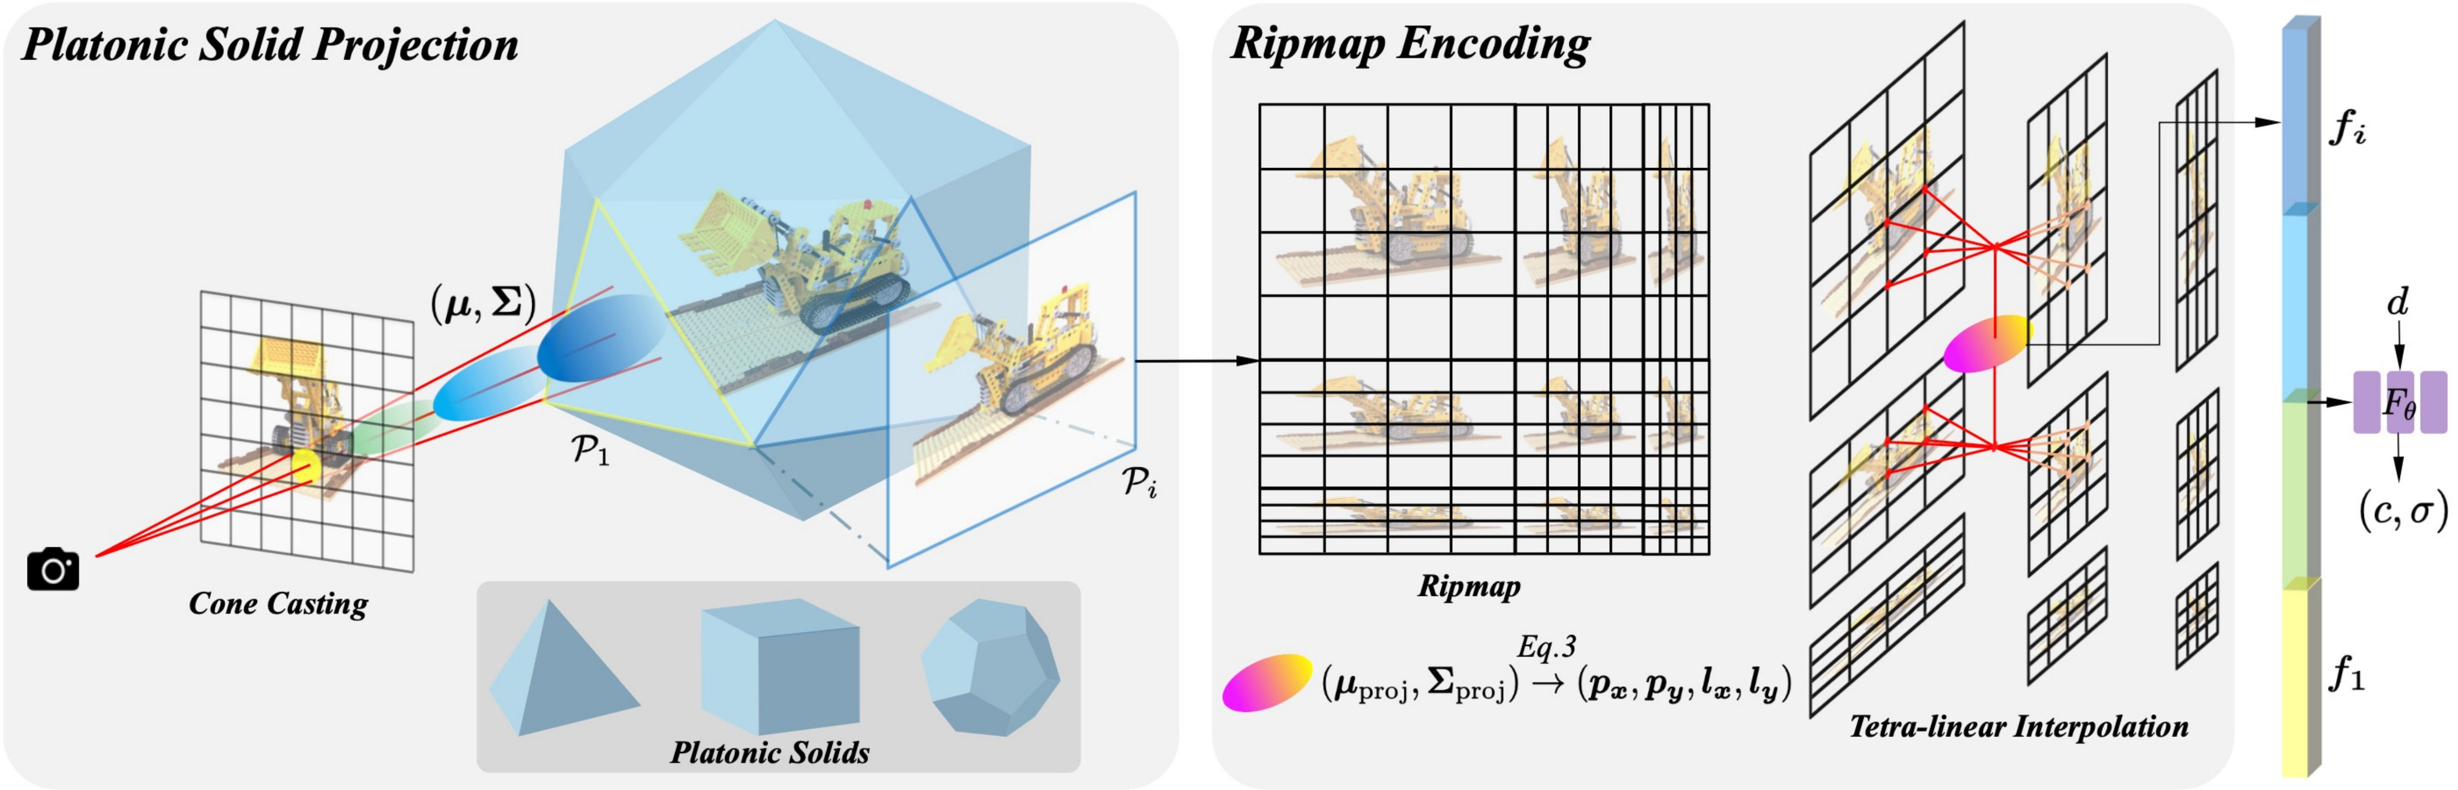

Rip-NeRF

- Rip-NeRF: Anti-aliasing Radiance Fields with Ripmap-Encoded Platonic Solids

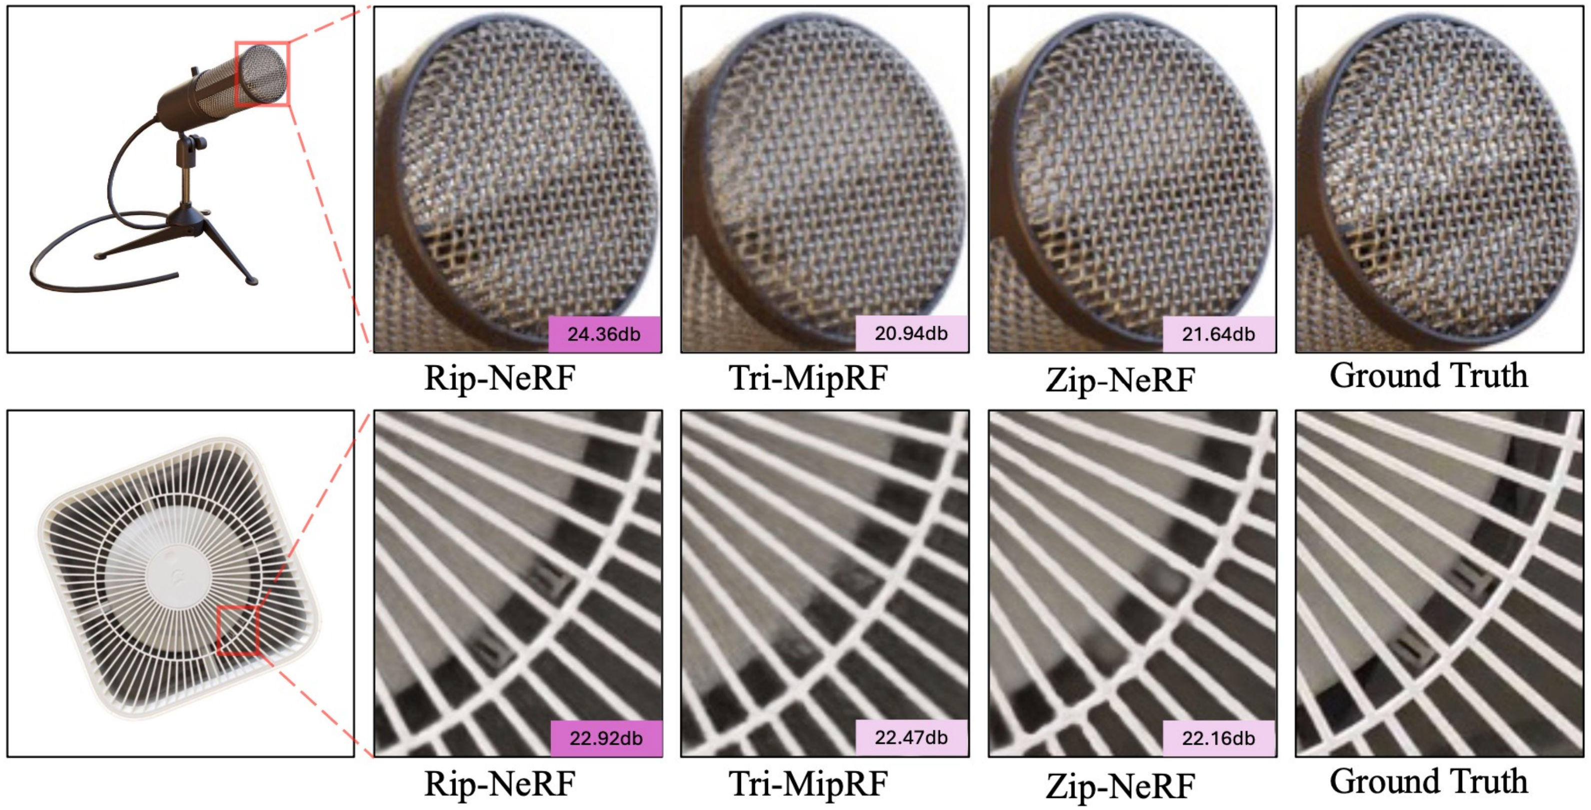

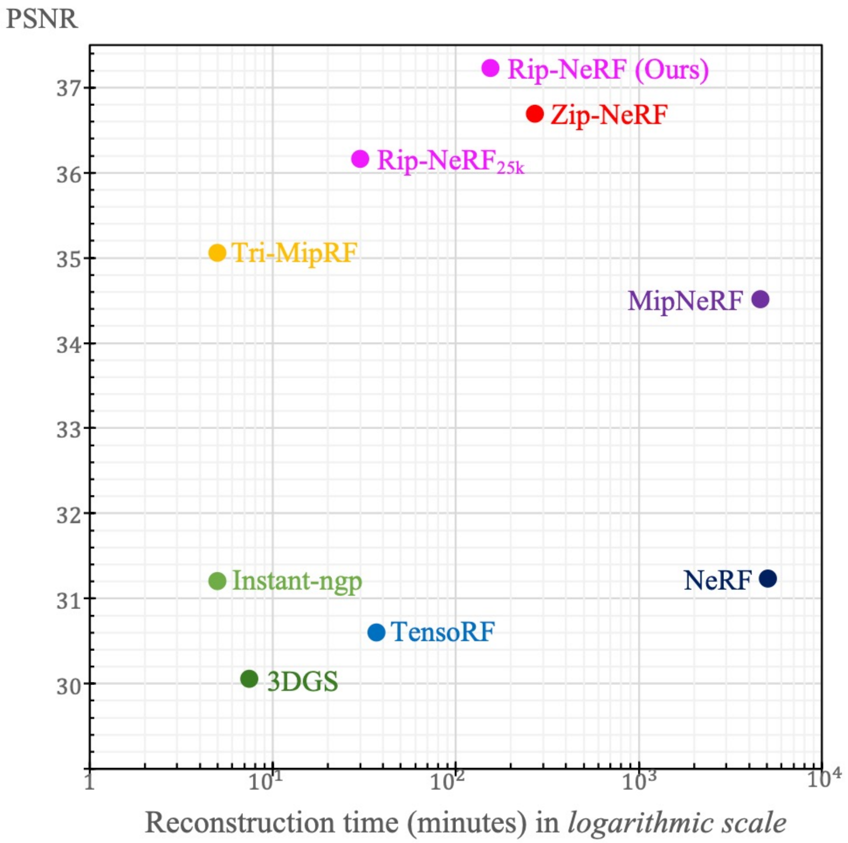

Rip-NeRF

- Quantitative results on ms-Blender

	Train ↓	Size ↓	PSNR ↑					SSIM ↑					LPIPS ↓				
			Full Res.	1/2 Res.	1/4 Res.	1/8 Res.	Avg.	Full Res.	1/2 Res.	1/4 Res.	1/8 Res.	Avg.	Full Res.	1/2 Res.	1/4 Res.	1/8 Res	Avg.
NeRF w/o $\mathcal{L}_{\text{area}}$	3 days	5.00 MB	31.20	30.65	26.25	22.53	27.66	0.950	0.956	0.930	0.871	0.927	0.055	0.034	0.043	0.075	0.052
NeRF [Mildenhall et al. 2020]	3 days	5.00 MB	29.90	32.13	33.40	29.47	31.23	0.938	0.959	0.973	0.962	0.958	0.074	0.040	0.024	0.039	0.044
TensoRF [Chen et al. 2022]	19 mins	71.8 MB	32.11	33.03	30.45	26.80	30.60	0.956	0.966	0.962	0.939	0.956	0.056	0.038	0.047	0.076	0.054
Instant-NGP [Müller et al. 2022]	5 mins	64.1 MB	30.00	32.15	33.31	29.35	31.20	0.939	0.961	0.974	0.963	0.959	0.079	0.043	0.026	0.040	0.047
Mip-NeRF [Barron et al. 2021]	3 days	2.50 MB	32.63	34.34	35.47	35.60	34.51	0.958	0.970	0.979	0.983	0.973	0.047	0.026	0.017	0.012	0.026
Tri-MipRF [Hu et al. 2023]	5.5 mins	48.0 MB	33.57	35.21	35.96	36.46	35.30	0.962	0.975	0.982	0.987	0.976	0.052	0.029	0.019	0.013	0.028
Zip-NeRF [Barron et al. 2023]	4.5 hrs	592 MB	34.21	36.55	37.88	38.13	36.69	0.974	0.985	0.990	0.992	0.985	0.036	0.019	0.014	0.015	0.021
3DGS [Kerbl et al. 2023]	7.5 mins	27.0 MB	29.00	30.94	32.06	28.21	30.05	0.946	0.965	0.976	0.964	0.963	0.064	0.037	0.024	0.030	0.039
Rip-NeRF$_{25k}$	32 mins	160 MB	34.30	35.94	36.92	37.47	36.16	0.966	0.978	0.984	0.989	0.979	0.045	0.025	0.016	0.011	0.024
Rip-NeRF (Ours)	2.6 hrs	160 MB	35.30	37.01	38.07	38.54	37.23	0.973	0.983	0.988	0.991	0.984	0.037	0.019	0.011	0.008	0.019

Zip-NeRF

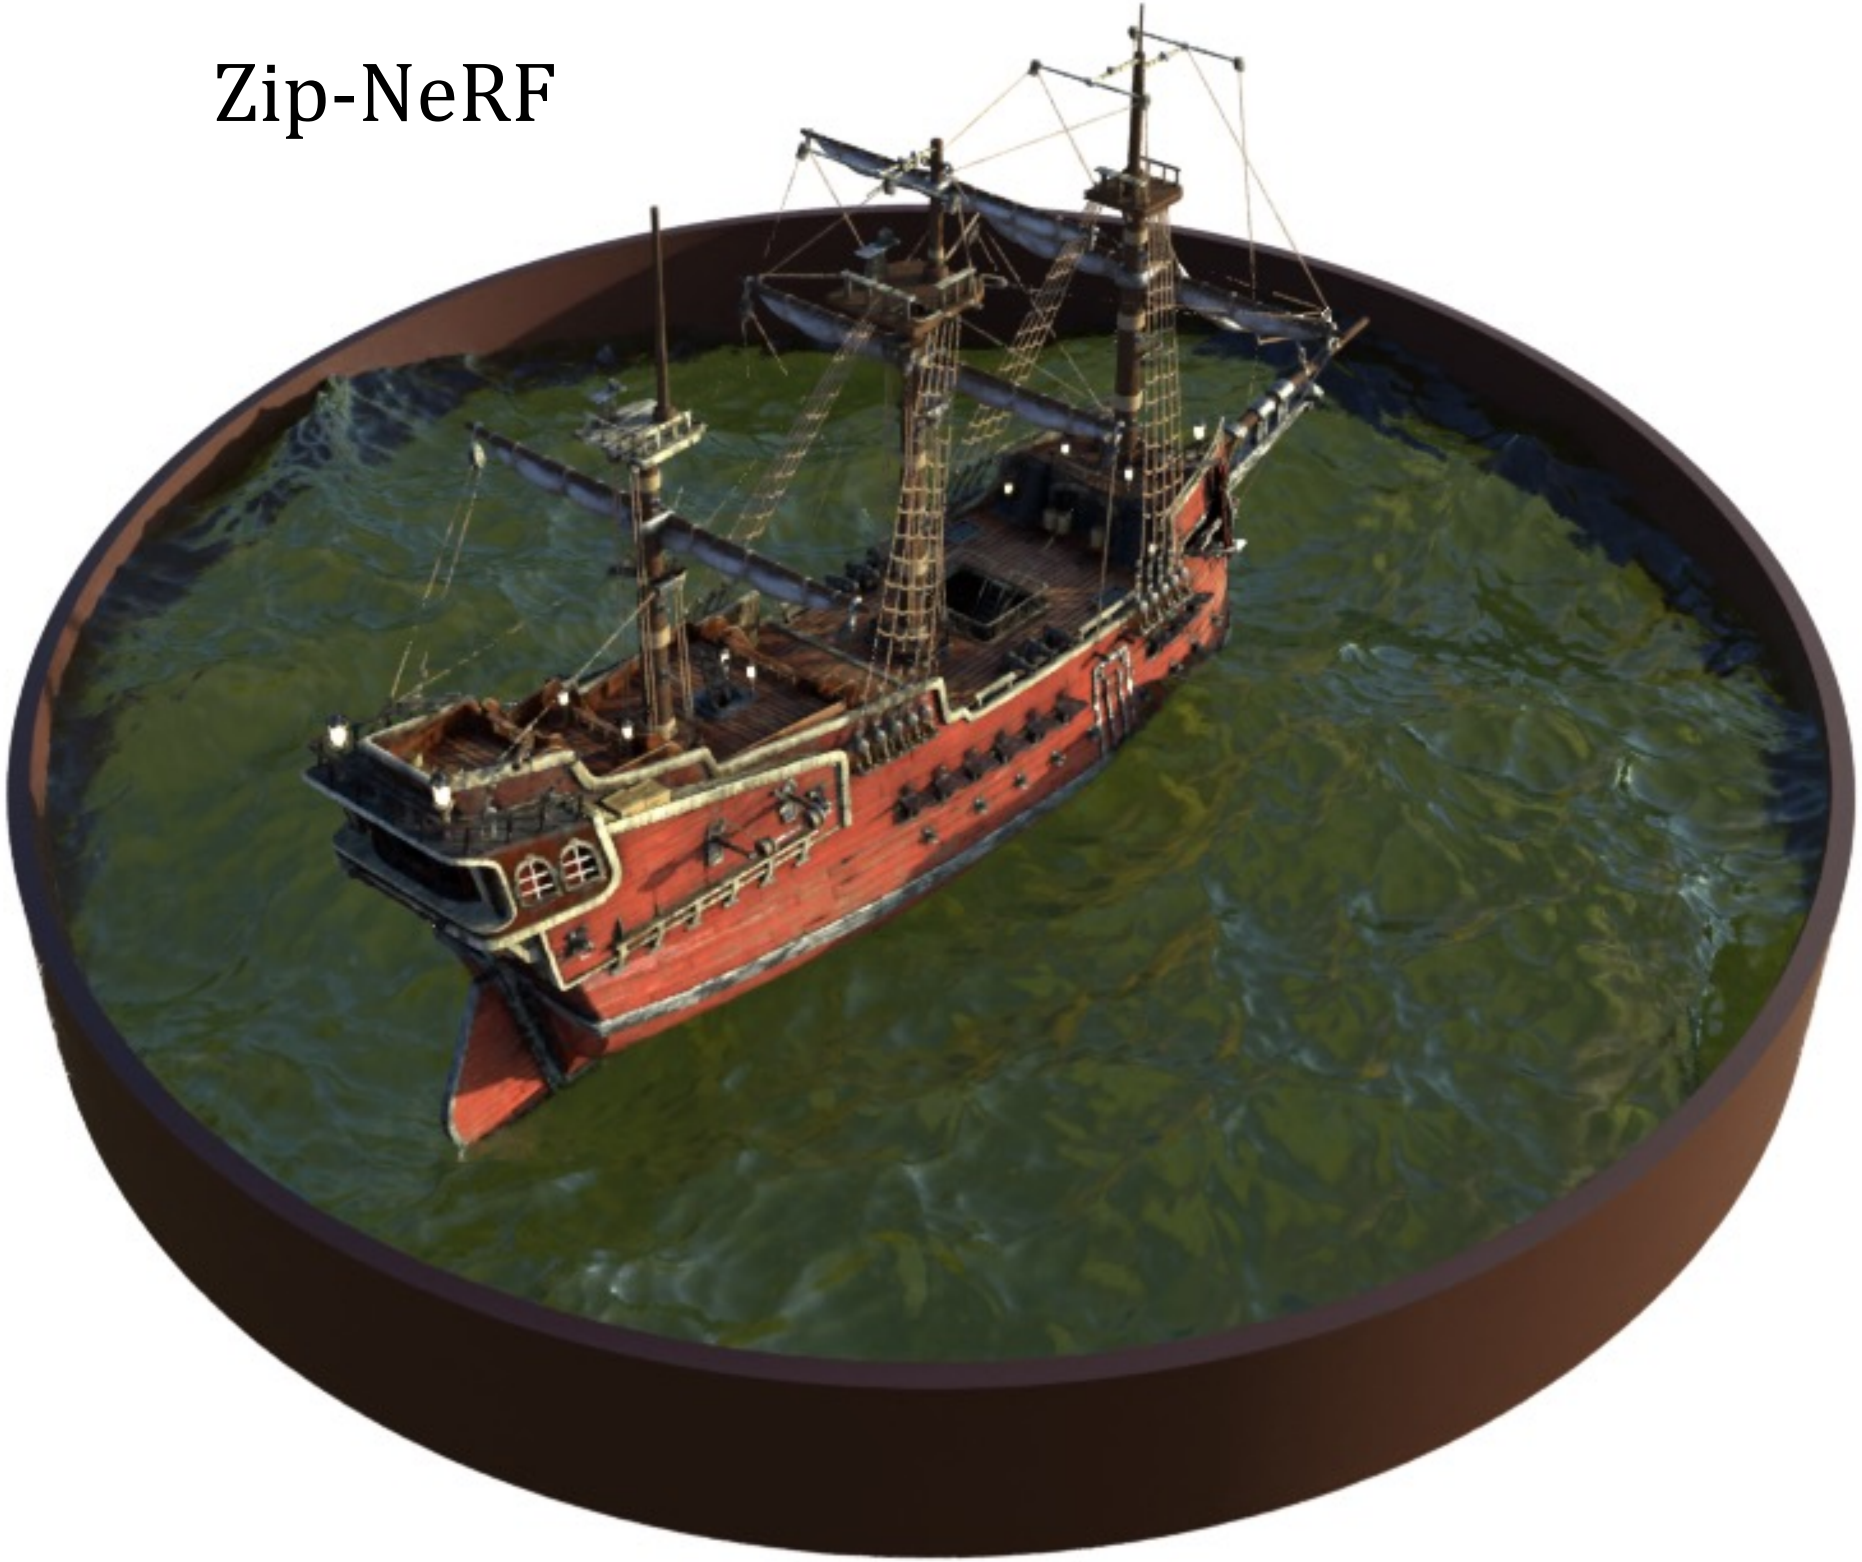

Rip-NeRF

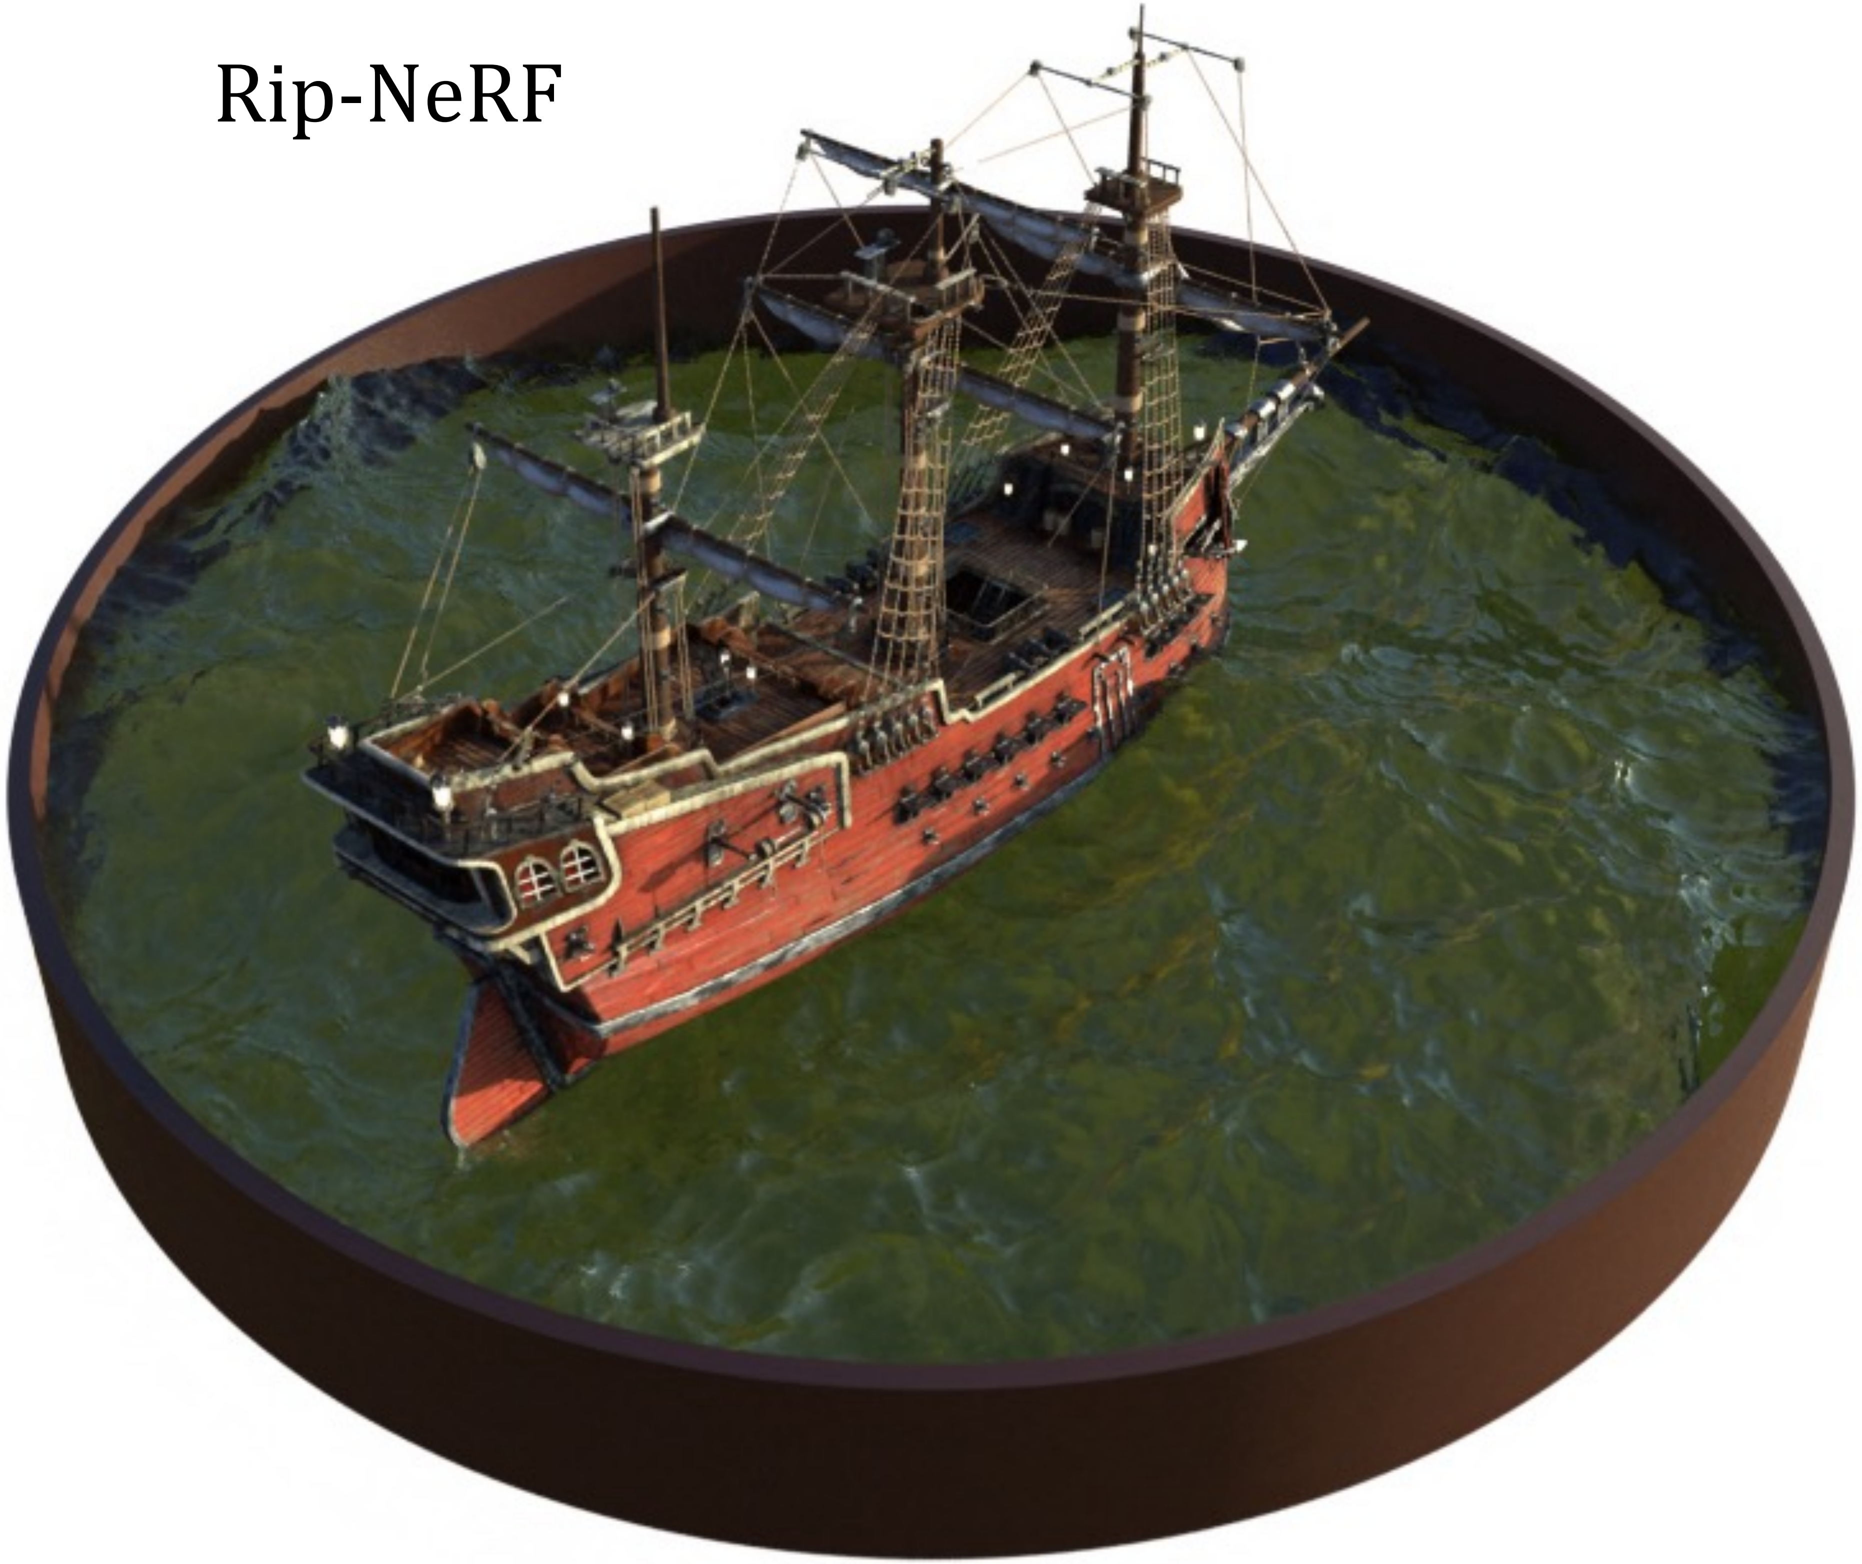

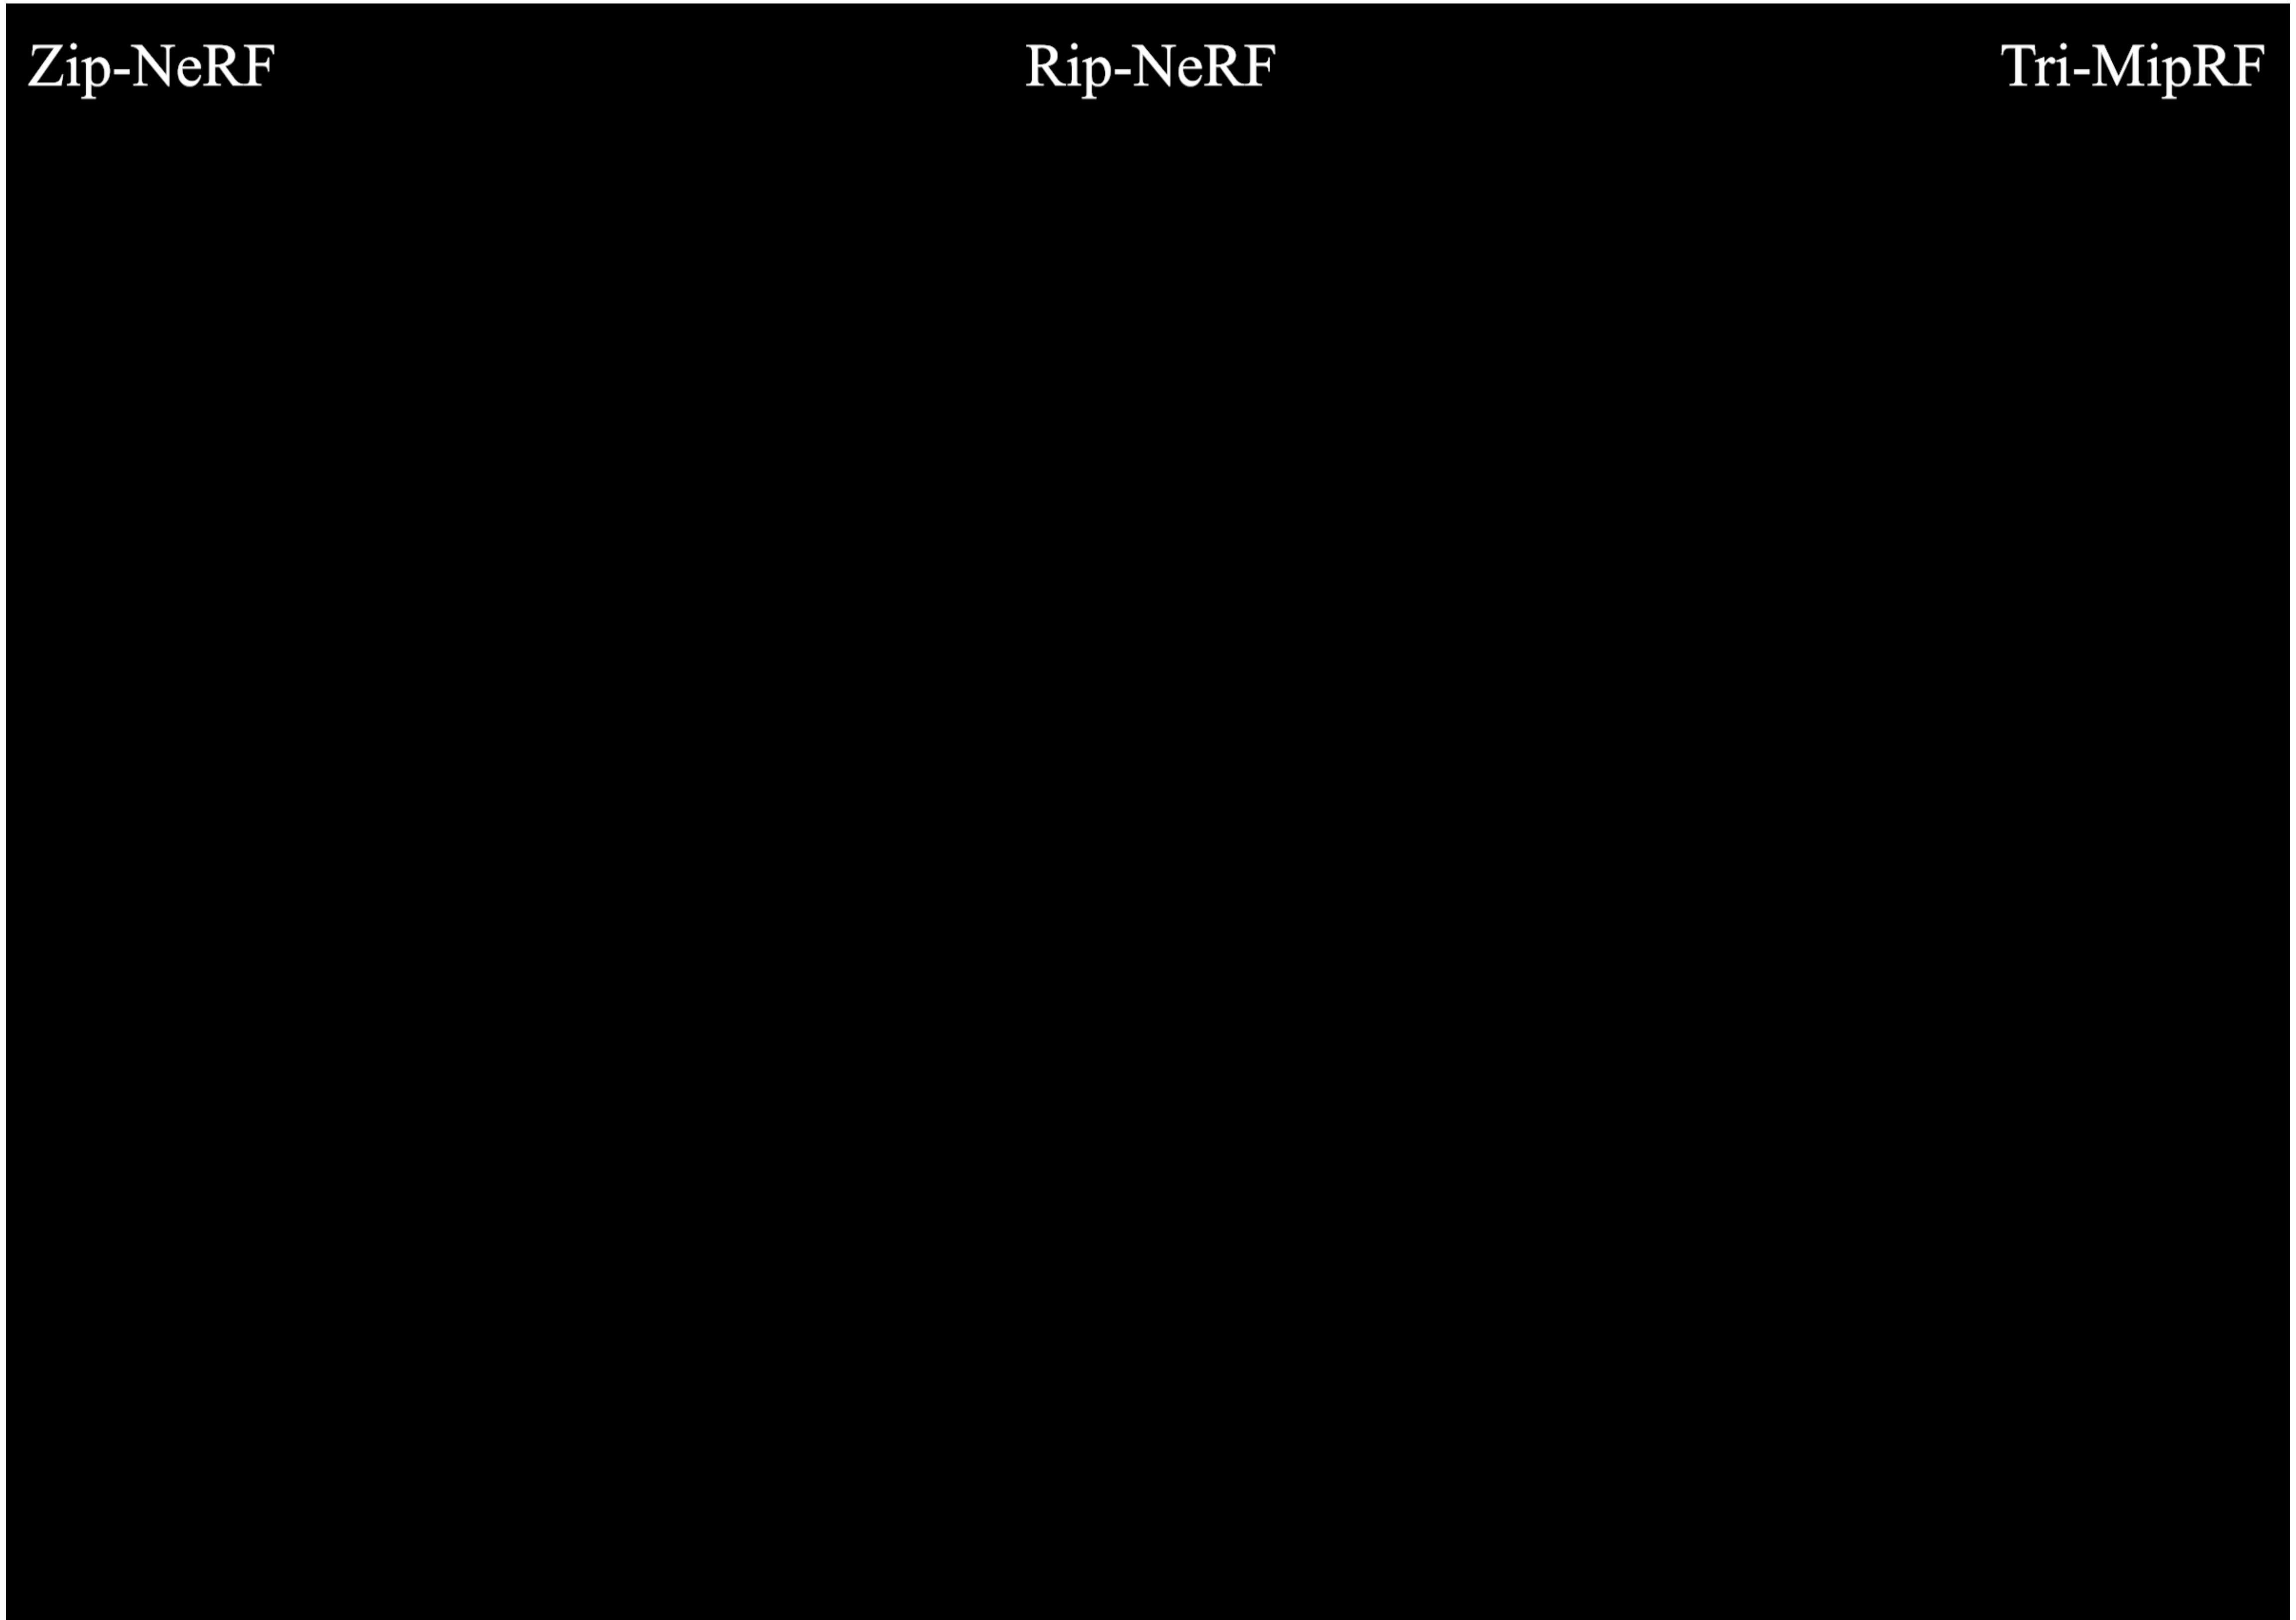
Zip-NeRF
Rip-NeRF
Tri-MipRF

Rip-NeRF

- Anti-Aliasing performance

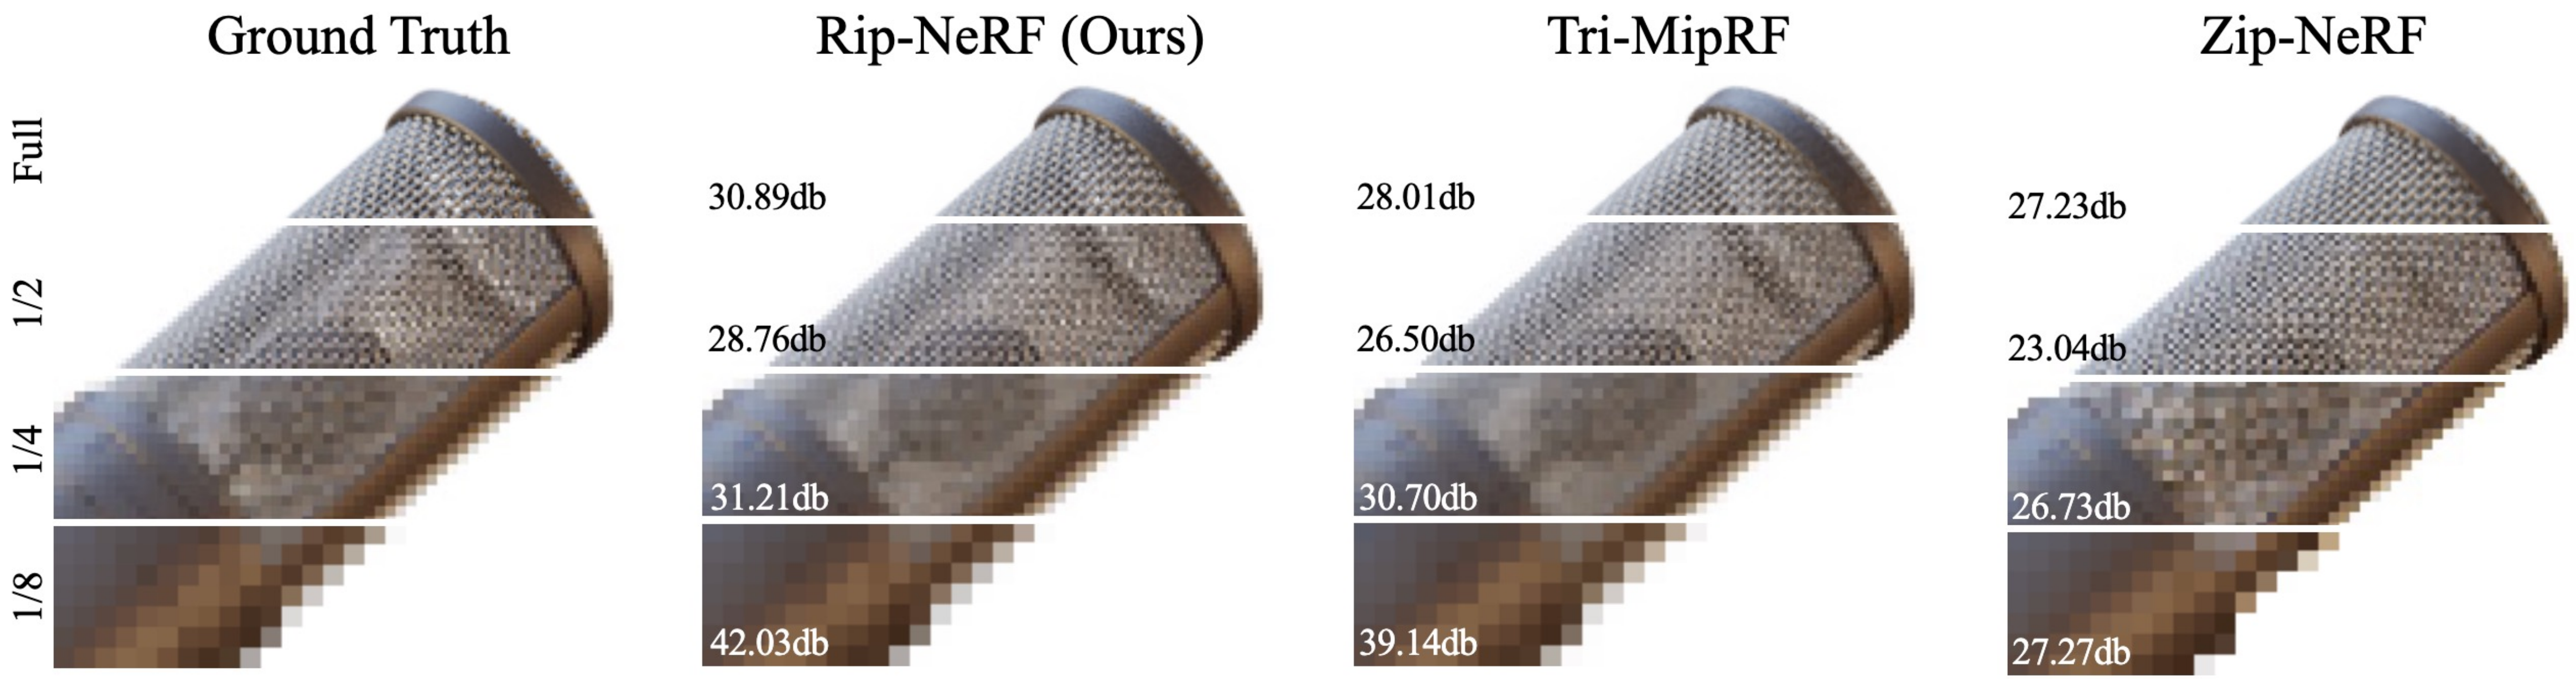

Rip-NeRF

- Ablations

	Train ↓	Size ↓	PSNR↑					SSIM↑					LPIPS↓				
			Full Res.	1/2 Res.	1/4 Res.	1/8 Res.	Avg.	Full Res.	1/2 Res.	1/4 Res.	1/8 Res.	Avg.	Full Res.	1/2 Res.	1/4 Res.	1/8 Res.	Avg.
Rip-NeRF, PS3 (w/o PSP)	25 min	48 MB	33.42	35.00	35.81	36.28	35.14	0.961	0.974	0.981	0.986	0.976	0.053	0.030	0.020	0.013	0.029
Rip-NeRF, PS4	25.5 min	64 MB	32.84	34.28	34.79	34.89	34.20	0.955	0.967	0.972	0.974	0.967	0.063	0.039	0.029	0.025	0.039
Rip-NeRF, PS6	26.5 min	96 MB	33.85	35.49	36.45	36.97	35.69	0.963	0.975	0.983	0.988	0.977	0.050	0.028	0.018	0.012	0.027
Rip-NeRF w/o RE	8.5 min	160 MB	33.64	35.26	36.13	36.68	35.43	0.962	0.974	0.981	0.987	0.976	0.052	0.030	0.020	0.013	0.029
Rip-NeRF$_{25K}$	32 min	160 MB	34.30	35.94	36.92	37.47	36.16	0.966	0.978	0.984	0.989	0.979	0.045	0.025	0.016	0.011	0.024
Rip-NeRF (ours)	2.6 h	160 MB	35.30	37.01	38.07	38.54	37.23	0.973	0.983	0.988	0.991	0.984	0.037	0.019	0.011	0.008	0.019

Rip-NeRF

- Limitations: weak for unbounded scenes

- Reasons:
  - non-vaguely-convex shapes lead to information from self-occluded locations being projected onto the same 2D area
  - space warping encourages more locations along a non-linear curve to be projected onto the same 2D area

- Potential exploration direction:
  - advanced 3D-to-2D mapping (potentially manifold-based) function

Rip-NeRF

- Summary
  - Anisotropic area-sampling is crucial for the quality of anti-aliasing
  - We propose the Platonic projection and Ripmap encoding to address the anisotropy in 3D and 2D space, respectively
  - Rip-NeRF achieves SOTA quality on bounded scenes with a relative faster training and more compact memory consumption compared against Zip-NeRF

Anti-aliasing in 3DGS

- Does 3DGS suffer from aliasing?   YES
- Mip-Splatting: Alias-free 3D Gaussian Splatting, *Yu et al. 2024, CVPR'24* (Oral)

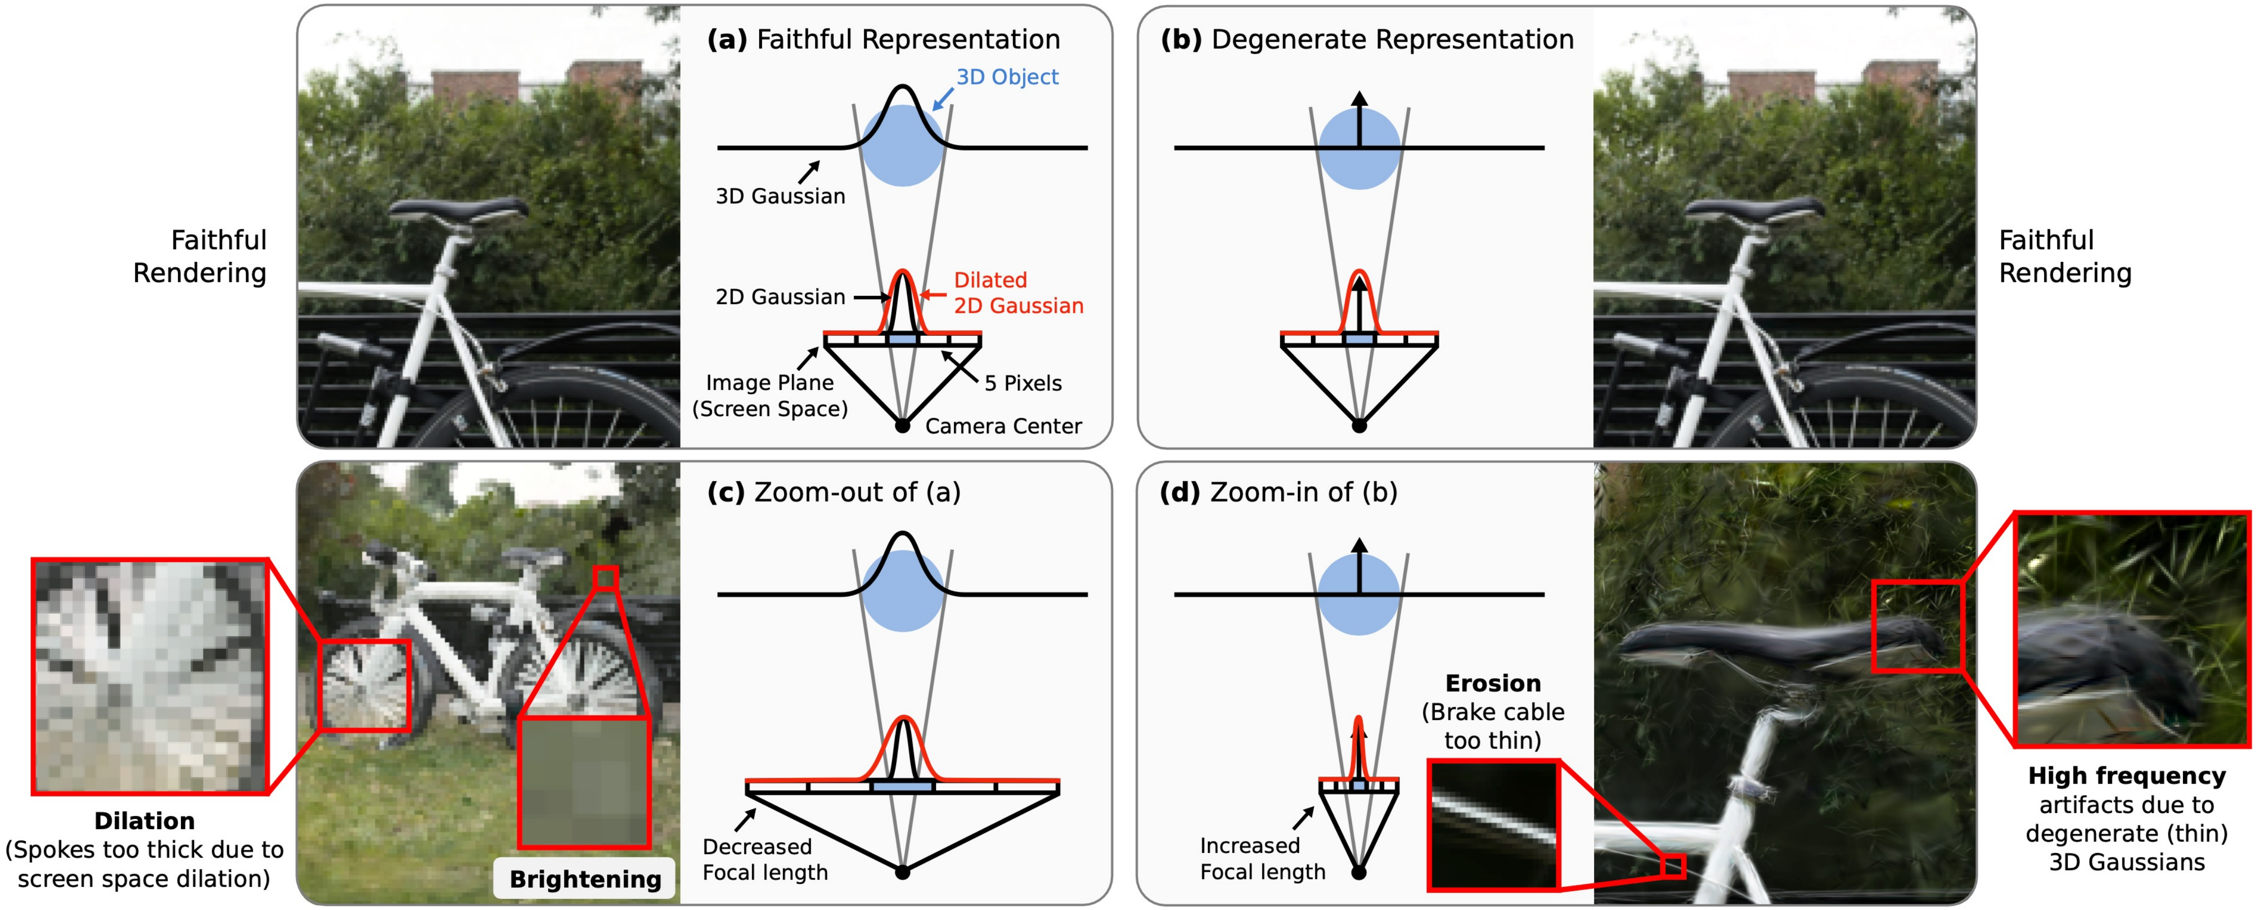

Anti-aliasing in 3DGS

- Our solution: Anti-aliasing via **Analytic Integration**

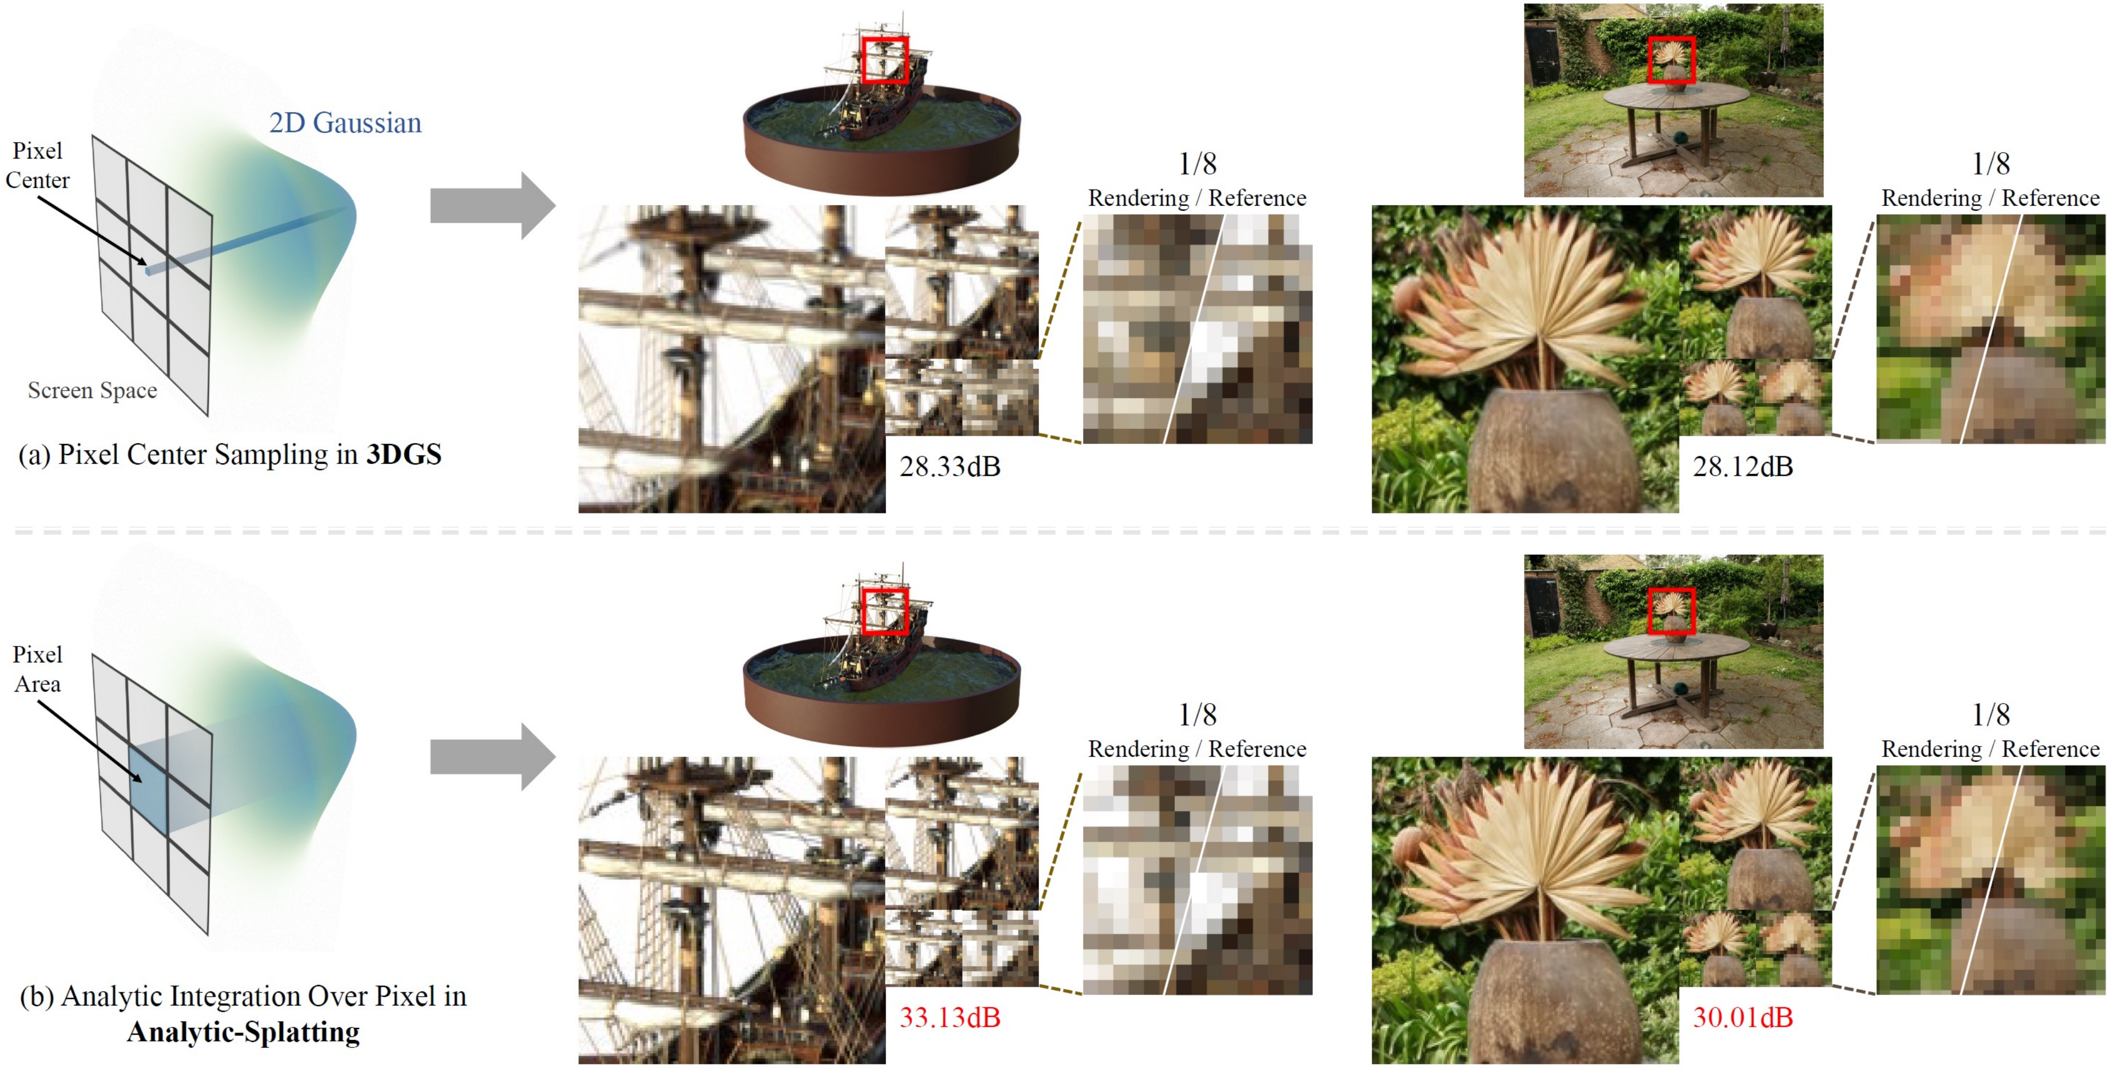

Analytic-Splatting

- Analytic-Splatting: Anti-Aliased 3D Gaussian Splatting via Analytic Integration
  - Core idea: integral over the pixel footprint
  - Challenges:
    - Non-axis-aligned 2D Gaussian integration

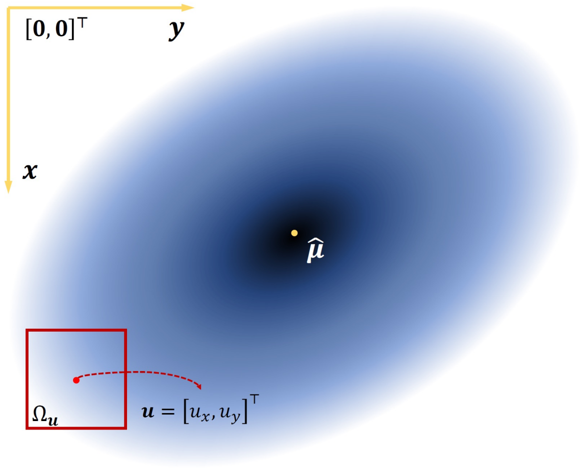

(a) Integral over the original pixel area.

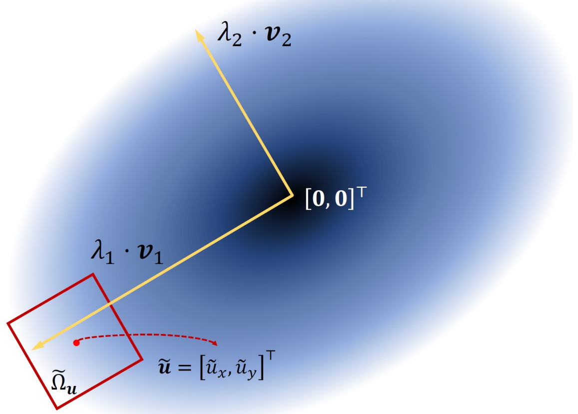

(b) Integral over the rotated pixel area.

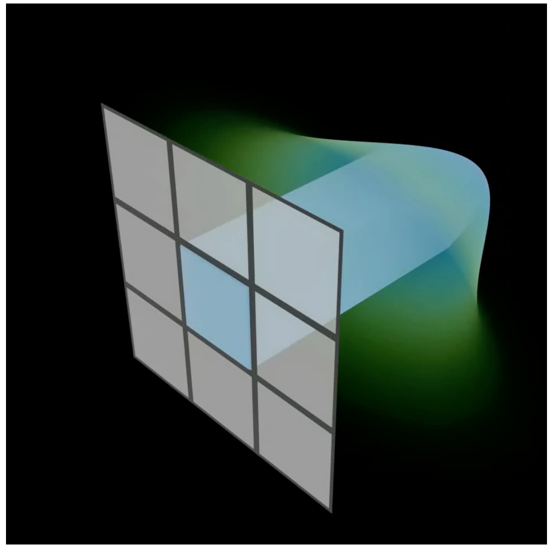

Analytic-Splatting

- Analytic-Splatting: Anti-Aliased 3D Gaussian Splatting via Analytic Integration
  - Core idea: integral over the pixel footprint
  - Challenges:
    - Non-axis-aligned 2D Gaussian integration
    - Efficient analytic integration

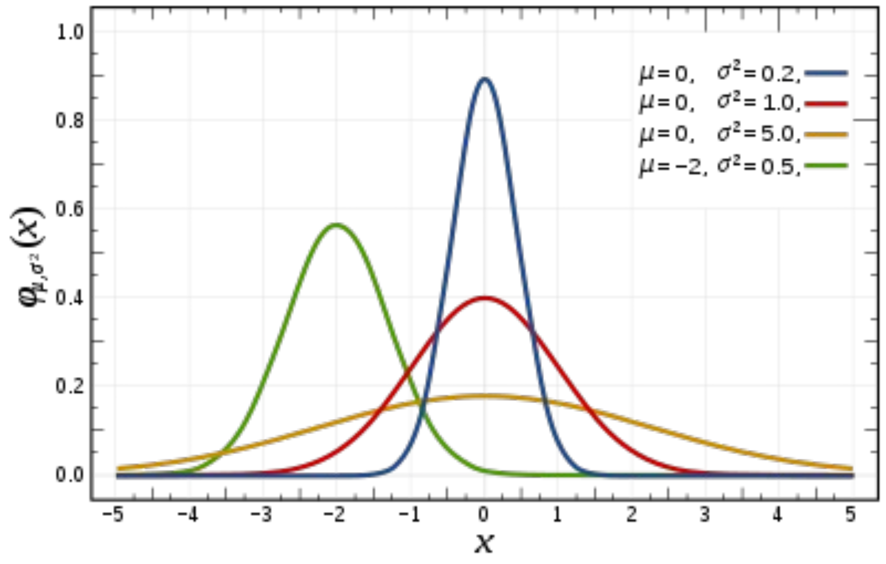

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}$$

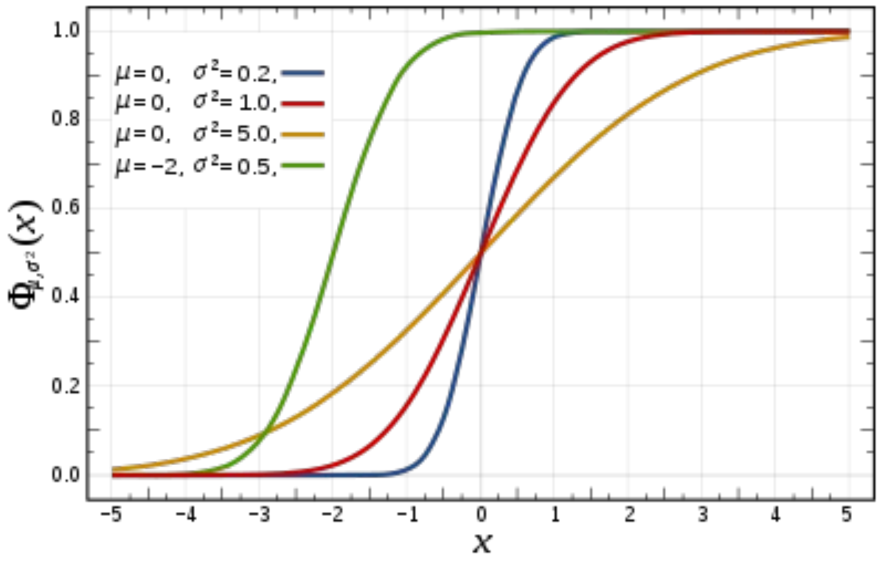

$$\Phi\left(\frac{x-\mu}{\sigma}\right) = \frac{1}{2}\left[1 + \operatorname{erf}\left(\frac{x-\mu}{\sigma\sqrt{2}}\right)\right]$$

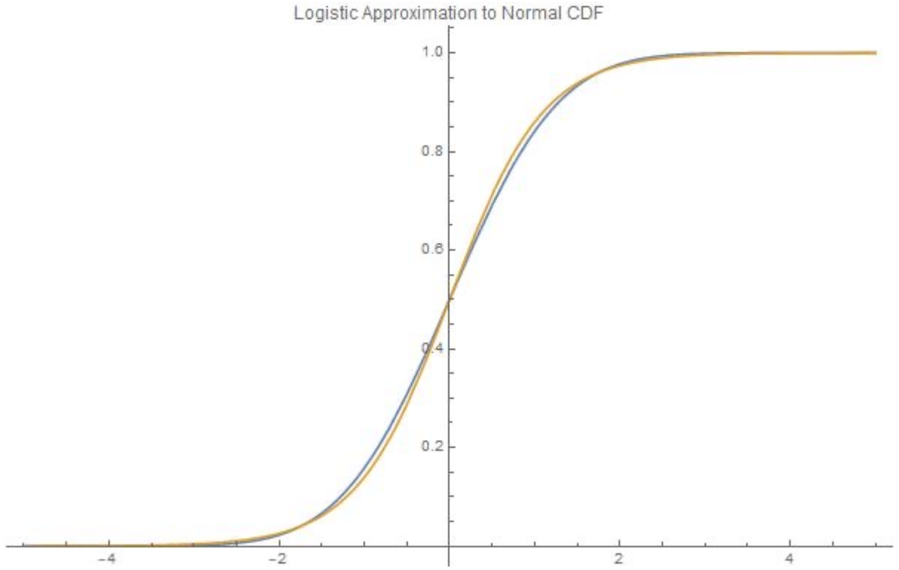

$$S(x) = \frac{1}{1+\exp\left(-1.6 \cdot x - 0.07 \cdot x^3\right)},$$

Analytic-Splatting

- Comparisons among different methods

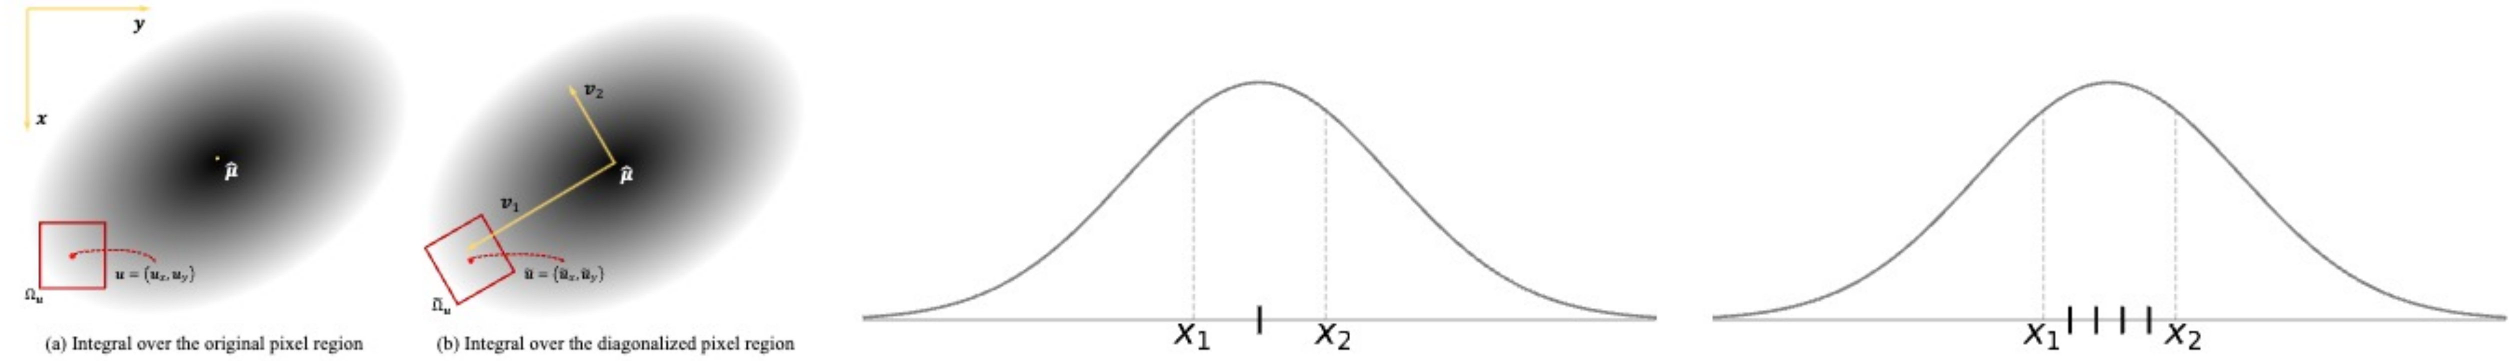

(a) Integration (Ground-truth) **(b)** Baseline (3DGS) **(c)** Super Sampling (3DGS-SS)

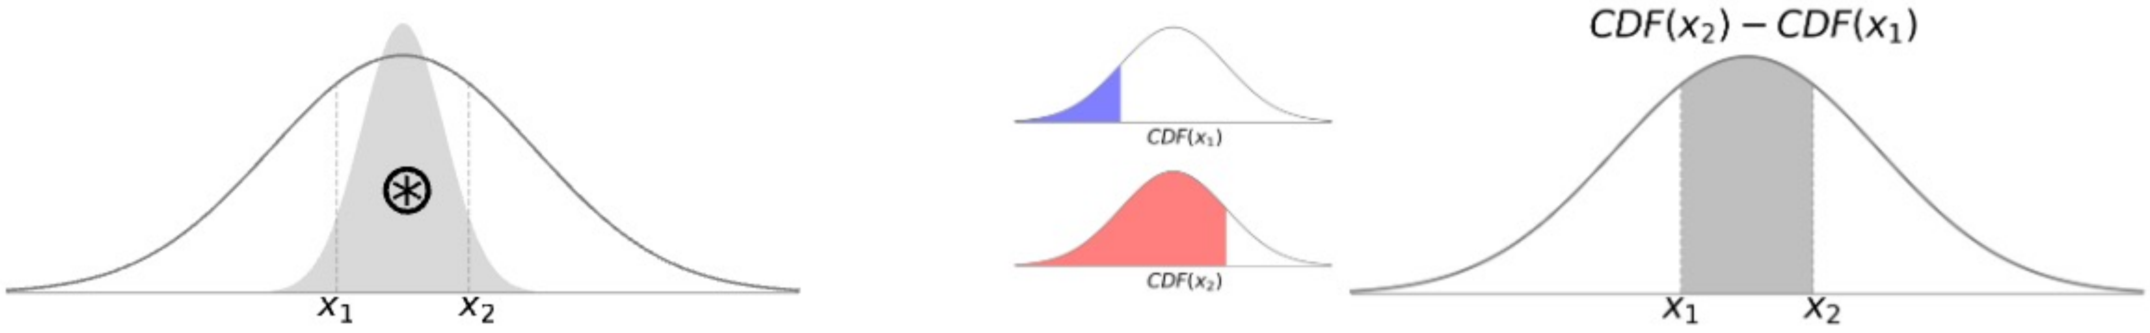

(d) Filtering (Mip-Splatting) **(e)** Analytic Approximation (Ours)

Analytic-Splatting

- Error amylysis

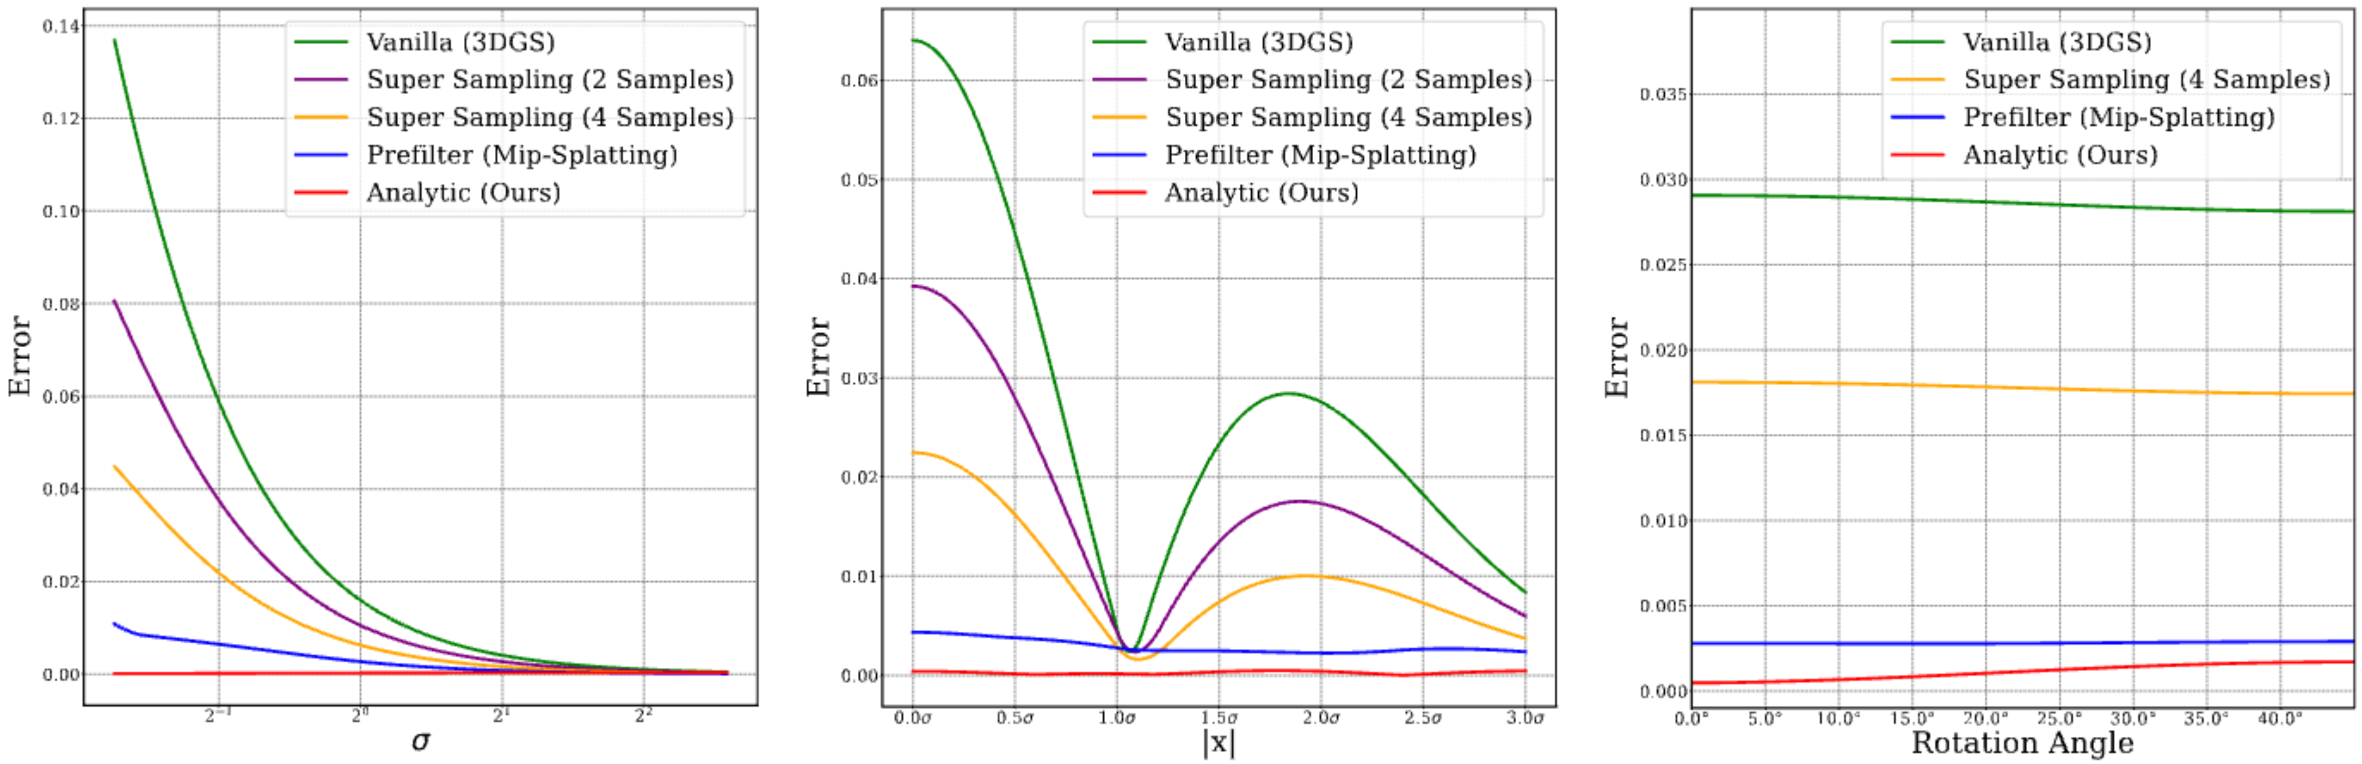

(a) Approximation errors referring to different standard variations.

(b) Approximation errors for different variable distribution.

(c) Approximation errors of rotating the integral domain by different angles in Fig. 3.

Analytic-Splatting: results

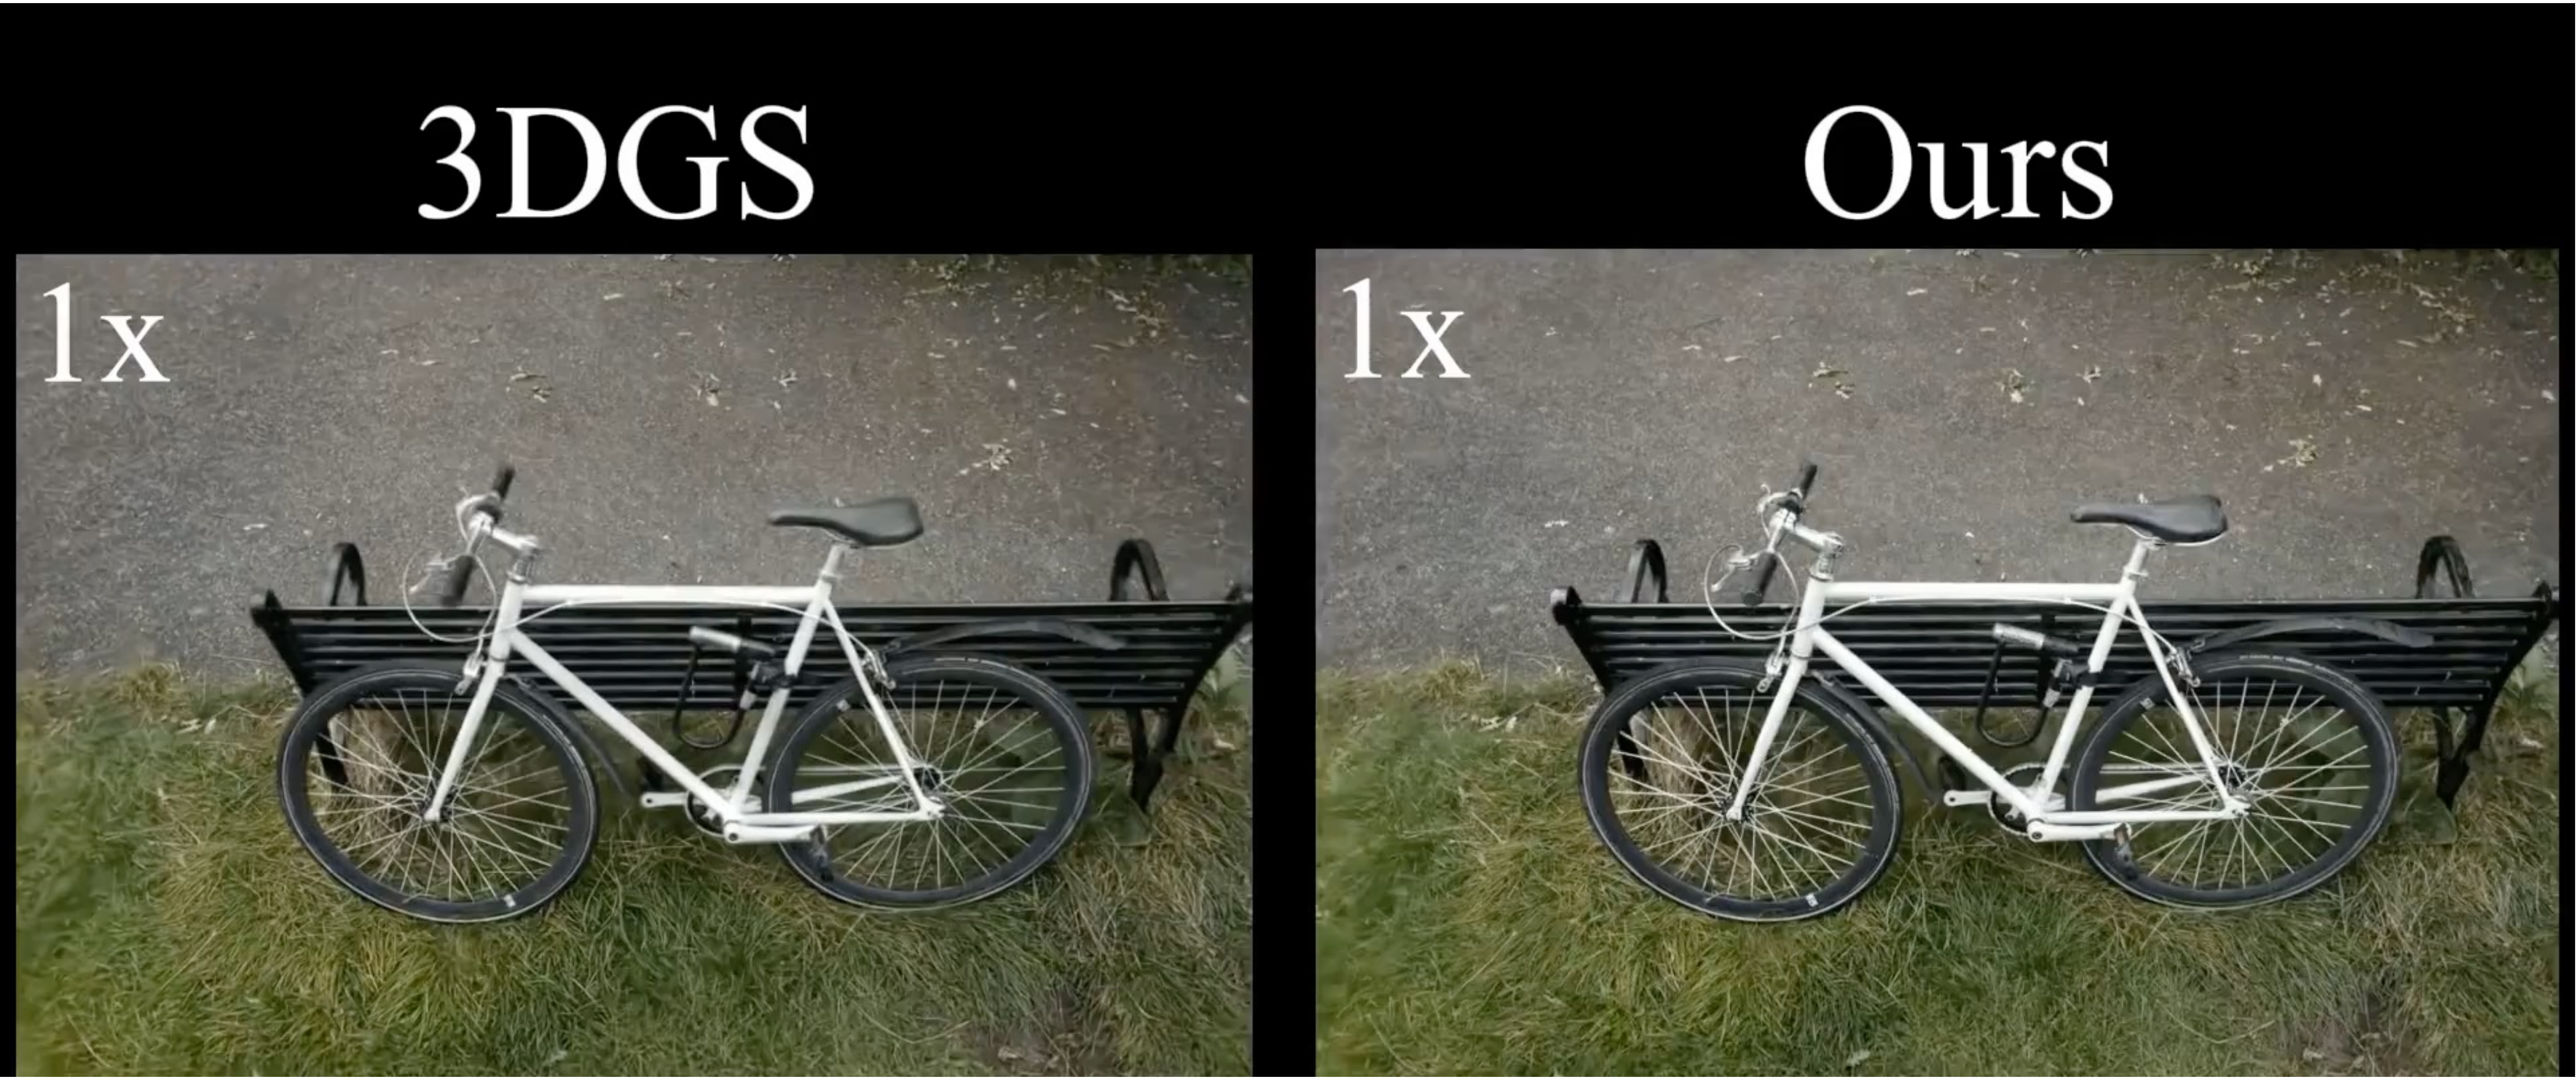

Analytic-Splatting: results

Analytic-Splatting: results

- MS-Mip-NeRF 360 dataset

	PSNR ↑					SSIM ↑					LPIPS ↓				
	Full Res.	1/2 Res.	1/4 Res.	1/8 Res.	Avg.	Full Res.	1/2 Res.	1/4 Res.	1/8 Res.	Avg.	Full Res.	1/2 Res.	1/4 Res.	1/8 Res.	Avg.
Mip-NeRF 360 [3]	27.50	29.19	30.45	30.86	29.50	0.778	0.864	0.912	0.931	0.871	0.254	0.136	0.077	0.058	0.131
Mip-NeRF 360 + iNGP	26.46	27.92	27.67	25.58	26.91	0.773	0.855	0.866	0.804	0.824	0.253	0.142	0.117	0.159	0.167
Zip-NeRF [4]	28.25	30.01	31.56	32.52	30.58	0.822	0.891	0.933	0.955	0.900	0.198	0.099	0.056	0.038	0.098
3DGS [16]	26.55	28.00	28.51	27.45	27.63	0.779	0.854	0.891	0.888	0.853	0.274	0.162	0.102	0.087	0.156
3DGS-SS [16]	27.20	28.75	29.89	29.71	28.89	0.800	0.871	0.914	0.928	0.878	0.246	0.138	0.081	0.061	0.131
Mip-Splatting [36]	27.20	28.74	29.90	30.66	29.12	0.802	0.870	0.915	0.944	0.883	0.244	0.146	0.090	0.056	0.134
Ours	27.50	28.99	30.35	31.21	29.51	0.808	0.874	0.919	0.945	0.887	0.231	0.132	0.077	0.051	0.123

Beyond Anti-Aliasing

- Technologies developed in anti-aliasing are useful beyond the anti-aliasing, *i.e.* cases that requires integration
- BRDF integration

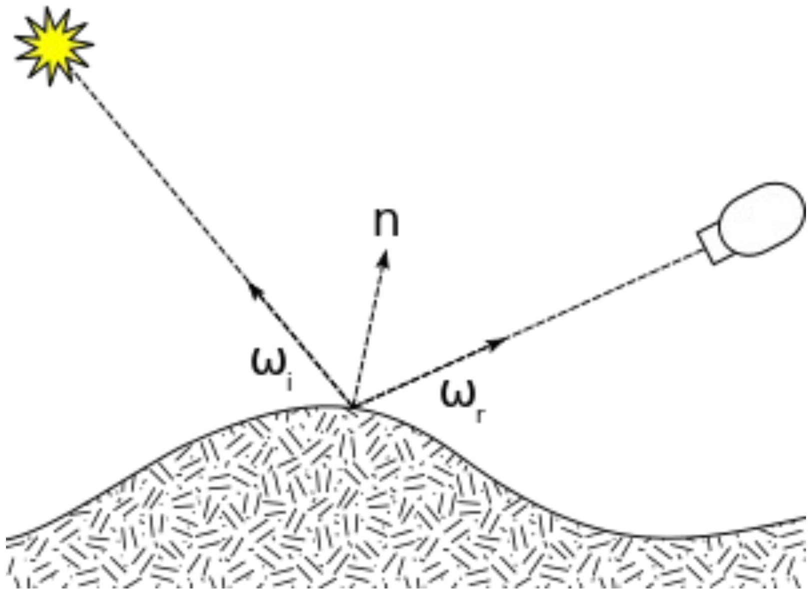

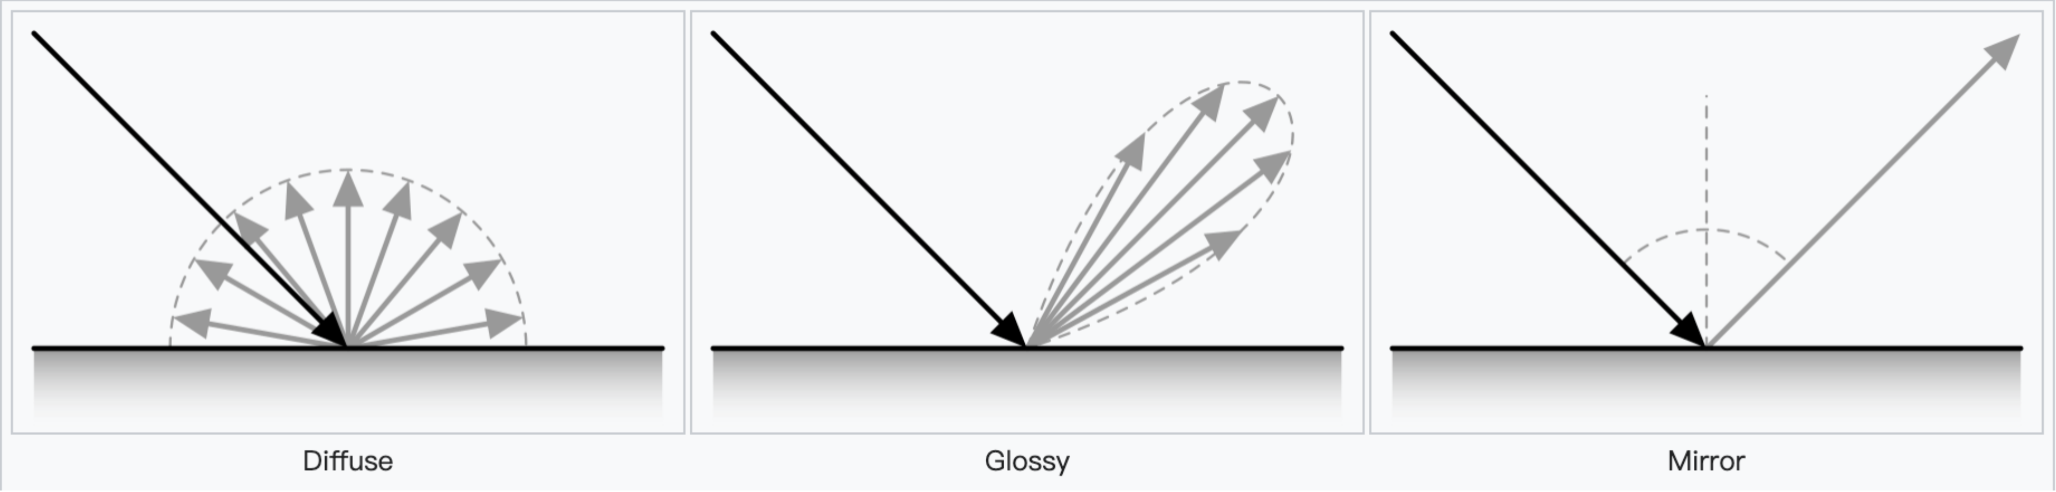

Inverse rendering

- Inverse Rendering of Glossy Objects via the Neural Plenoptic Function and Radiance Fields, CVPR'24

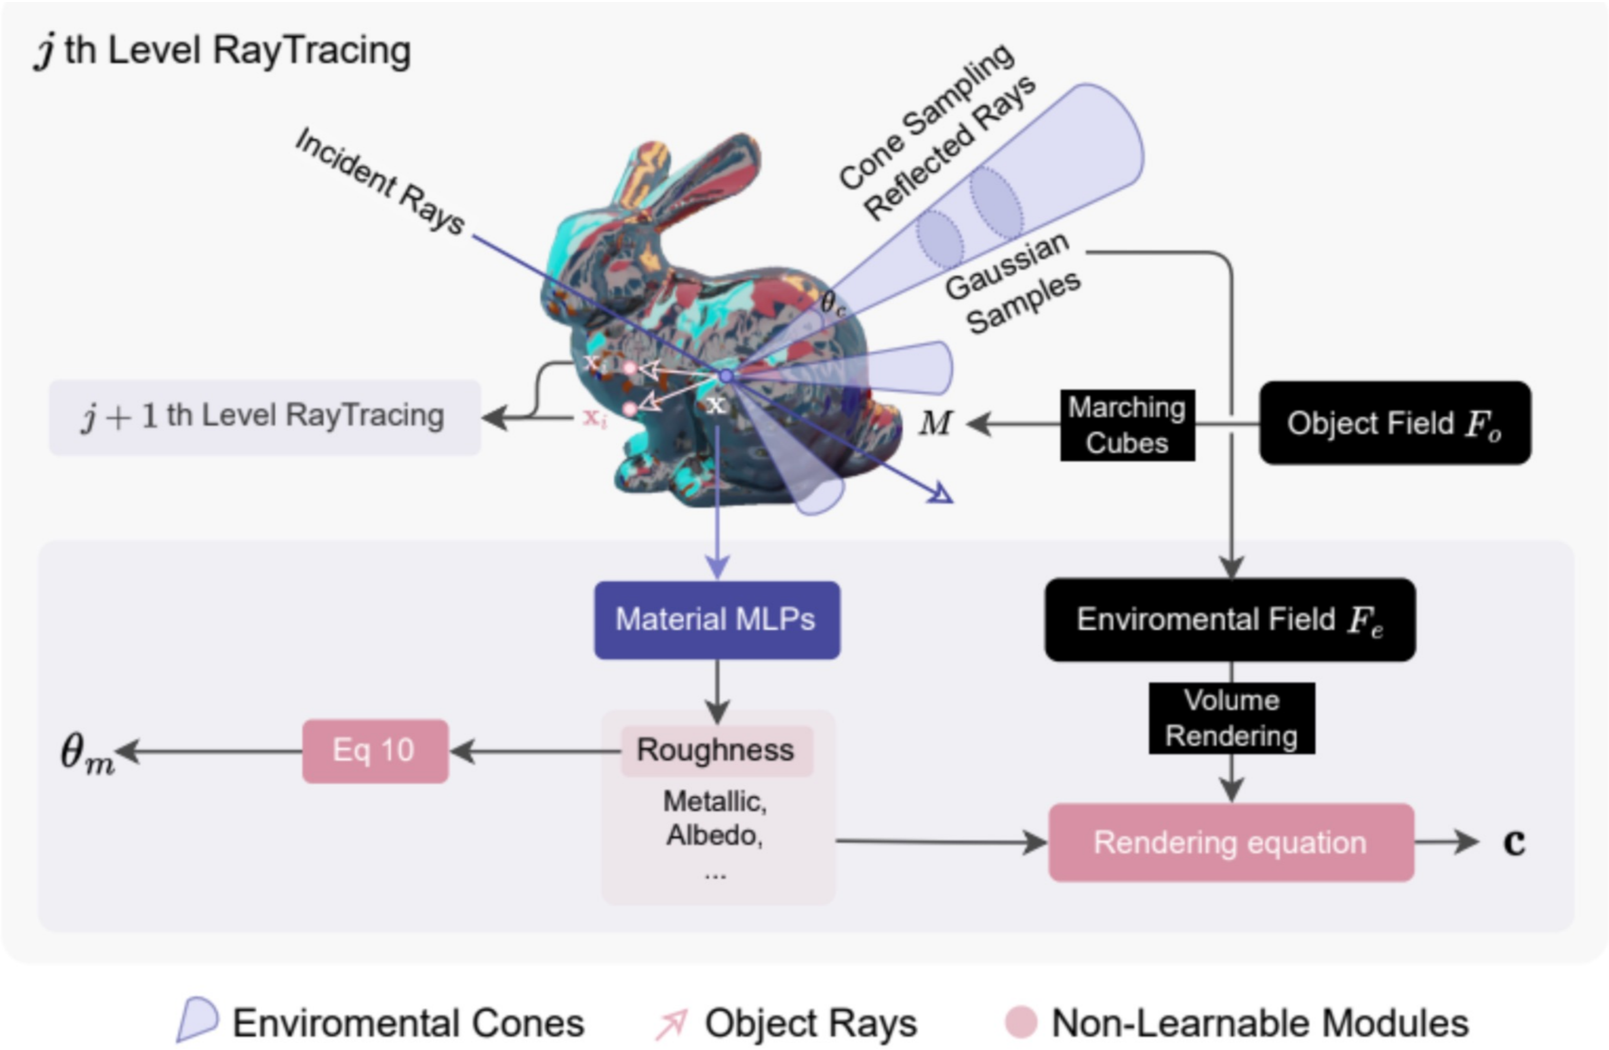

RGB

NeRO Ours


Roughness
NeRO
Ours
GT

Metallic

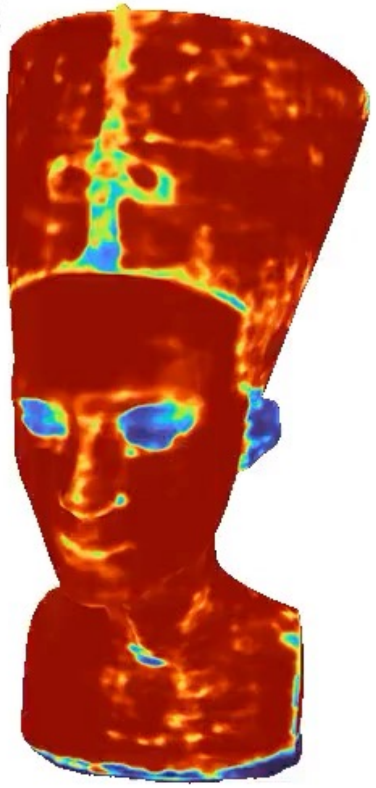

NeRO

Ours

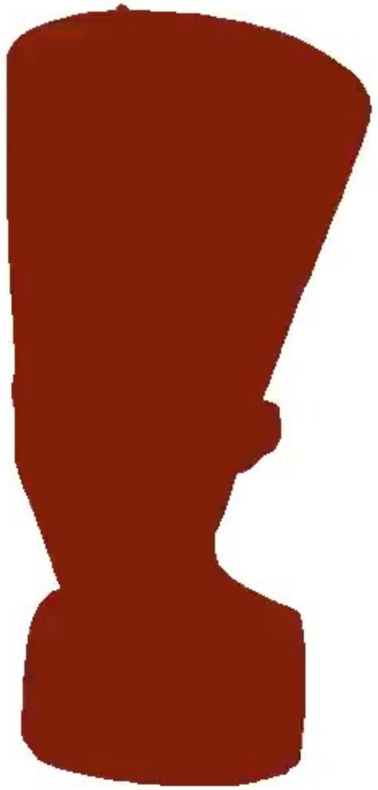

GT

Albedo

Diffuse Color

NeRO

Ours

Normal

NeRO

Ours

Lighting

NeRO

Ours

Summary

- Anti-Aliasing is crucial for the rendering quality in both NeRF and 3DGS
- Multi-sampling and Area-sampling are two main streams of strategies for anti-aliasing
- Anisotropy should be considered in the anti-aliasing
- Anti-aliasing technologies can be useful for other integration-related tasks

Thank You!

More resources are available on my webpage