

Dynamic Novel View Synthesis in High Dynamic Range

<https://github.com/prinasi/HDR-4DGS>

Kaixuan Zhang¹, Zhipeng Xiong¹, Minxian Li^{1 2}, Mingwu Ren^{1 2}, Jiankang Deng³, Xiatian Zhu⁴

¹Nanjing University of Science and Technology

²State Key Laboratory of Intelligent Manufacturing of Advanced Construction Machinery

³Imperial College London

⁴University of Surrey

➤ Introduction

Low Dynamic Range (LDR) images typically refers to images that can only represent a **limited range of luminance values** (usually, pixel values ranging from 0-255).

High Dynamic Range (HDR) images typically refers to images that represent a much wider range of luminance values—from the darkest shadows to the brightest highlights (usually ranges from 0 to $+\infty$).




Low Dynamic Range




High Dynamic Range

➤ Introduction

Novel View Synthesis (NVS) typically reconstructs *static* 3D scene representations from multi-view LDR images to enable photorealistic *LDR* rendering from arbitrary viewpoints.

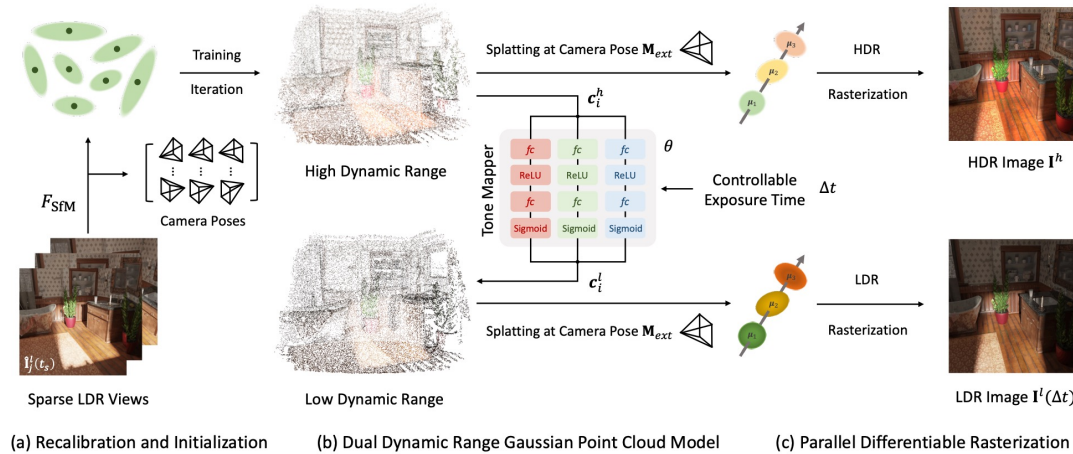


Dynamic Novel View Synthesis (DNVS) typically reconstructs *dynamic* 4D scene representations from multi-view LDR videos to enable photorealistic *LDR* rendering from arbitrary viewpoints and timestamps.

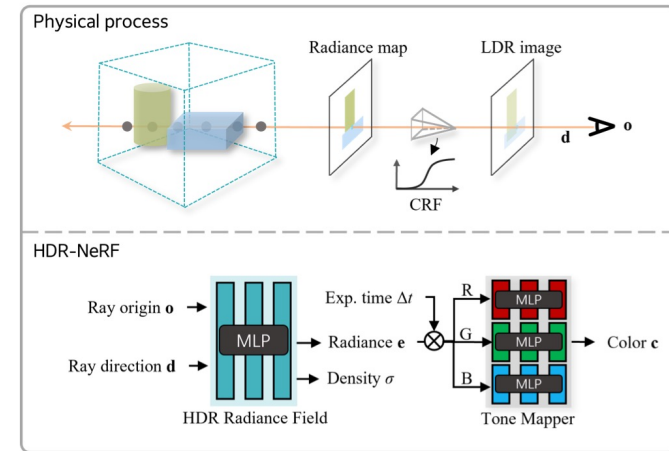


➤ Introduction

HDR Novel View Synthesis (HDR NVS) typically reconstructs *static* HDR 3D scene representations from multi-view and multi-exposure LDR images to enable photorealistic HDR rendering from arbitrary viewpoints.



Pipeline of HDR-GS.



Pipeline of HDR-NeRF.


They cannot reconstruct *dynamic* scenes!

➤ Problem Formulation

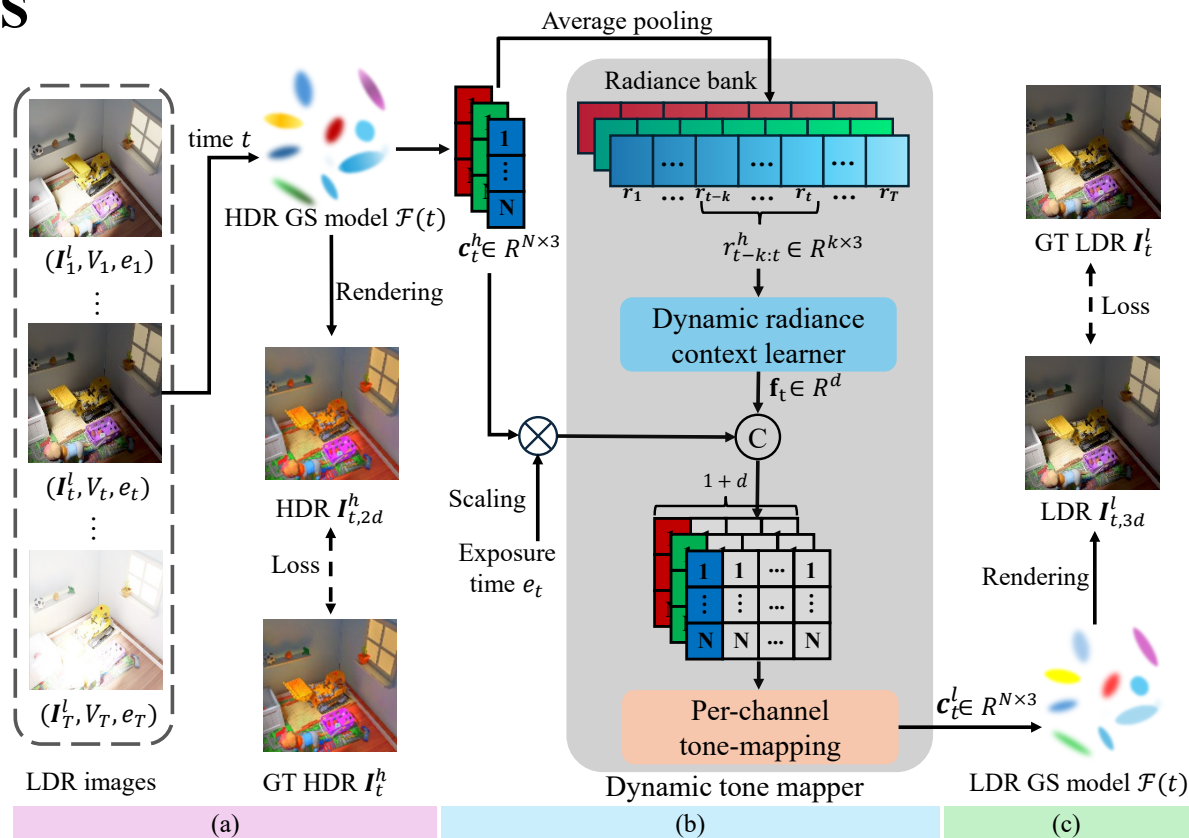
We introduce *High Dynamic Range Dynamic Novel View Synthesis (HDR DNVS)*, a new, more practical task that seeks to reconstruct temporally coherent HDR radiance fields and dynamic geometry from sparse, time-varying LDR inputs. HDR DNVS demands the joint modeling of evolving scene structure and HDR radiance, posing both geometric and photometric challenges absent in prior static or LDR-constrained settings.

➤ Definition

In HDR DNVS, we aim to learn an HDR 3D model \mathcal{F} for a target dynamic scene G , $\mathcal{F} : (t', V') \rightarrow \mathbf{I}_{t', V'}^h$, that would render an HDR image $\mathbf{I}_{t', V'}^h$ for any timestamp t' and viewpoint V' . To that end, we capture a set of multi-exposure LDR training images $\mathbf{I}^l = \{\mathbf{I}_1^l, \dots, \mathbf{I}_t^l, \dots, \mathbf{I}_T^l\}$ each \mathbf{I}_t^l associated with the exposure time e_t , and the camera viewpoint V_t . We may have optional access to coupled HDR training data $\mathbf{I}^h = \{\mathbf{I}_1^h, \dots, \mathbf{I}_t^h, \dots, \mathbf{I}_T^h\}$.



➤ Overview of HDR-4DGS



(a) Given a sequence of LDR training images, each annotated with camera pose and exposure time, the goal is to learn an HDR Gaussian Splatting model \mathcal{F} . To address scene dynamics and LDR-HDR domain differences, (b) a Dynamic Tone Mapper (DTM) is introduced. DTM bridges the HDR and LDR color domains by modeling temporally varying radiance through a per-channel tone mapping strategy. Central to DTM is a dynamic radiance context learner (DRCL) that extracts a temporal context embedding f_t based on the radiance distribution over a k -frame sliding time window. For each color channel, the tone mapping input is constructed by scaling the HDR attributes by e_t and concatenating with the context embedding f_t . The DTM then predicts the corresponding LDR variant suitable for image rendering. (c) During training, losses are computed on the synthesized LDR outputs and optionally on HDR reconstructions. At inference time, given a specific timestamp and viewpoint, \mathcal{F} directly renders HDR images. To obtain LDR renderings under a specific exposure, the pretrained DTM, along with the radiance bank, converts HDR attributes into LDR color values for domain translation.

➤ Datasets


Due to no benchmarks for HDR DNVS, we introduce two complementary datasets: HDR-4D-Syn and HDR-4D-Real. HDR-4D-Syn consists of 8 synthetic dynamic scenes adapted from the dataset released by HDR-HexPlane. The real-world dataset, HDR-4D-Real, captures 4 dynamic indoor scenes in real-world settings. Ground truth HDR images are generated UltraFusion, ensuring realistic and high-quality HDR reconstructions.

Table 7: Statistics of HDR-4D-Syn, HDR / LDR means the number of HDR / LDR images.

Scenes	Training			Testing			Cameras	Resolution	Exposure Time (s)
	Frames	HDR	LDR	Frames	HDR	LDR			
Airplane	280	280	280	70	70	70	1	800×800	0.125/2/32
Deer	80	80	80	20	20	20	1	800×800	0.125/2/32
Hook	28	280	840	4	40	120	10	800×800	0.125/2/32
Jump	21	210	630	3	30	90	10	800×800	0.125/2/32
Lego	240	240	240	60	60	60	1	800×800	0.125/2/32
Mutant	135	405	405	14	42	42	3	800×800	0.125/2/32
Standup	51	255	765	8	40	120	5	800×800	0.125/2/32
Tank	136	136	136	34	34	34	1	800×800	0.125/2/32

Table 8: Statistics of HDR-4D-Real, HDR / LDR means the number of HDR / LDR images.

Scenes	Training			Testing			Cameras	Resolution	Exposure Time (s)
	Frames	HDR	LDR	Frames	HDR	LDR			
Bed	22	132	317	22	79	79	6	4032×3024	0.004/0.008/0.017
Excavator	40	240	576	40	144	144	6	4032×3024	0.002/0.008/0.033
Tank	20	40	96	20	24	24	2	4032×3024	0.007/0.02/0.05
Toys	20	40	96	20	24	24	2	4032×3024	0.008/0.02/0.1



➤ Experiments

Table 1: Results on HDR-4D-Syn. *: HDR only supervision; †: LDR+HDR supervision.

Row	Method	Supervision	PSNR↑	HDR SSIM↑	LPIPS↓	PSNR↑	LDR SSIM↑	LPIPS↓	Training time (min)	Inference speed (fps)
1	HexPlane	LDR	-	-	-	14.20	0.564	0.551	28.45	1.60
2	HexPlane*	HDR	24.89	0.771	0.377	-	-	-	59.64	1.94
3	4DGS	LDR	-	-	-	13.92	0.549	0.281	116.25	75.82
4	4DGS*	HDR	22.40	0.650	0.345	-	-	-	44.58	172.61
5	HDR-NeRF	LDR	8.54	0.062	0.552	21.66	0.664	0.553	212.83	0.061
6	HDR-GS	LDR	4.64	0.158	0.645	6.45	0.272	0.423	13.88	380.38
7	HDR-GS†	LDR+HDR	14.33	0.360	0.482	10.84	0.372	0.378	22.50	255.21
8	HDR-HexPlane	LDR	14.70	0.649	0.287	32.59	0.912	0.145	37.83	1.61
9	HDR-HexPlane†	LDR+HDR	29.30	0.844	0.223	31.09	0.896	0.185	54.31	1.33
10	HDR-4DGS (Ours)	LDR	25.88	0.865	0.076	33.16	0.949	0.055	69.38	40.80
11	HDR-4DGS† (Ours)	LDR+HDR	30.40	0.914	0.097	30.69	0.927	0.097	76.86	48.63

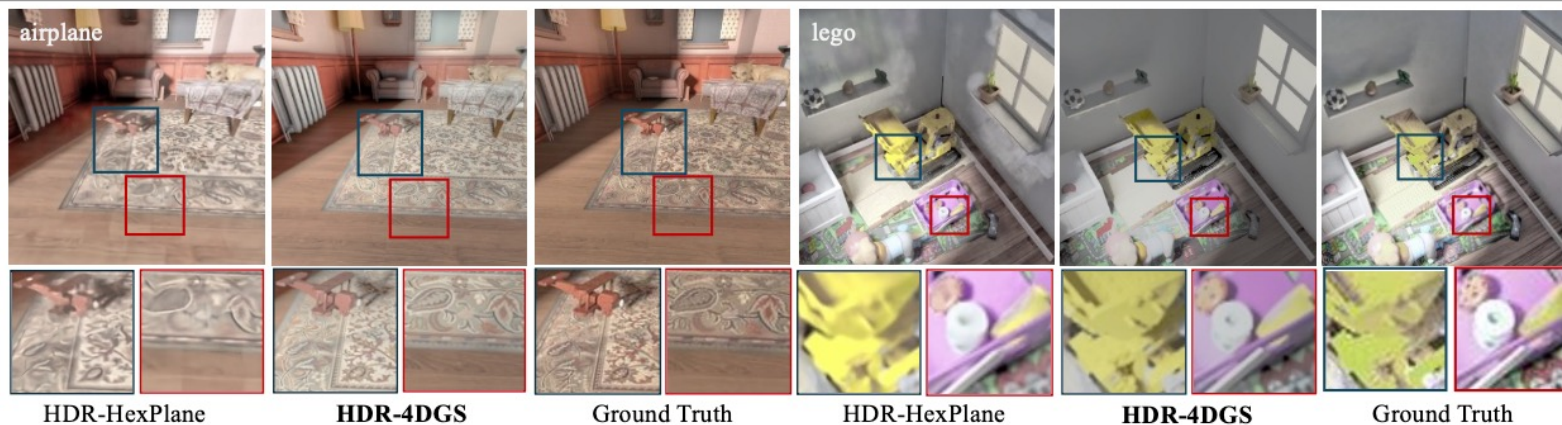


Figure 2: Visual comparison of HDR DNVS on HDR-4D-Syn.

➤ Experiments

Table 2: Results on HDR-4D-Real. *: HDR only supervision; †: LDR+HDR supervision.

Row	Method	Supervision	PSNR↑	HDR SSIM↑	LPIPS↓	PSNR↑	LDR SSIM↑	LPIPS↓	Training time (min)	Inference speed (fps)
1	HexPlane	LDR	-	-	-	13.82	0.551	0.576	25.50	0.44
2	HexPlane*	HDR	32.76	0.893	0.242	-	-	-	28.24	0.49
3	4DGS	LDR	-	-	-	7.99	0.072	0.620	42.50	290.57
4	4DGS*	HDR	7.85	0.220	0.534	-	-	-	40.75	307.85
5	HDR-NeRF	LDR	14.60	0.711	0.411	8.326	0.029	0.943	212.25	0.17
6	HDR-GS	LDR	13.27	0.783	0.261	20.52	0.840	0.148	38.25	73.30
7	HDR-GS†	LDR+HDR	29.40	0.936	0.097	20.85	0.834	0.182	56.50	64.88
8	HDR-HexPlane	LDR	9.306	0.672	0.353	27.44	0.748	0.353	36.98	0.35
9	HDR-HexPlane†	LDR+HDR	33.03	0.904	0.192	28.12	0.767	0.307	44.43	0.24
10	HDR-4DGS (Ours)	LDR	14.50	0.884	0.200	26.88	0.825	0.221	98.75	35.27
11	HDR-4DGS† (Ours)	LDR+HDR	25.13	0.909	0.162	30.69	0.927	0.097	76.86	48.63



Figure 3: Visual comparison of HDR DNVS on HDR-4D-Real.

➤ Experiments

PSNR prefers over-smooth images:




Figure 8: Additional visual comparison of HDR DNVs on HDR-4D-Syn.




Figure 9: Additional visual comparison of HDR DNVs on HDR-4D-Real.