

HDR-NeRF: High Dynamic Range Neural Radiance Fields

CVPR 2022

Xin Huang, Qi Zhang, Ying Feng, Hongdong Li, Xuan Wang, Qing Wang

First unit: Northwestern Polytechnical University

徐一伦

Website: <https://shsf0817.github.io/hdr-nerf/>

2022/10/16

Outline

- Authorship
- **Background**
- Method
- Experiment
- Conclusion

Background-Novel View Synthesis

Task description



Rendered view



Real scene



Captured views



Application



Virtual Reality



Sports Live



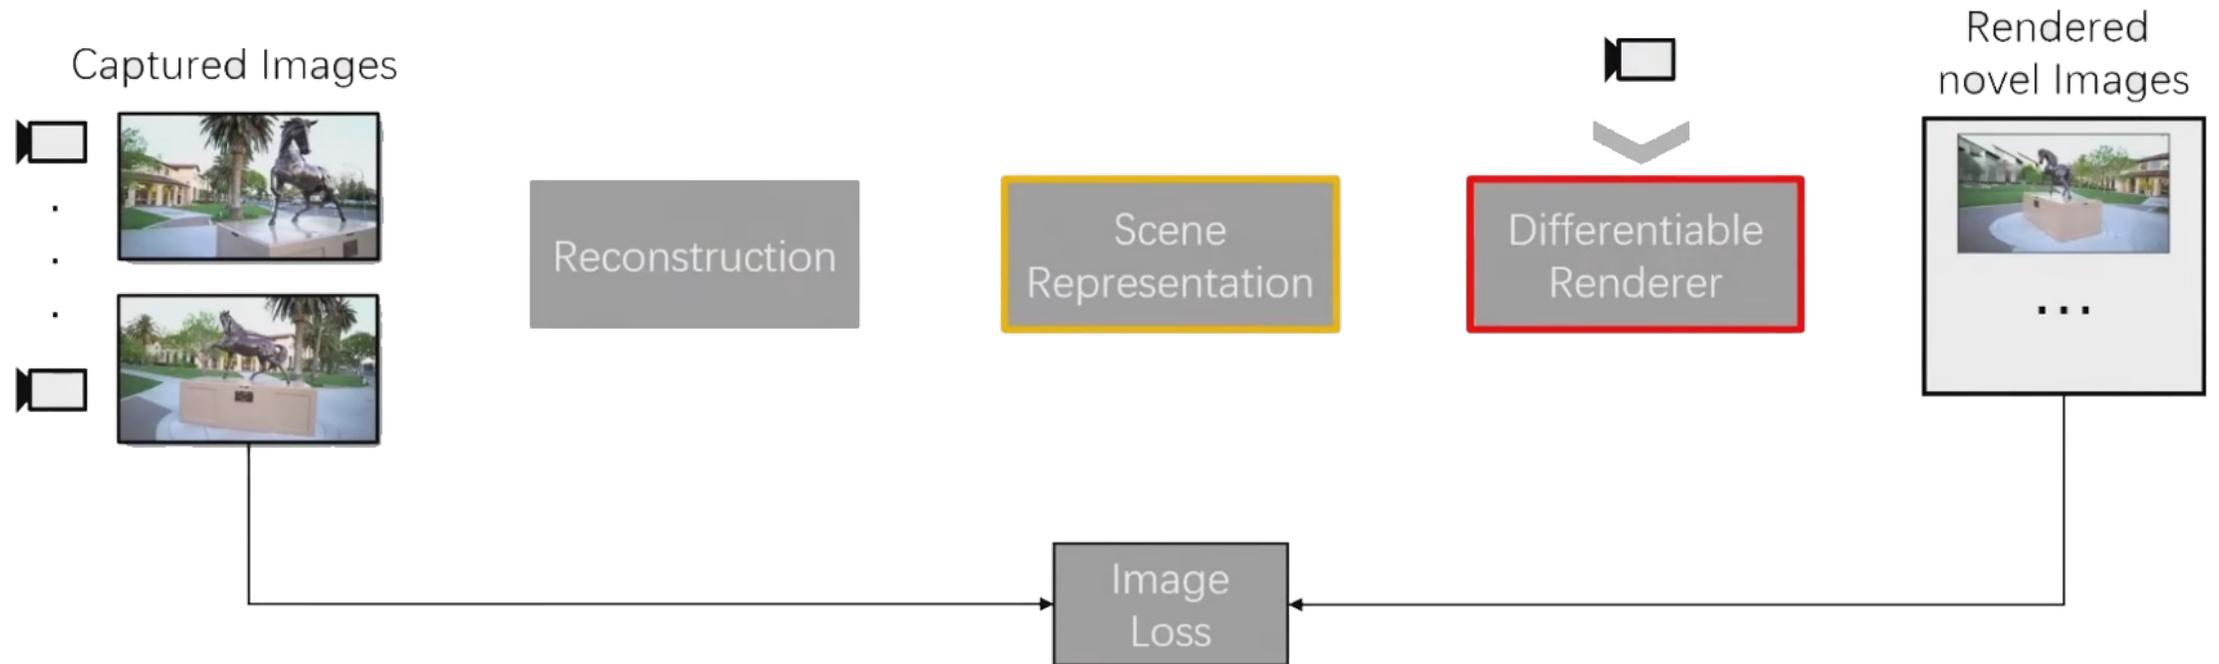
Facebook 3D Photo



Google Starline

Background-Novel View Synthesis

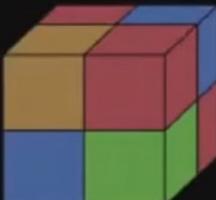
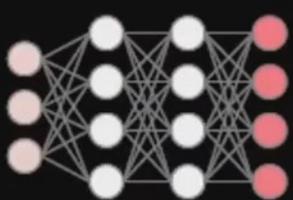
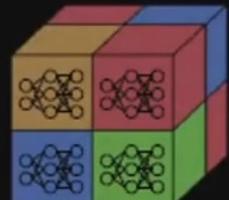
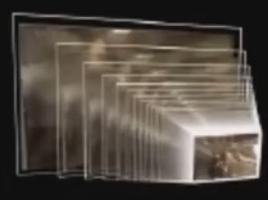
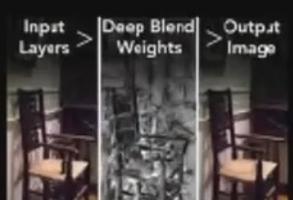
General pipeline



Scene representation + Differentiable renderer

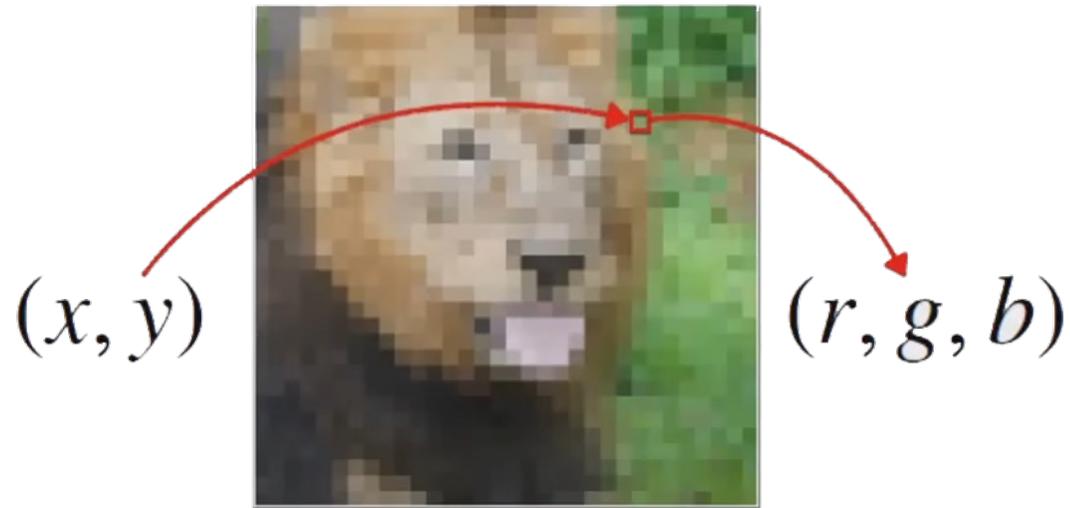
Background-Novel View Synthesis

Overview of Representations

					
Scene Representation	Voxelgrids	Implicit Function	Hybrid Implicit/Explicit	Multi-Plane Images	Image-based
Renderer	Volumetric Ray-based	Sphere-Tracing Volumetric	Volumetric	(Alpha) compositing	Rasterization / Volumetric
Pros	"True 3D" Fast	True 3D High quality	Significant Speedup Admits <i>local</i> priors	High-quality Fast	High-quality Fast
Cons	Memory $O(n^3)$	Extremely expensive, slow rendering	No <i>global</i> priors	Large Size Only 2.5D	Not compact: Needs source images.

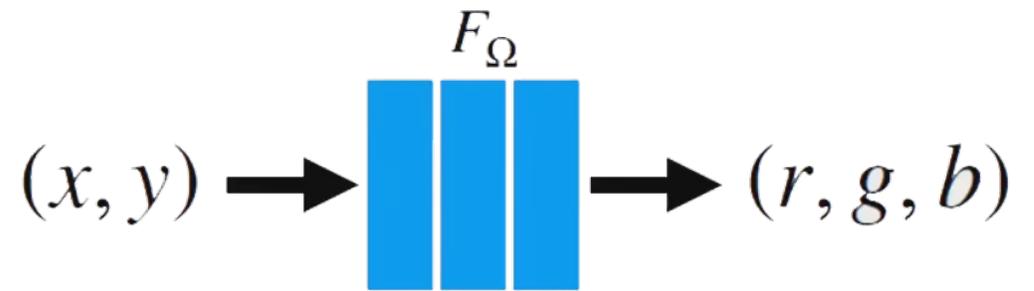
Background-Neural Radiance Fields (NeRF)

Explicit representation of 2D scene



Store an image as
a 2D grid of RGB color values

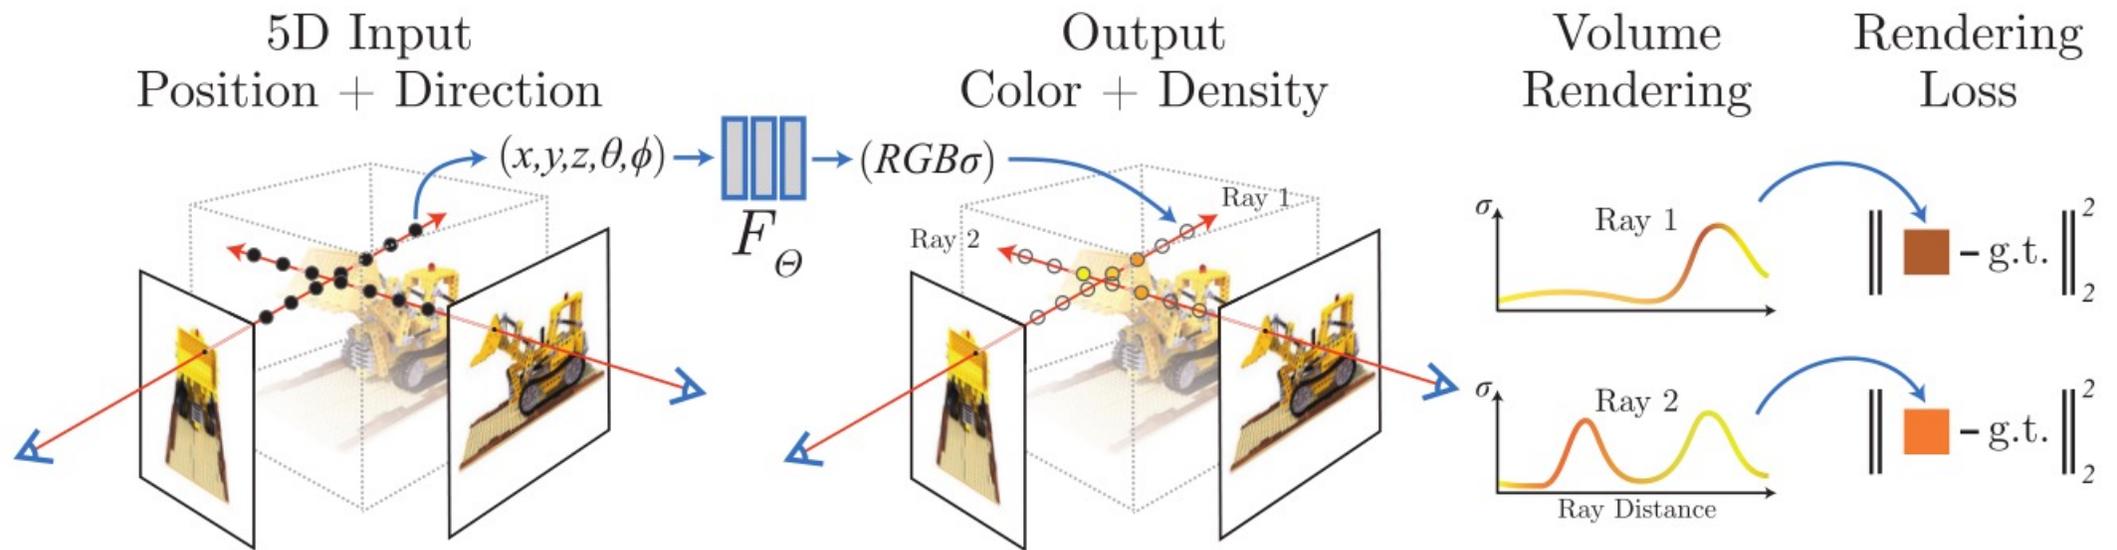
Implicit representation of 2D scene



Training a simple MLP to do this instead.

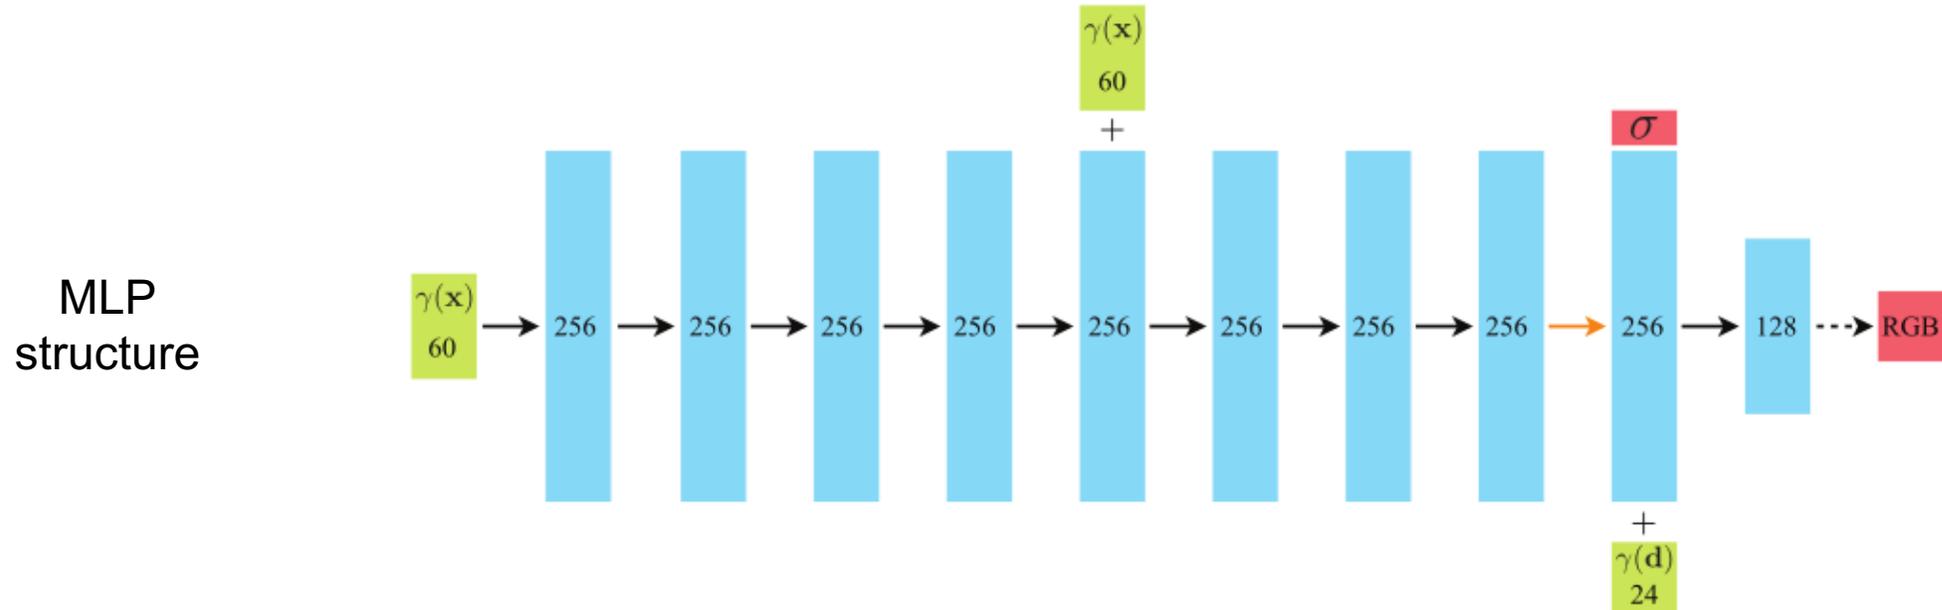
Background-Neural Radiance Fields (NeRF)

Build implicit representations of 3D scenes via MLP



Background-Neural Radiance Fields (NeRF)

Build implicit representations of 3D scenes via MLP

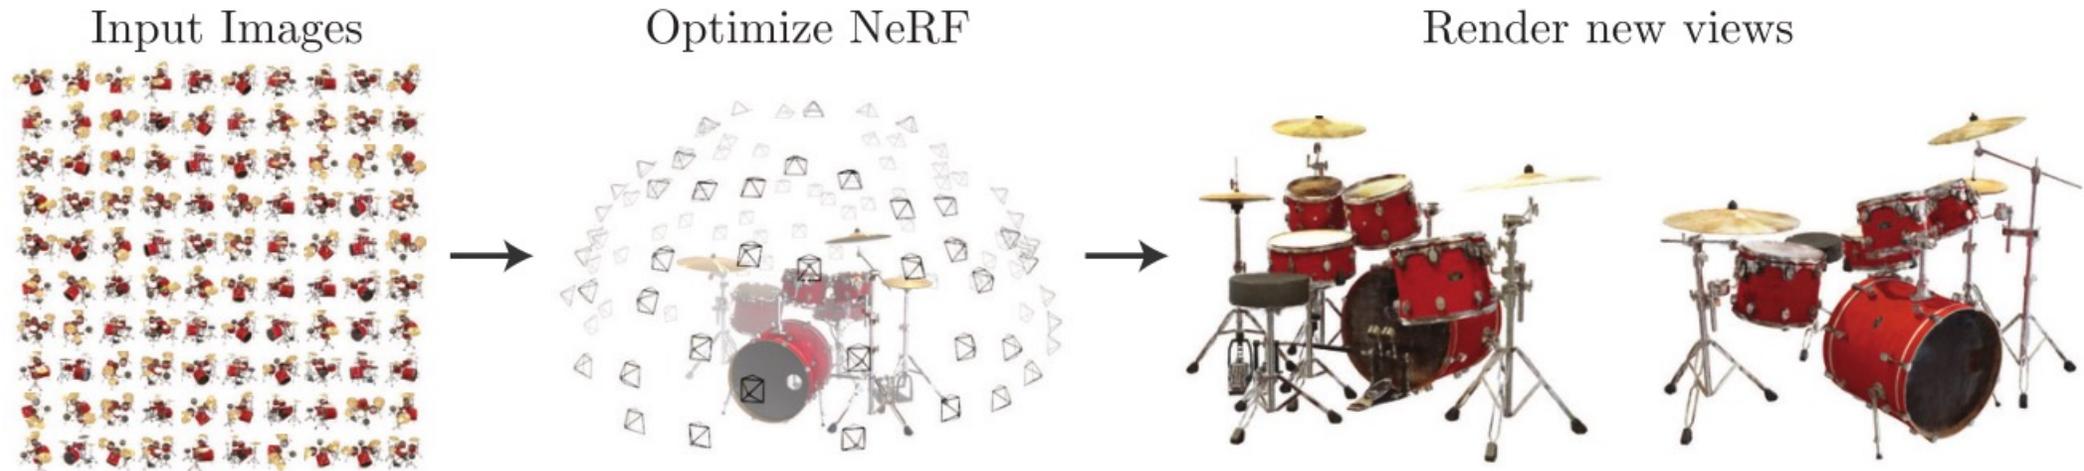


volume rendering

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

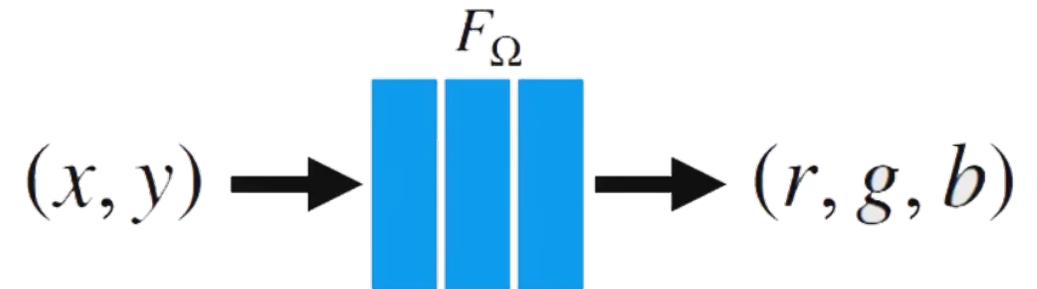
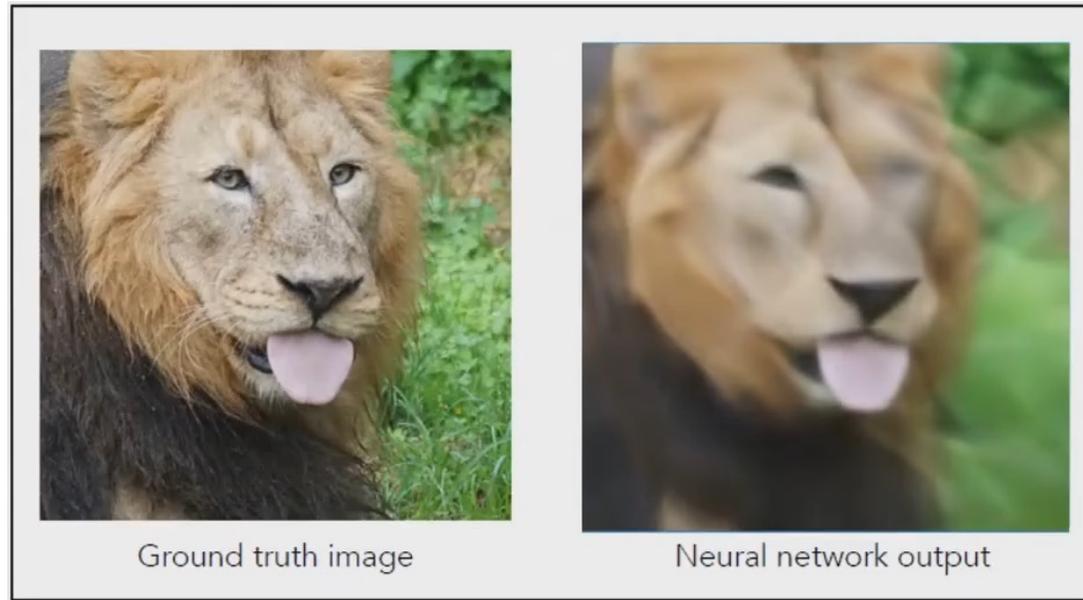
Background-Neural Radiance Fields (NeRF)

Overall pipeline



Background-Neural Radiance Fields (NeRF)

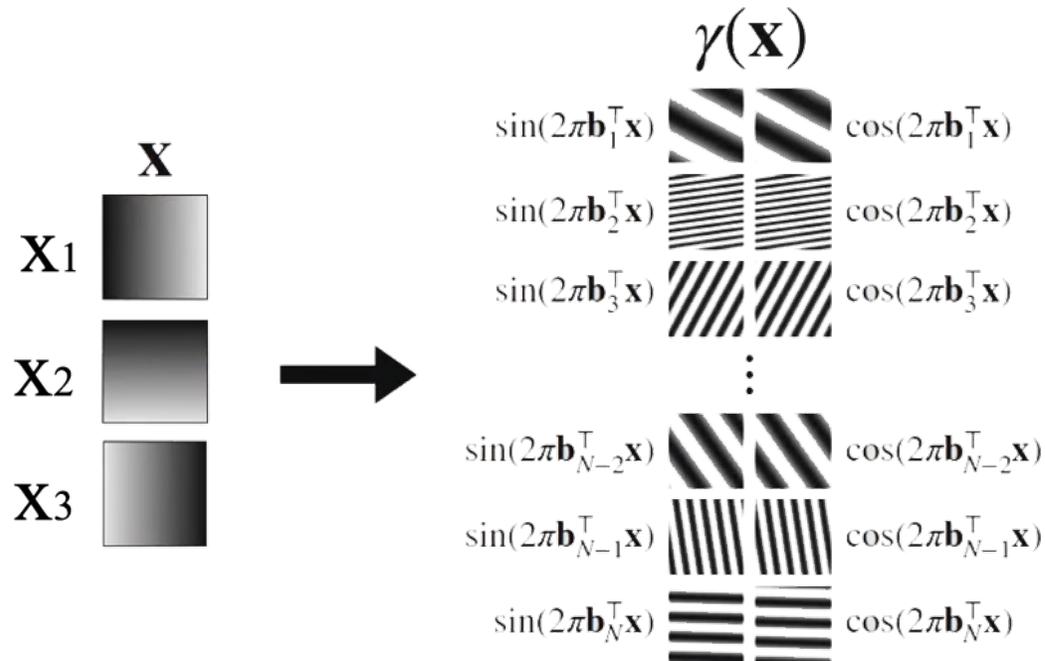
Optimizing the NerF: Positional encoding



How to learn high frequency information?

Background-Neural Radiance Fields (NeRF)

Optimizing the NerF: Positional encoding



$$\gamma(p) = (\sin(2^0\pi p), \cos(2^0\pi p), \dots, \sin(2^{L-1}\pi p), \cos(2^{L-1}\pi p))$$

Background-Neural Radiance Fields (NeRF)

Optimizing the NerF: Hierarchical volume sampling

“coarse” network

$$\hat{C}_c(\mathbf{r}) = \sum_{i=1}^{N_c} w_i c_i, \quad w_i = T_i(1 - \exp(-\sigma_i \delta_i))$$

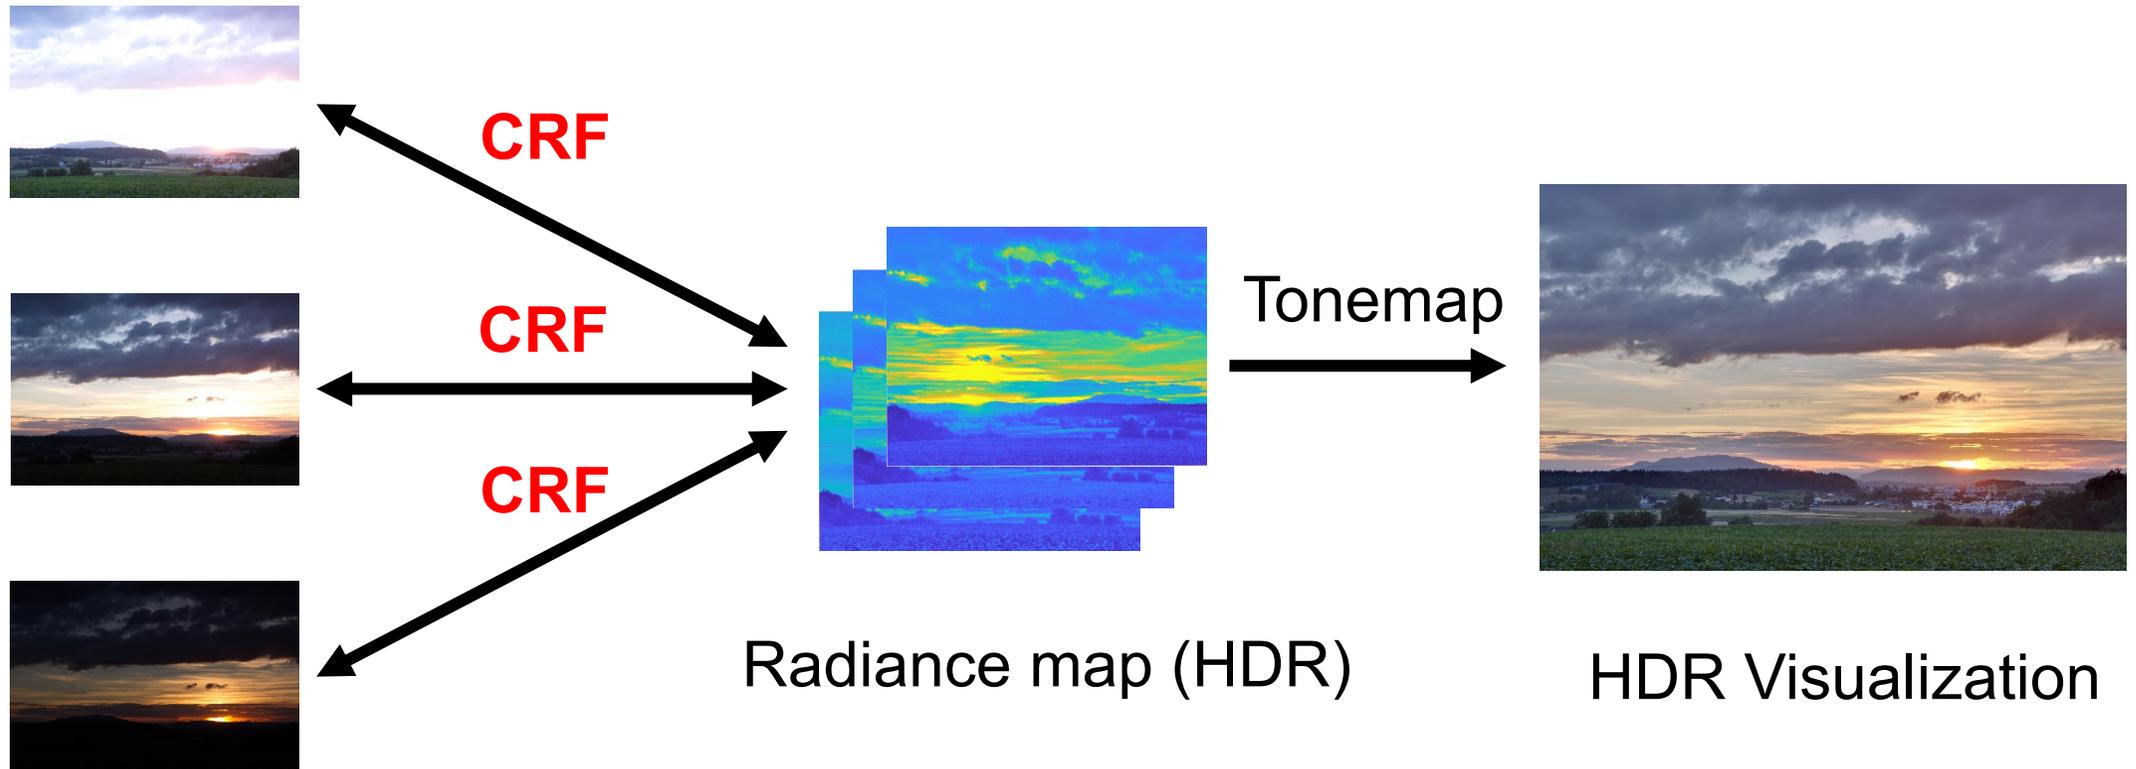
Speed up training

“coarse” network + “fine” network

$$\mathcal{L} = \sum_{\mathbf{r} \in \mathcal{R}} \left[\left\| \hat{C}_c(\mathbf{r}) - C(\mathbf{r}) \right\|_2^2 + \left\| \hat{C}_f(\mathbf{r}) - C(\mathbf{r}) \right\|_2^2 \right]$$

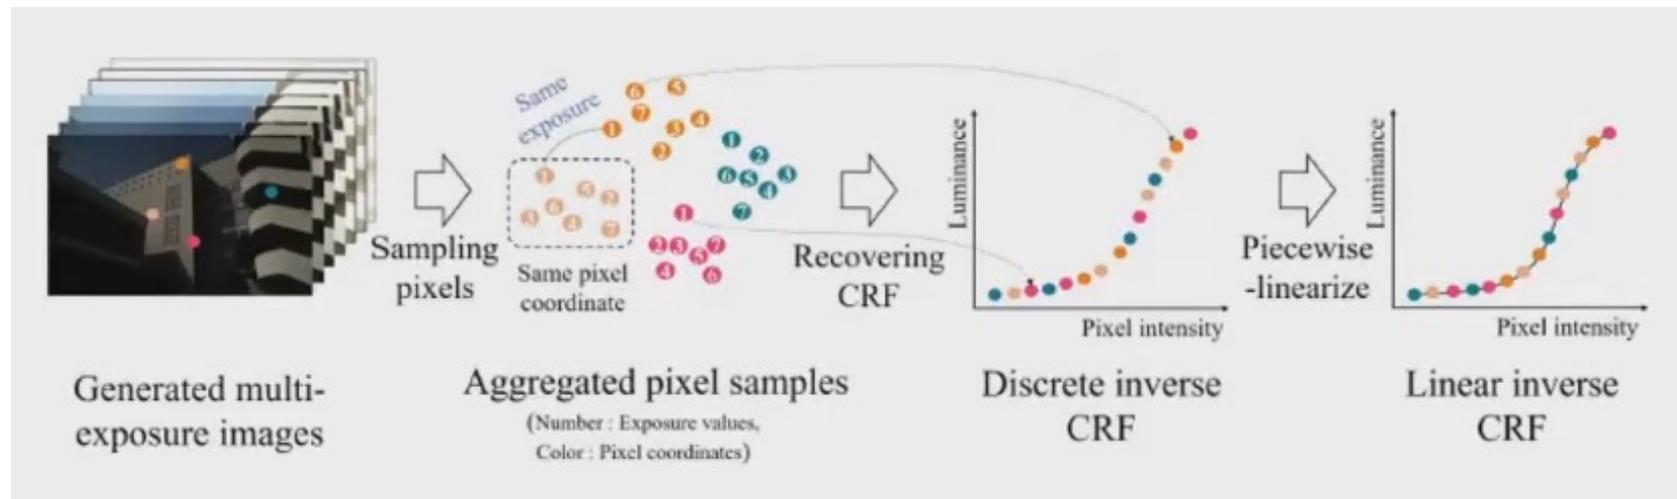
Background-High dynamic range imaging (HDRI)

A typical pipeline



Background-High dynamic range imaging (HDRI)

Traditional CRF Estimation Algorithms



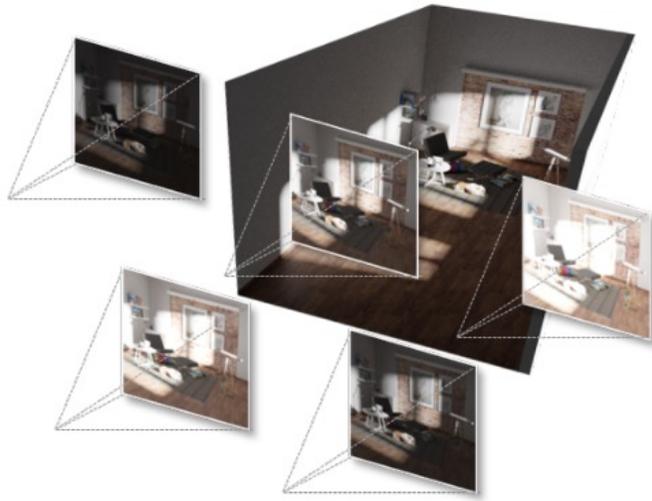
Limitation : Requires aligned multi-exposure images as input

Outline

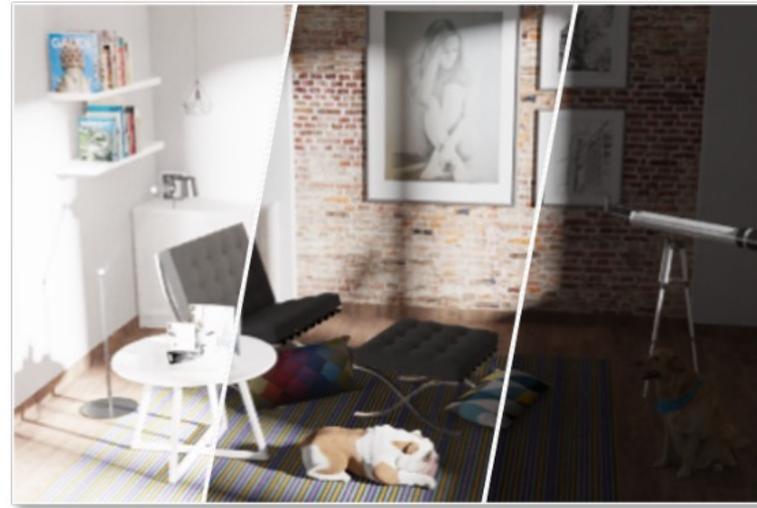
- Authorship
- Background
- **Method**
- Experiment
- Conclusion

Method

Task description



(a) Input views



(b) Novel LDR views

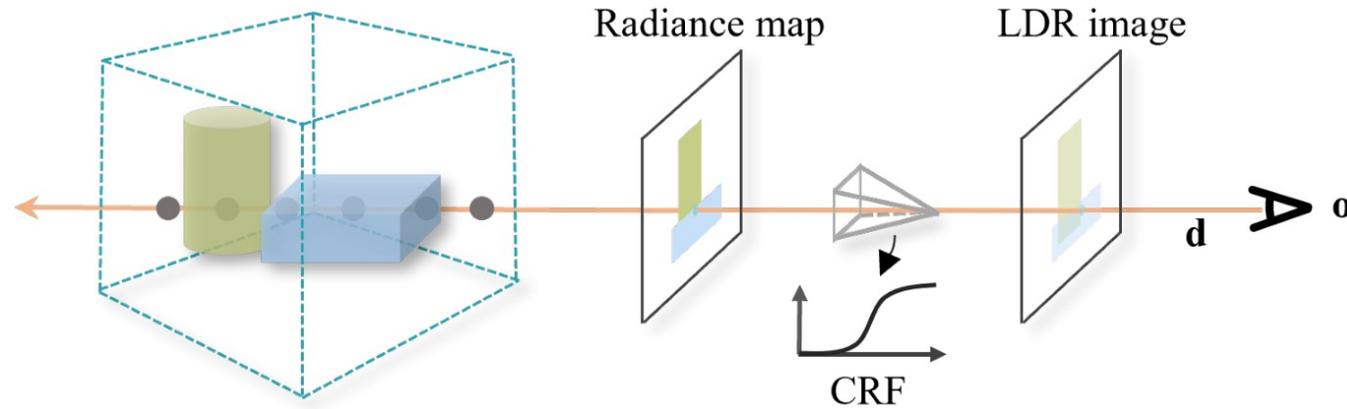


(c) One of the novel HDR views

Recover HDR radiance fields from LDR views with different exposure

Method

Radiance to color

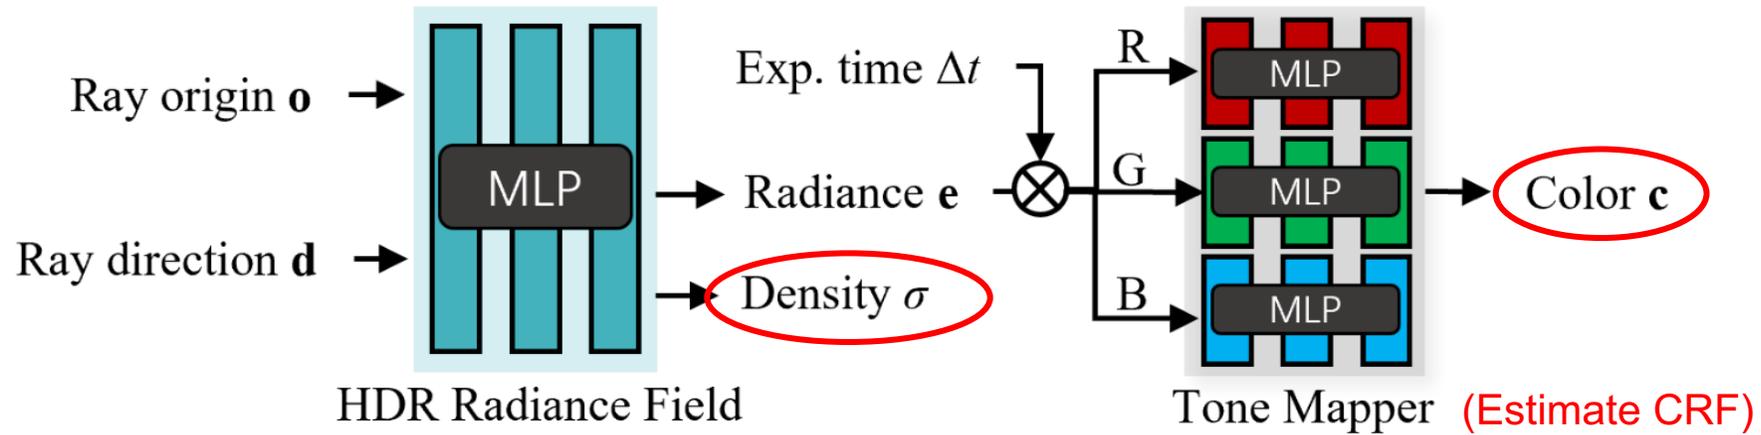


Combine volume rendering and tone-mapping

$$\left. \begin{aligned} \hat{C}(\mathbf{r}) &= \int_{s_n}^{s_f} T(s)\sigma(\mathbf{r}(s))\mathbf{c}(\mathbf{r}(s), \mathbf{d}) ds \\ \mathbf{c}(\mathbf{r}, \Delta t) &= g(\ln \mathbf{e}(\mathbf{r}) + \ln \Delta t(\mathbf{r})) \end{aligned} \right\} \hat{C}(\mathbf{r}, \Delta t) = \int_{s_n}^{s_f} T(s)\sigma(\mathbf{r}(s))g(\ln \mathbf{e}(\mathbf{r}(s)) + \ln \Delta t(\mathbf{r})) ds$$

Method

Pipeline

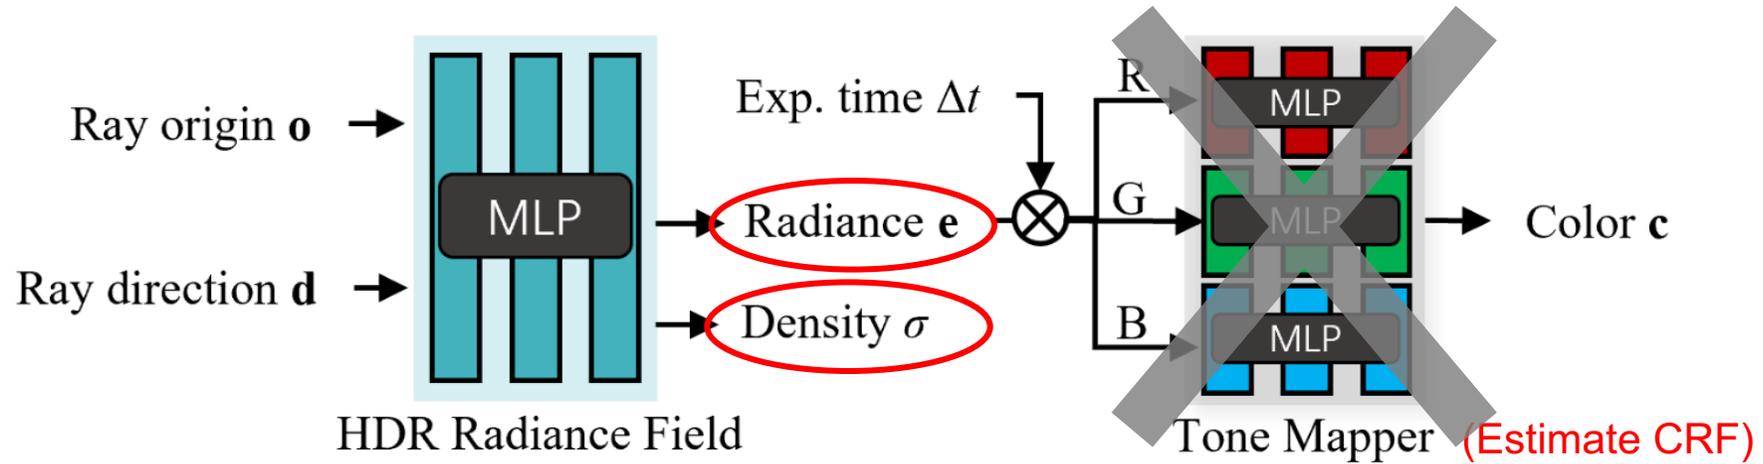


LDR Views Rendering

$$\hat{C}(\mathbf{r}, \Delta t) = \int_{s_n}^{s_f} T(s) \sigma(\mathbf{r}(s)) g(\ln \mathbf{e}(\mathbf{r}(s)) + \ln \Delta t(\mathbf{r})) ds$$

Method

Pipeline



HDR Views Rendering

$$\hat{E}(\mathbf{r}) = \int_{s_n}^{s_f} T(s) \sigma(\mathbf{r}(s)) \mathbf{e}(\mathbf{r}(s)) ds$$

Method

Loss Function

Color reconstruction loss

$$\mathcal{L}_c = \sum_{\mathbf{r} \in \mathcal{R}(\mathbf{P})} \|\hat{C}_c(\mathbf{r}, \Delta t) - C(\mathbf{r}, \Delta t)\|_2^2 + \|\hat{C}_f(\mathbf{r}, \Delta t) - C(\mathbf{r}, \Delta t)\|_2^2.$$

Unit exposure loss $\mathcal{L}_u = \|g(0) - C_0\|_2^2$ C_0 : the midway of the pixel value



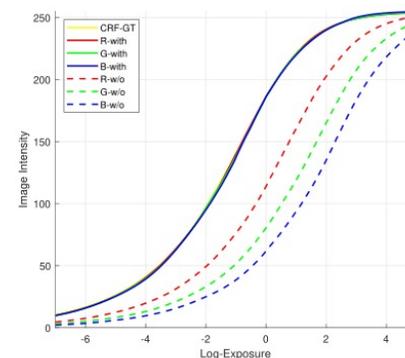
(a) with \mathcal{L}_u



(b) w/o \mathcal{L}_u



(c) GT



(d) CRFs

Outline

- Authorship
- Background
- Method
- **Experiment**
- Conclusion

Experiment

Novel Views Rendering



LDR GT



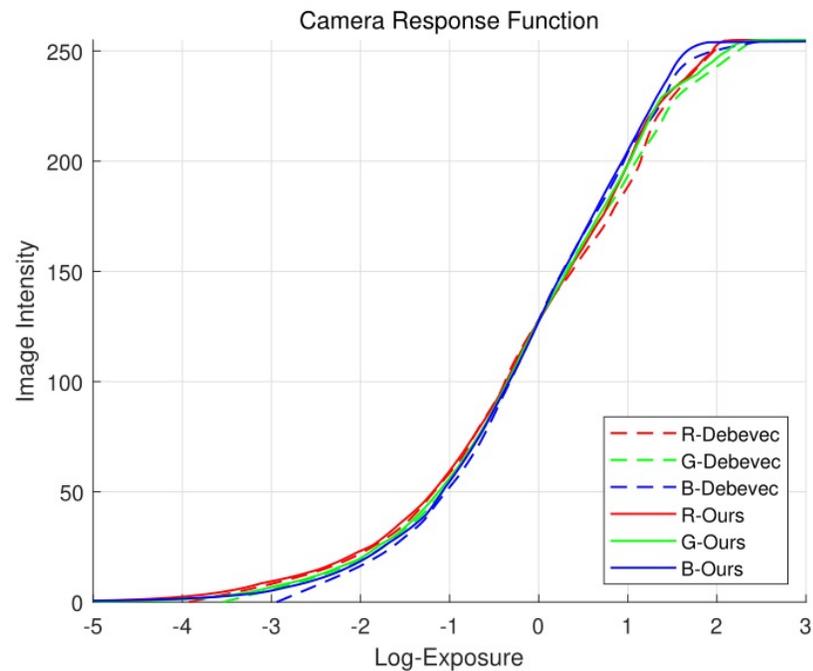
LDR output



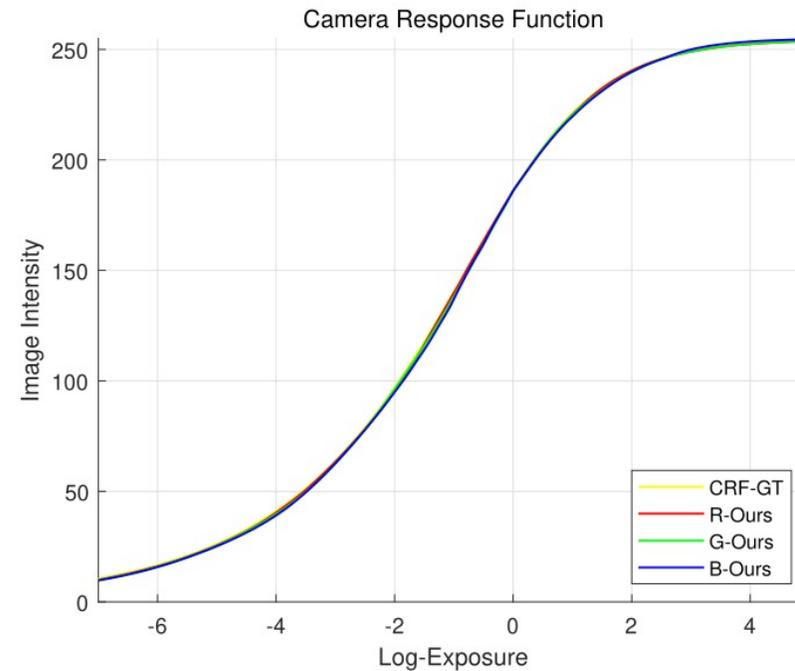
HDR output

Experiment

Estimated CRF by MLP



Comparisons with Debevec's method
in real scene



Comparisons with GT CRF
in synthetic scene

Experiment

Quantitative comparisons with baseline methods
on synthetic and real scenes

		LDR-OE (t_1, t_3, t_5)			LDR-NE (t_2, t_4)			HDR		
		PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
NeRF [42]	Syn.	13.97	0.555	0.376	—	—	—	—	—	—
	Real	14.95	0.661	0.308	—	—	—	—	—	—
NeRF-W ¹ [37]	Syn.	29.83	0.936	0.047	29.22	0.927	0.050	—	—	—
	Real	28.55	0.927	0.094	28.64	0.923	0.089	—	—	—
NeRF-GT ² [42]	Syn.	37.66	0.965	0.028	35.87	0.955	0.032	37.80	0.964	0.029
	Real	34.55	0.958	0.057	34.59	0.956	0.051	—	—	—
Ours [†]	Syn.	—	—	—	—	—	—	—	—	—
	Real	30.37	0.944	0.075	29.37	0.938	0.078	—	—	—
Ours	Syn.	39.07	0.973	0.026	37.53	0.966	0.024	36.40	0.936	0.018
	Real	31.63	0.948	0.069	31.43	0.943	0.069	—	—	—

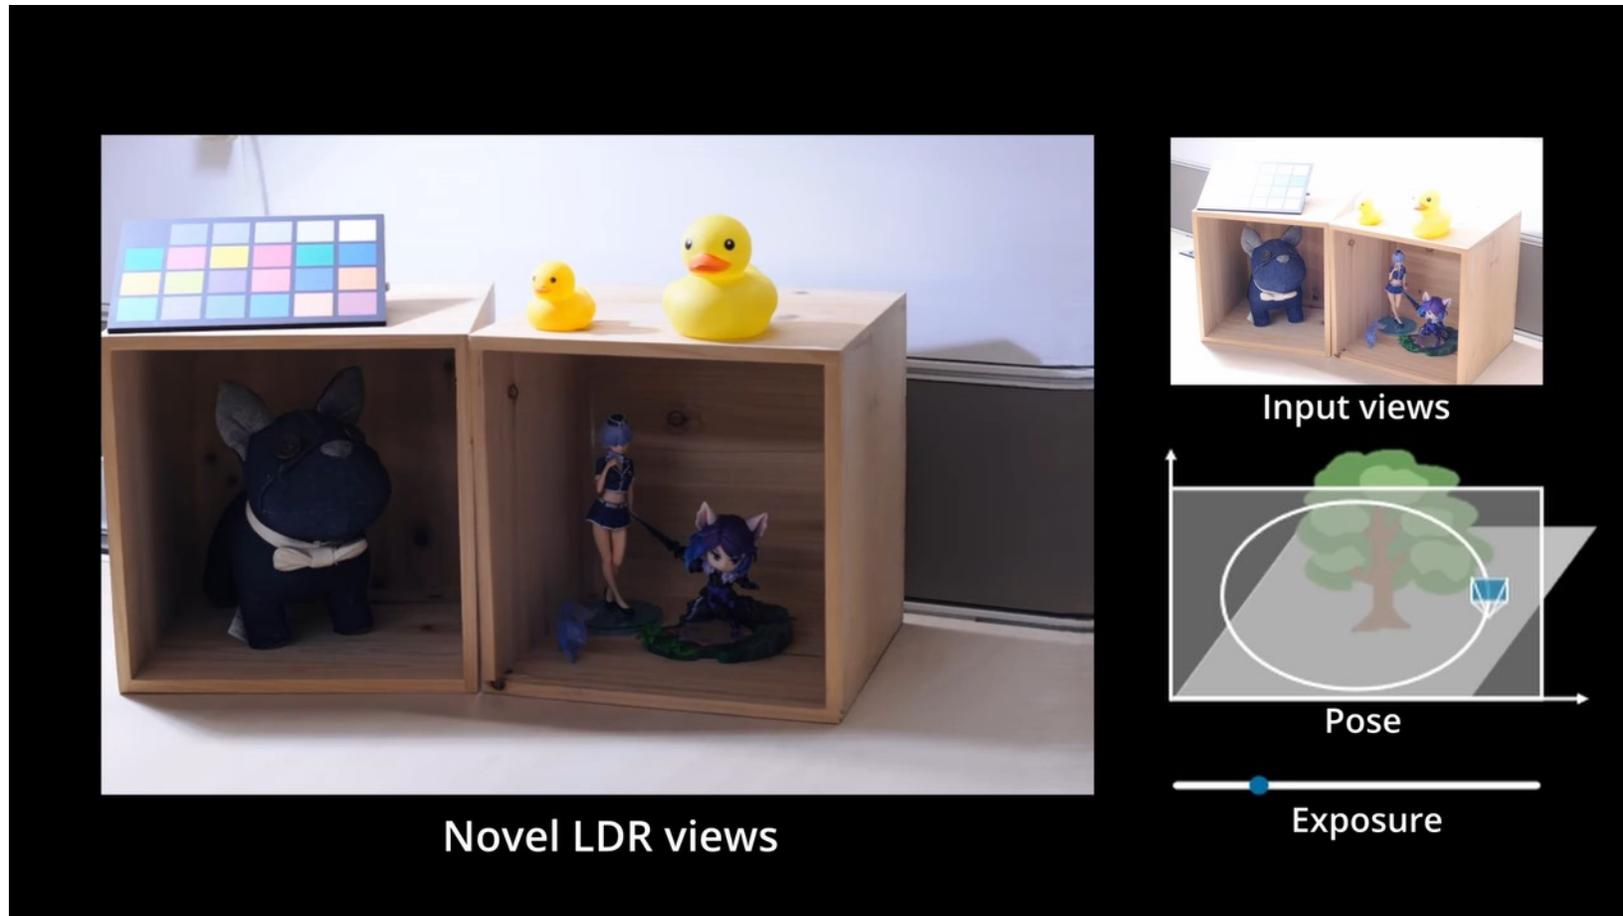
¹ The exposures of input views for NeRF-W are randomly selected from all five exposures to learn five appearance vectors for testing.

² A version of NeRF (as the upper bound of our method) that is trained from LDR images with consistent exposures or HDR images.

[†] An ablation study of our method that models the tone-mapping operations of RGB channels with a single MLP.

Experiment

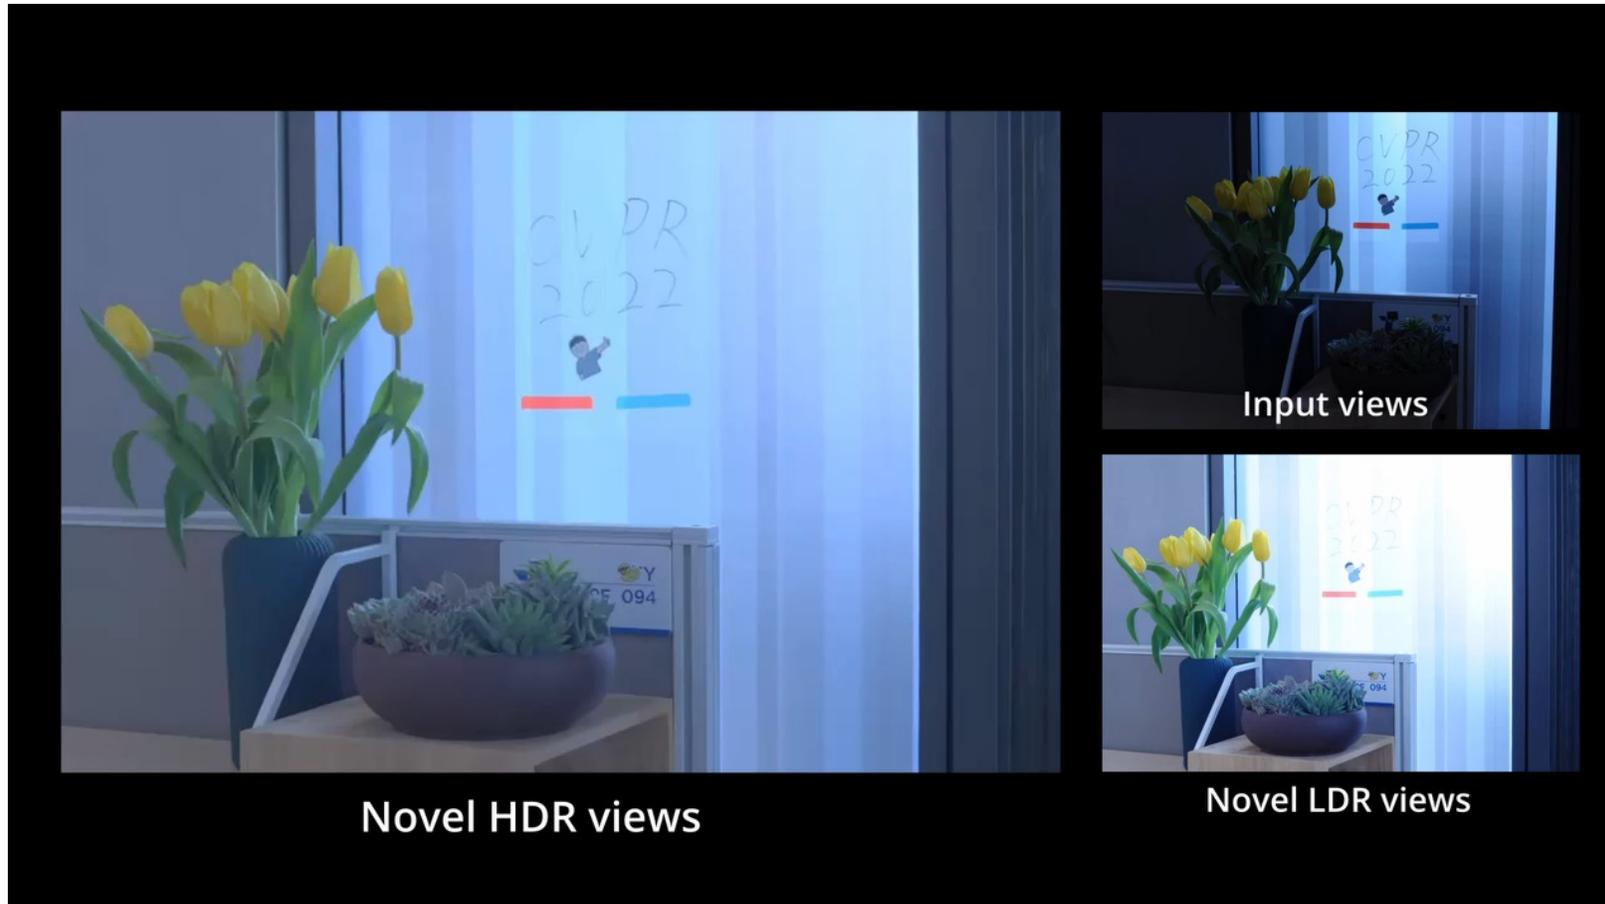
Novel LDR views



Video website: <https://shsf0817.github.io/hdr-nerf/images/ldr1.mp4>

Experiment

Novel HDR views (Tone-mapped)



Video website: <https://shsf0817.github.io/hdr-nerf/images/hdr1.mp4>

Outline

- Authorship
- Background
- Method
- Experiment
- **Conclusion**

Conclusion

Contribution:

- Proposing a novel method to recover the high dynamic range neural radiance field from a set of LDR views with different exposures.
- Rendering novel HDR views without ground-truth HDR supervision.
- Producing high-fidelity LDR views with specified exposures.

Future work:

- Dynamic 3D scene
- HDR Video
- Joint with denoise in raw data

Thanks for watching.

徐一伦

yi lunxu_buaa@163. com