

CYLINDERSPLAT: 3D GAUSSIAN SPLATTING WITH CYLINDRICAL TRIPLANES FOR PANORAMIC NOVEL VIEW SYNTHESIS

Qiwei Wang, Xianghui Ze, Jingyi Yu, Yujiao Shi


{shiyj2}@shanghaiitech.edu.cn



Paper

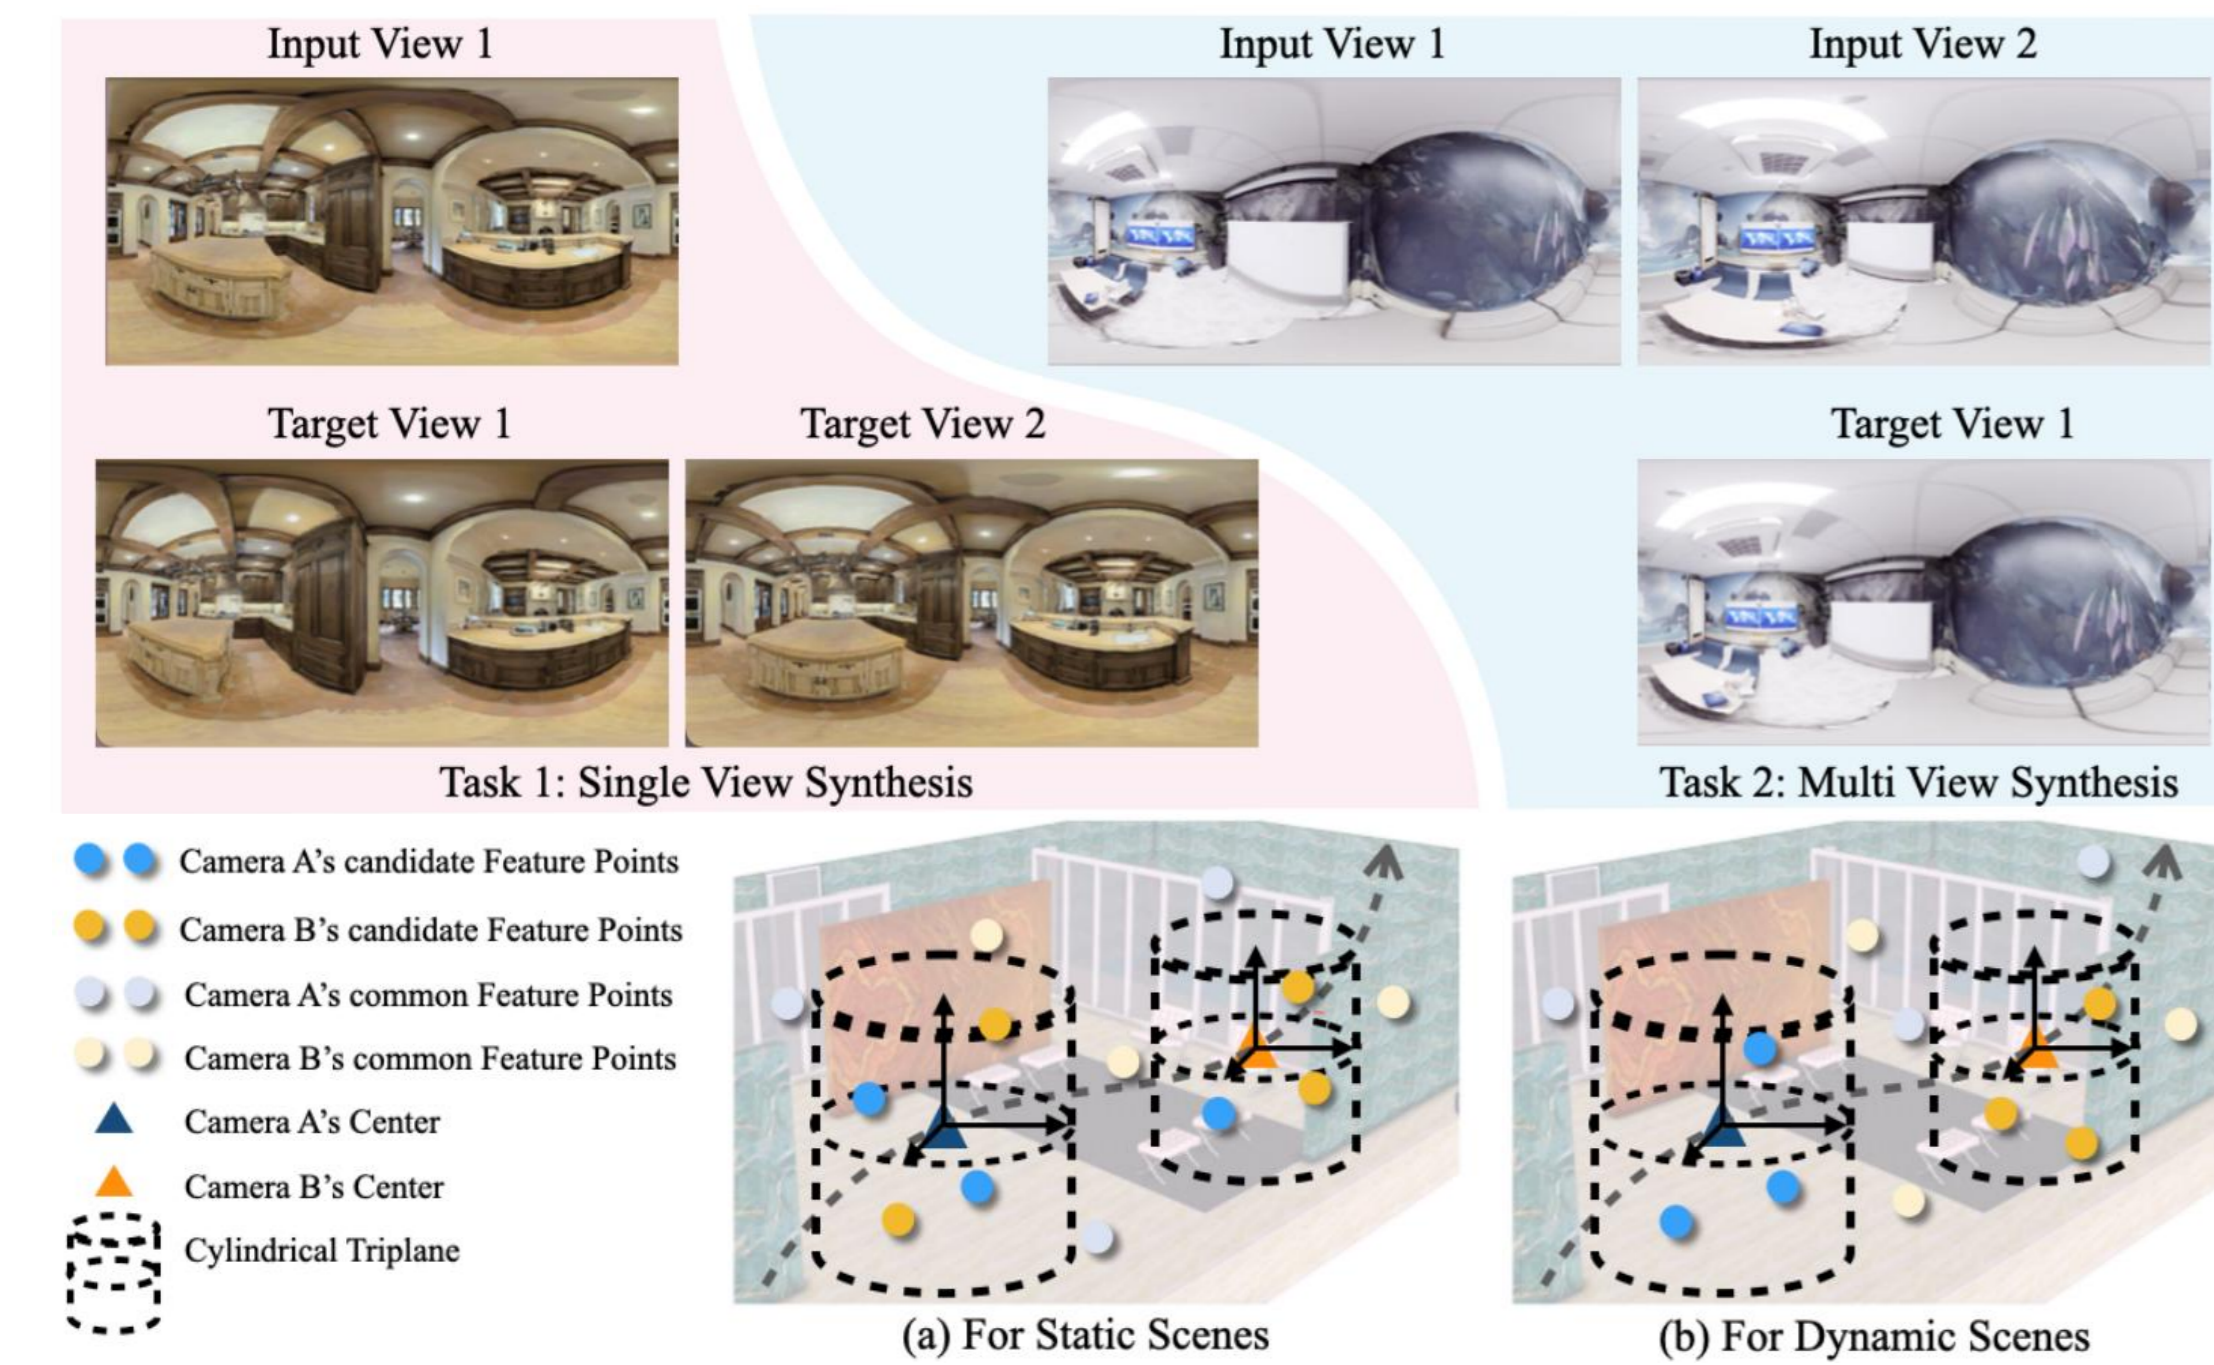


Code



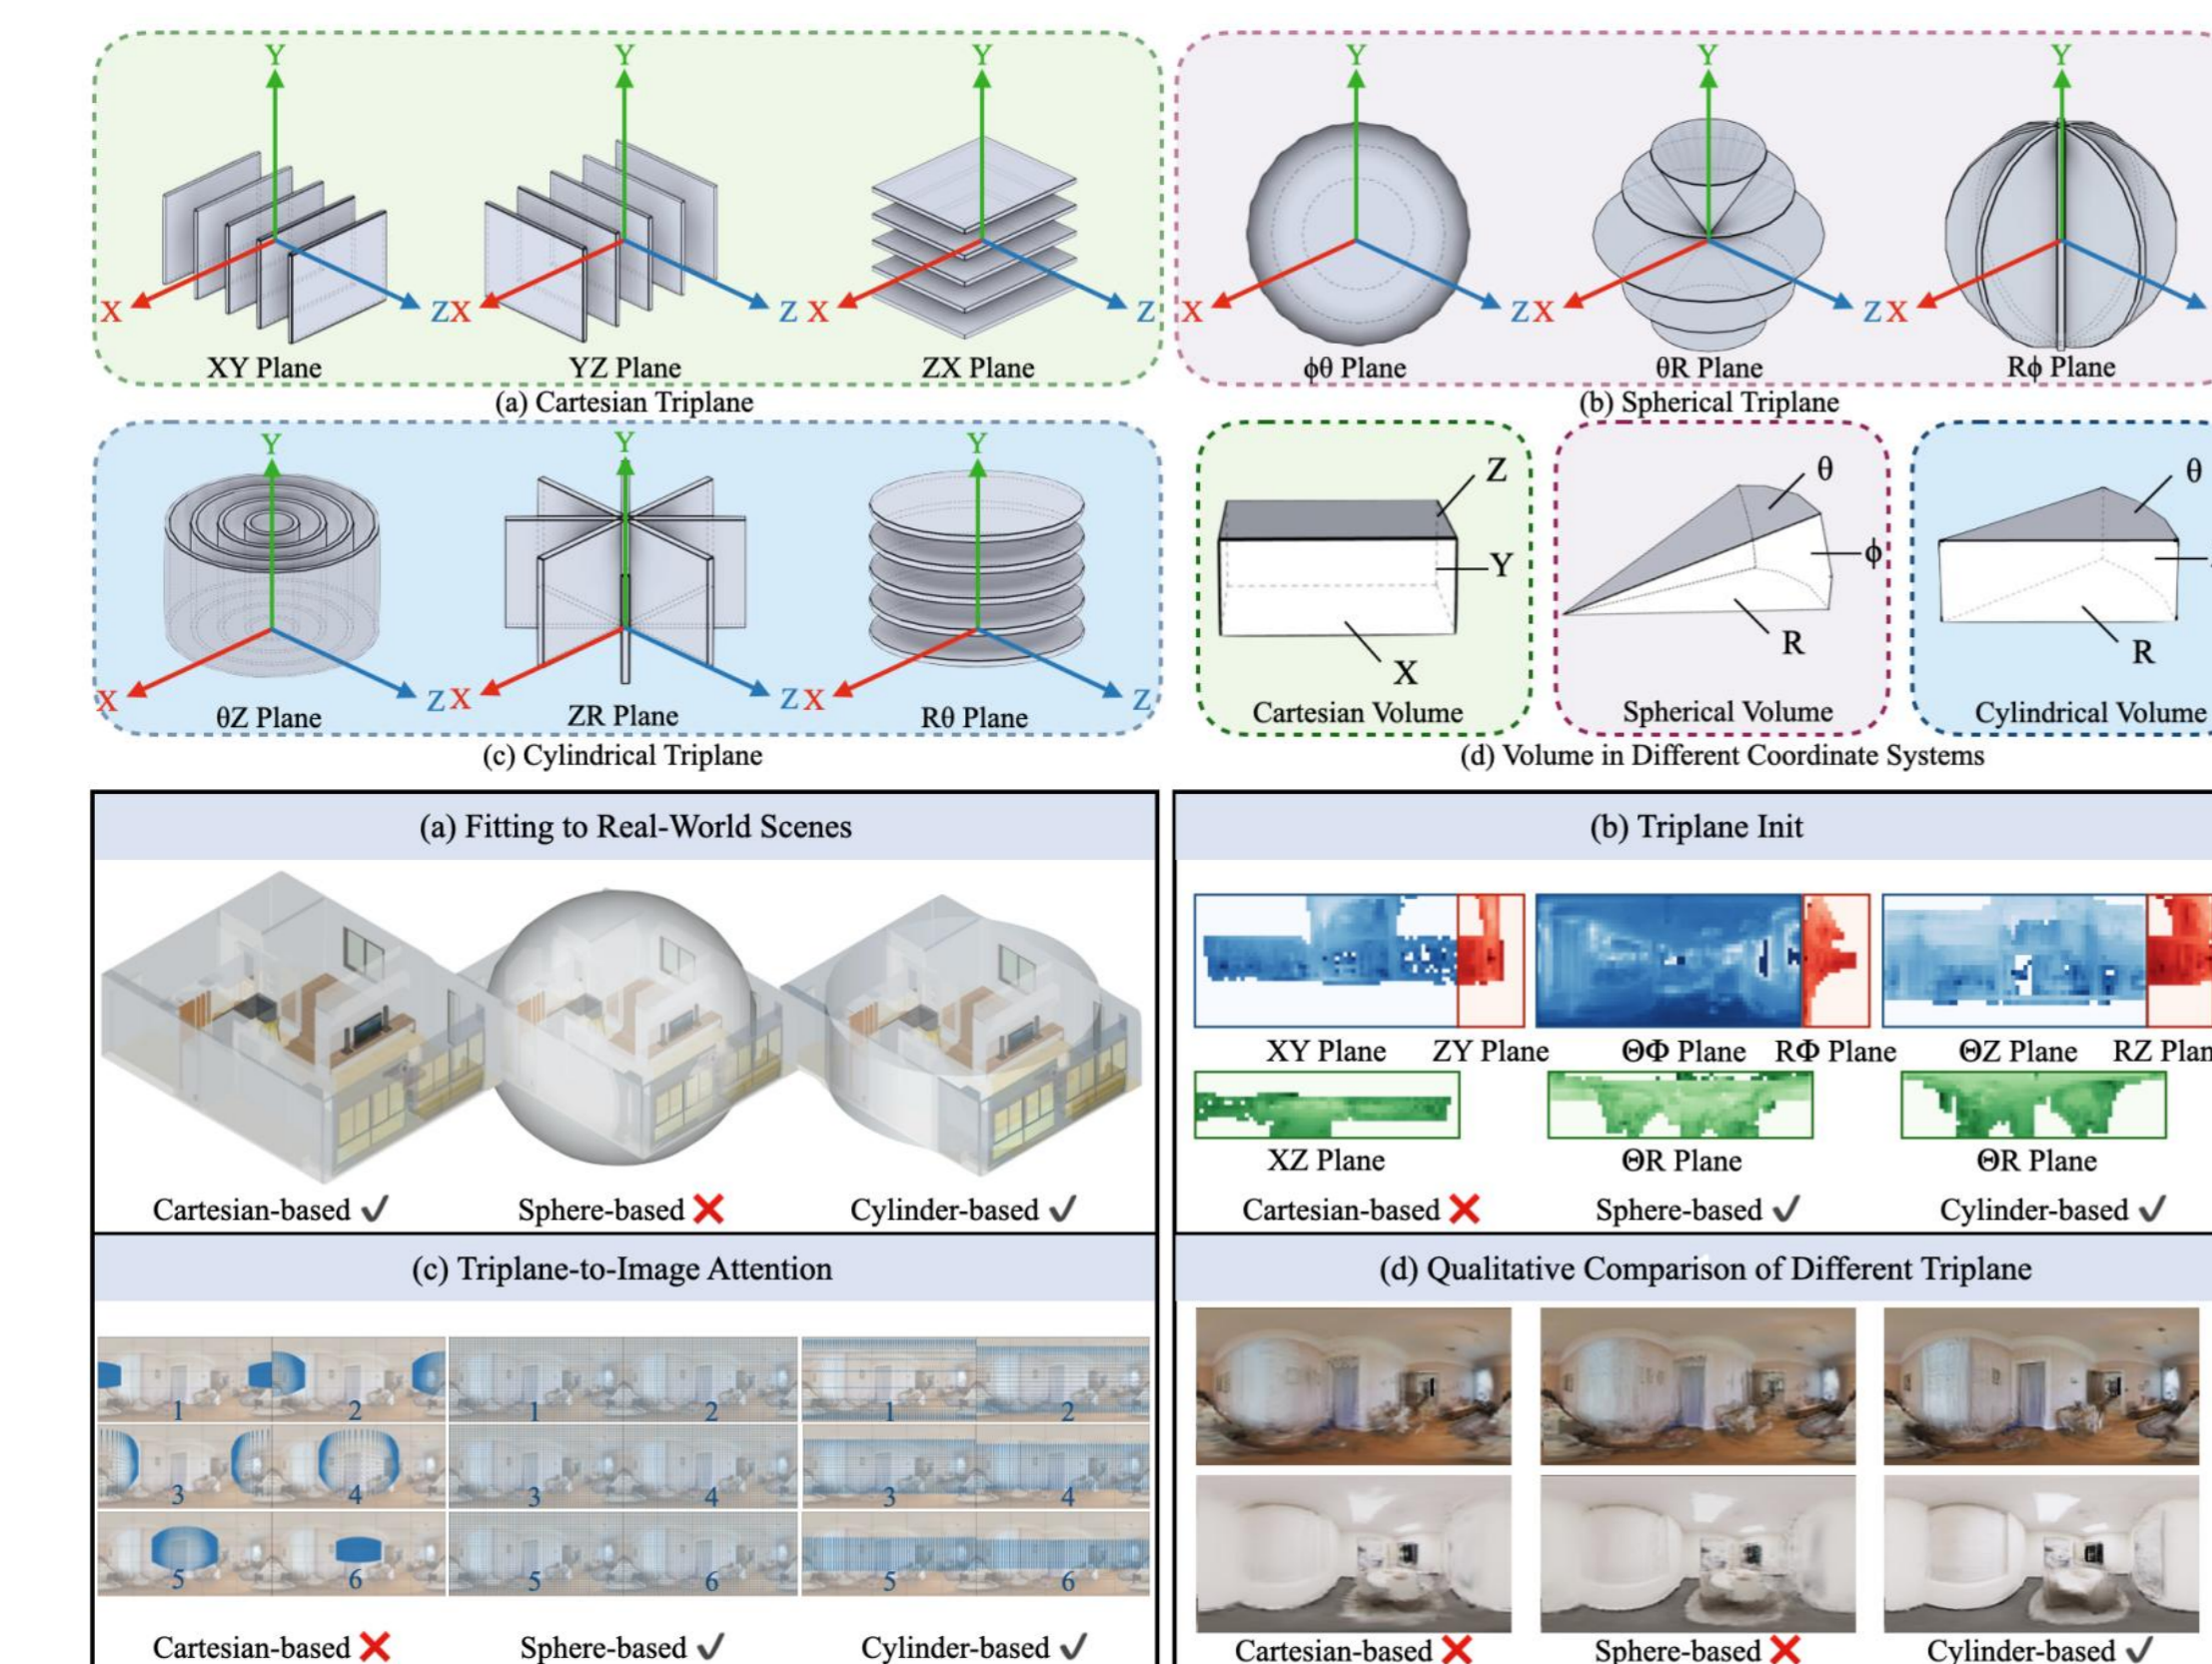
ICLR

Introduction



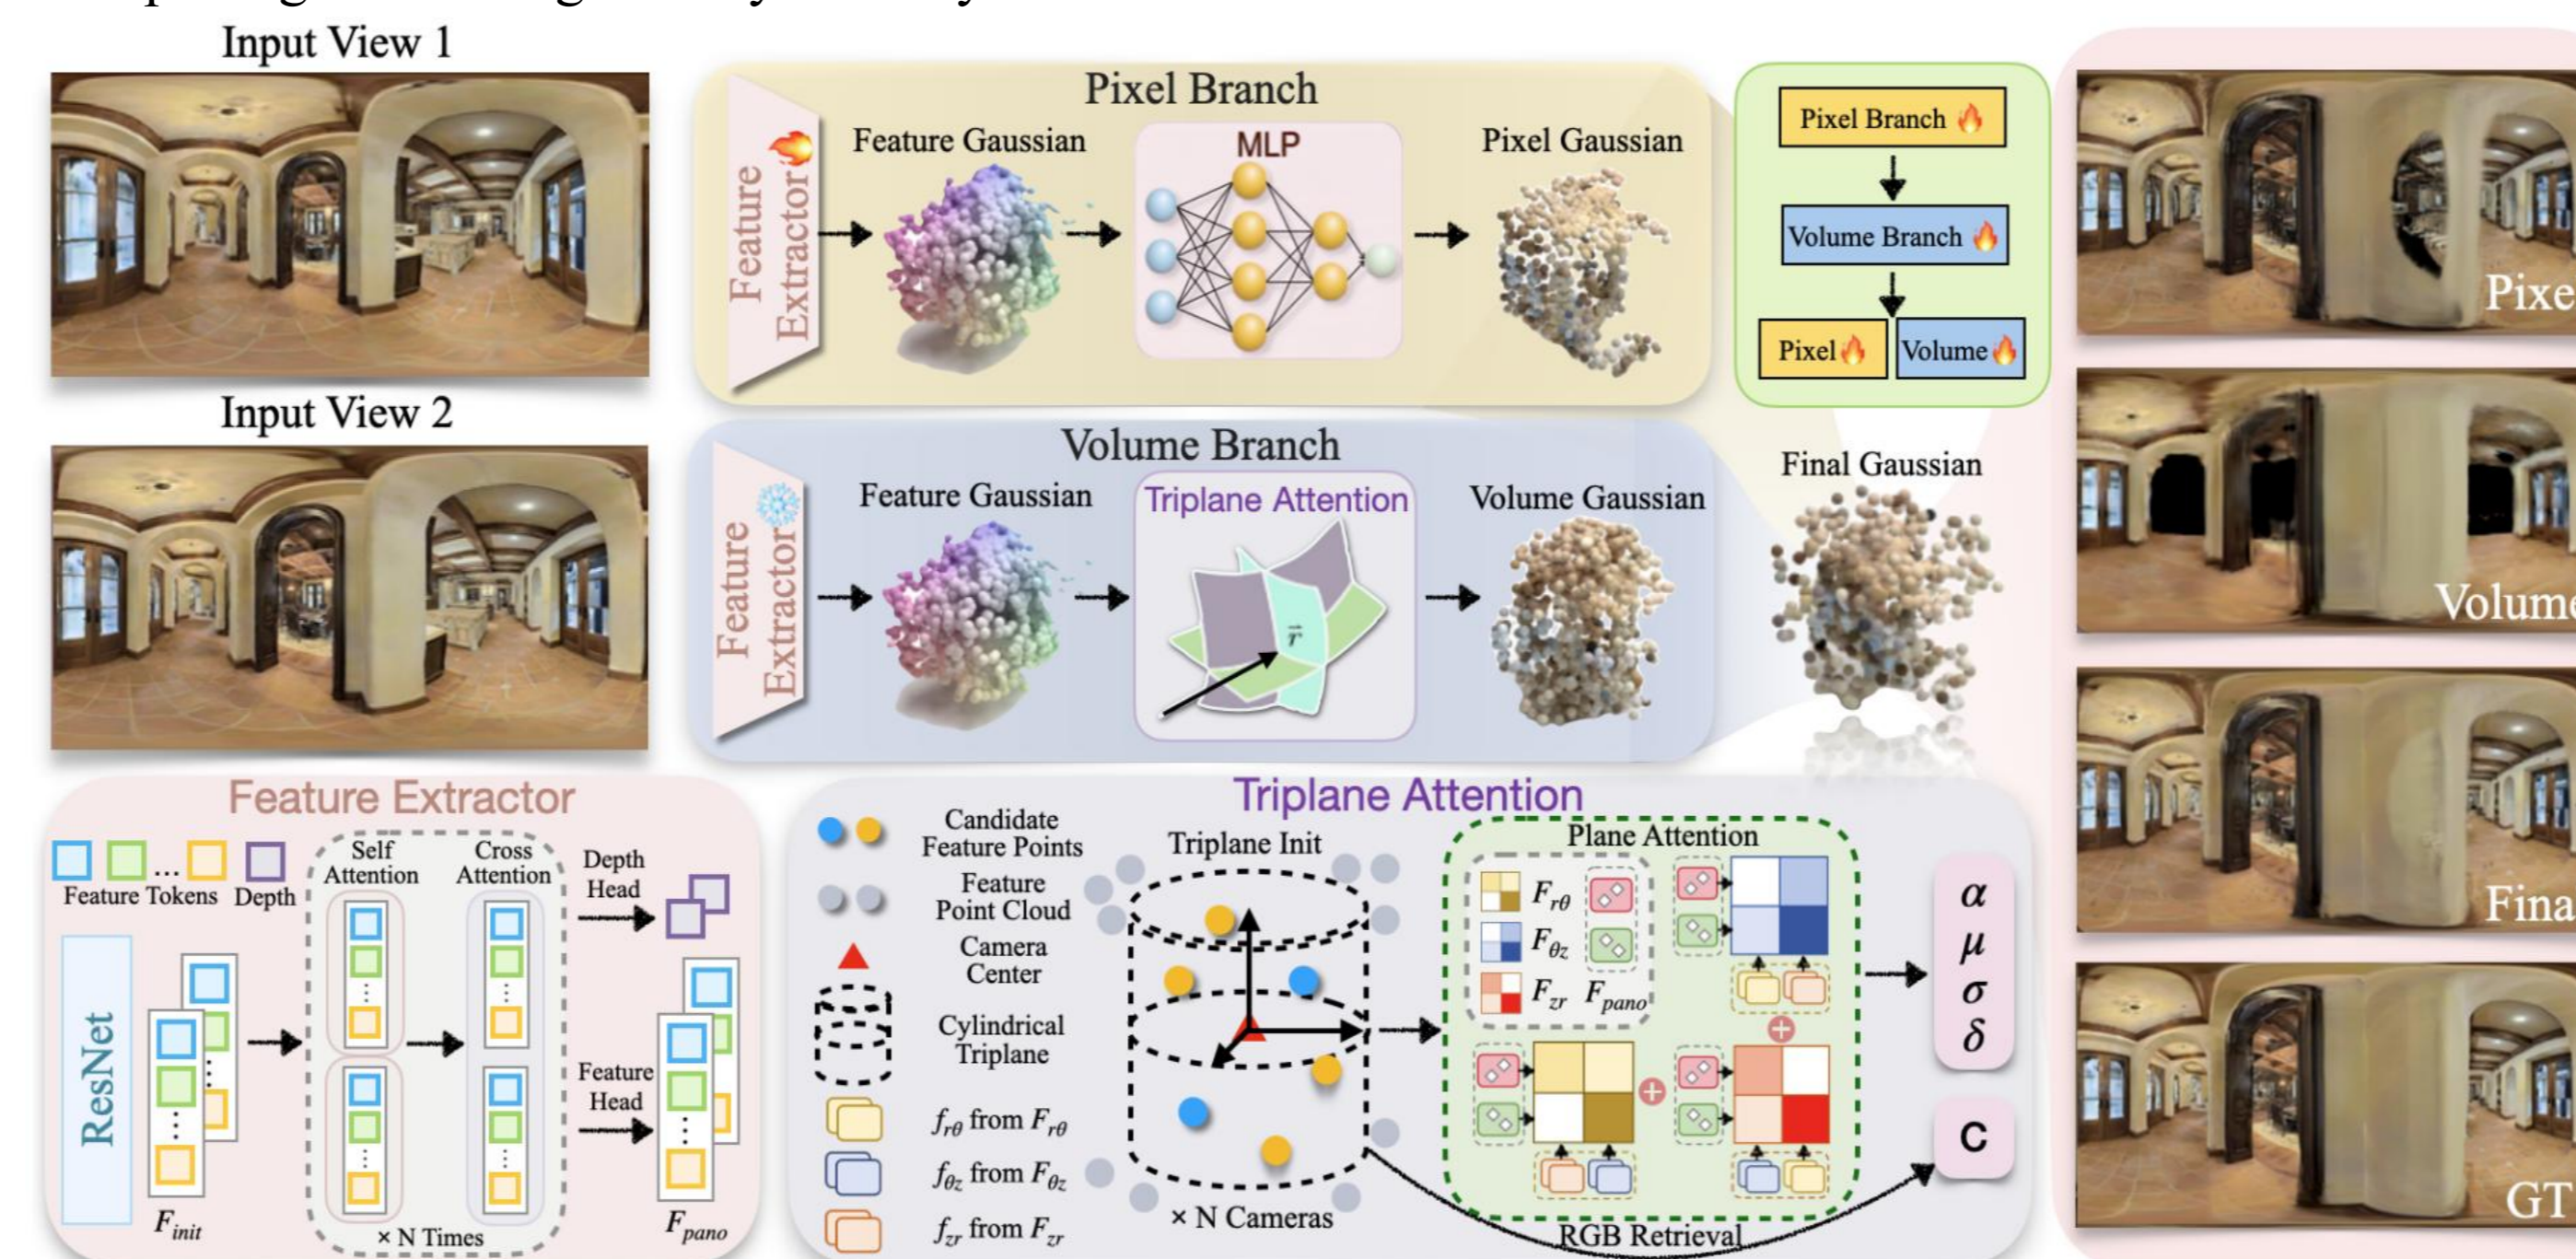
Applying feed-forward 3D Gaussian Splatting to panoramas encounters issues with geometric distortion and occlusion artifacts. We present CylinderSplat, which uses a new Cylindrical Triplane aligned with 360° geometry. Combined with a dual-branch architecture (pixel- and volume-based) to address sparse views and occlusions, it enables panoramic novel view synthesis in both single- and multi-view scenarios.

Motivation



Method

Our method uses a dual-branch architecture trained via a three-stage curriculum. The pixel branch uses a multi-view attention mechanism to generate high-quality Gaussians for well-observed regions. The volume branch is designed to fill the gaps by lifting features into our cylindrical triplane representation, thereby completing the scene geometry robustly.



Results

Table 1: Quantitative comparison for the two-view reconstruction on the Matterport3D, Replica, and Residential datasets. The first, second, and third best results are highlighted. Methods marked with * indicate that we reimplemented them using their official code.

Dataset	Matterport3D				Replica				Residential			
	2.0m		1.5m		1.0m		1.0m		about 0.3m		about 0.3m	
Method	PCC↑	WS-PSNR↑	SSIM↑	LPIPS↓	PCC↑	WS-PSNR↑	SSIM↑	LPIPS↓	PCC↑	WS-PSNR↑	SSIM↑	LPIPS↓
MVSplat	—	13.31	0.595	0.554	—	21.82	0.807	0.230	—	28.19	0.912	0.105
PanSplat*	—	20.96	0.701	0.352	—	23.38	0.811	0.282	—	27.12	0.876	0.195
OmniScene*	0.732	22.75	0.707	0.241	0.788	23.73	0.749	0.204	0.846	25.43	0.797	0.151
Splatter360*	0.684	21.31	0.741	0.285	0.791	24.68	0.801	0.199	0.887	26.71	0.828	0.141
PanSplat	0.716	20.56	0.777	0.265	0.779	24.09	0.849	0.181	0.819	28.83	0.935	0.091
Ours	0.851	23.76	0.835	0.175	0.867	25.91	0.873	0.128	0.923	28.89	0.937	0.081

Table 2: Quantitative comparison for the single-view reconstruction task.

Dataset	Matterport3D				Replica				Residential			
	2.0m		1.5m		1.0m		1.0m		about 0.3m		about 0.3m	
Method	PCC↑	WS-PSNR↑	SSIM↑	LPIPS↓	PCC↑	WS-PSNR↑	SSIM↑	LPIPS↓	PCC↑	WS-PSNR↑	SSIM↑	LPIPS↓
OmniScene*	0.681	22.52	0.707	0.244	0.757	23.50	0.749	0.206	0.816	25.22	0.798	0.152
Splatter360*	0.628	21.81	0.650	0.380	0.648	22.40	0.690	0.350	0.653	24.15	0.738	0.308
PanSplat*	0.596	19.72	0.625	0.334	0.578	20.84	0.668	0.292	0.621	22.43	0.717	0.223
Ours	0.821	23.75	0.822	0.175	0.856	25.13	0.854	0.136	0.903	27.01	0.915	0.089

Table 3: Quantitative comparison for the two-view reconstruction task on the 360Loc dataset.

Dataset	360Loc (avg. 1.40m baseline)			
	PCC↑	WS-PSNR↑	SSIM↑	LPIPS↓
MVSplat	—	24.67	0.823	0.170
OmniScene*	0.830	27.41	0.831	0.153
Splatter360*	0.866	28.14	0.860	0.127
PanSplat	0.188	28.24	0.873	0.108
Ours	0.884	28.35	0.896	0.095

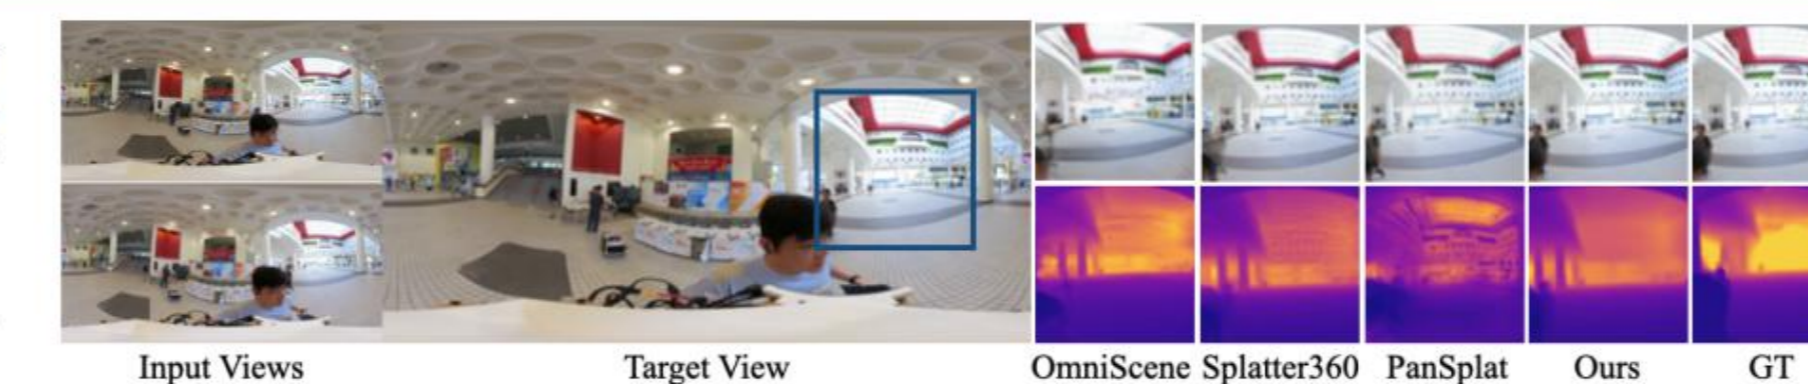


Figure 4: Qualitative comparison of 360Loc (two-view). Left: input/target views. Right: zoomed-in novel views and depth maps (warm = far, cool = near). Our ground truth (GT) depth is obtained from DepthAnywhere (Wang & Liu, 2024), which serves as a reference for calculating PCC.

Qualitative comparison

