Lid-Lab-NeRF: Generating Temporally Consistent, Labelled LiDAR Point Clouds using Neural Radiance Fields

Shrestha Srivastava[1], Vaibhav Kumar[1,2]

[1]GeoAI4Cities Lab, [2]SimDAAS Autonomy Pvt. Ltd.

Lid-Lab-NeRF

- Generating Novel, Labeled LiDAR Scans from Unseen Viewpoints
- Uses a NeRF architecture to learn an implicit 3D representation from existing LiDAR scans.
- Outputs completely novel and fully labelled scans from any new viewpoint.

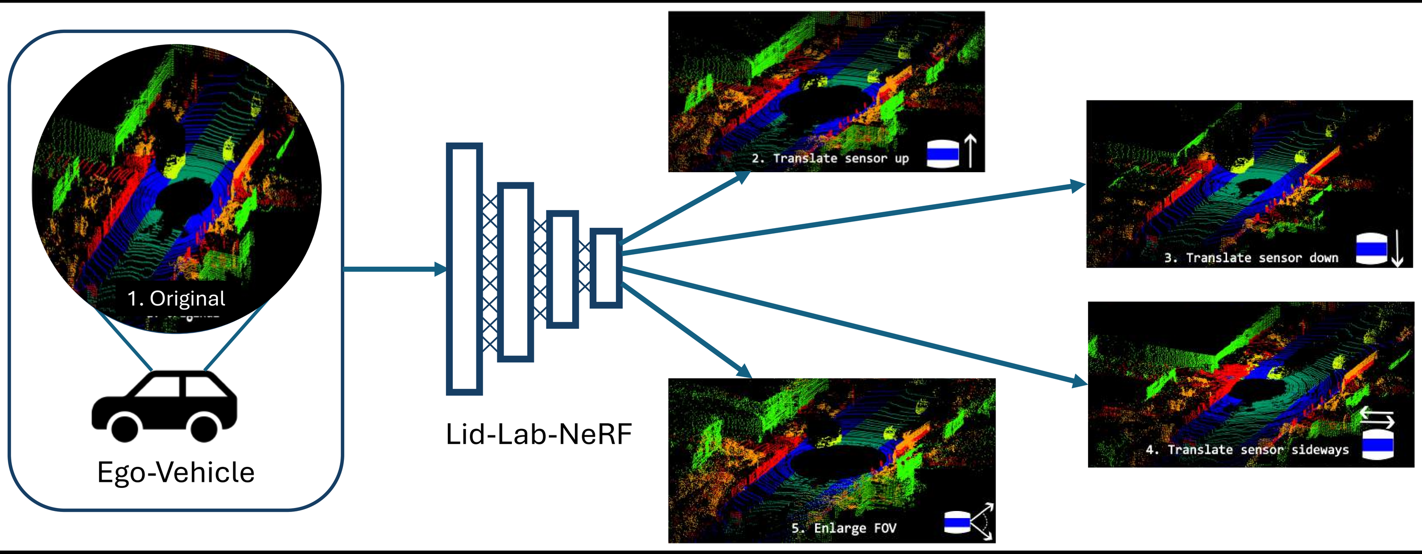

1 Methodology

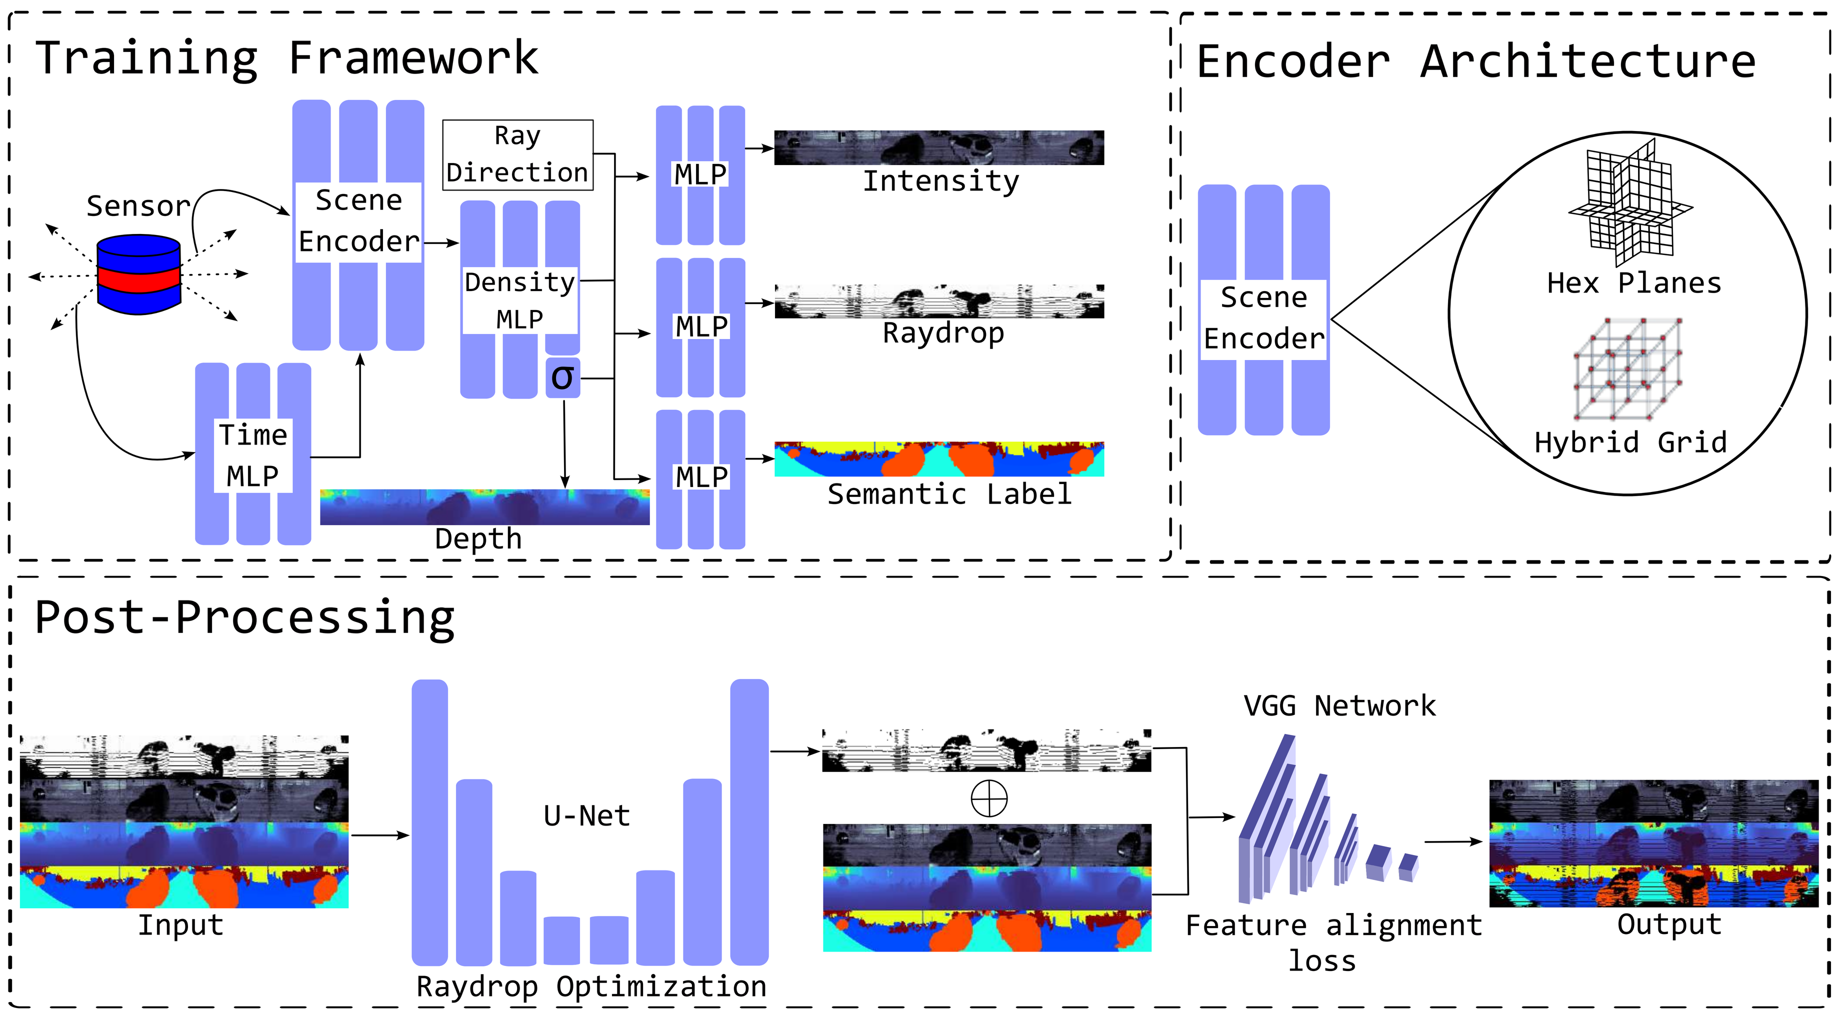

Architecture Components

- K-Plane & Grid Encoder: Handles large-scale, sparse scenes.
- Time MLP: Models moving objects.
- Prediction MLPs: Output Depth, Intensity, Label, & Ray-Drop.
- U-Net (VGG Loss): Ensures realistic scan patterns.

2 Qualitative Results

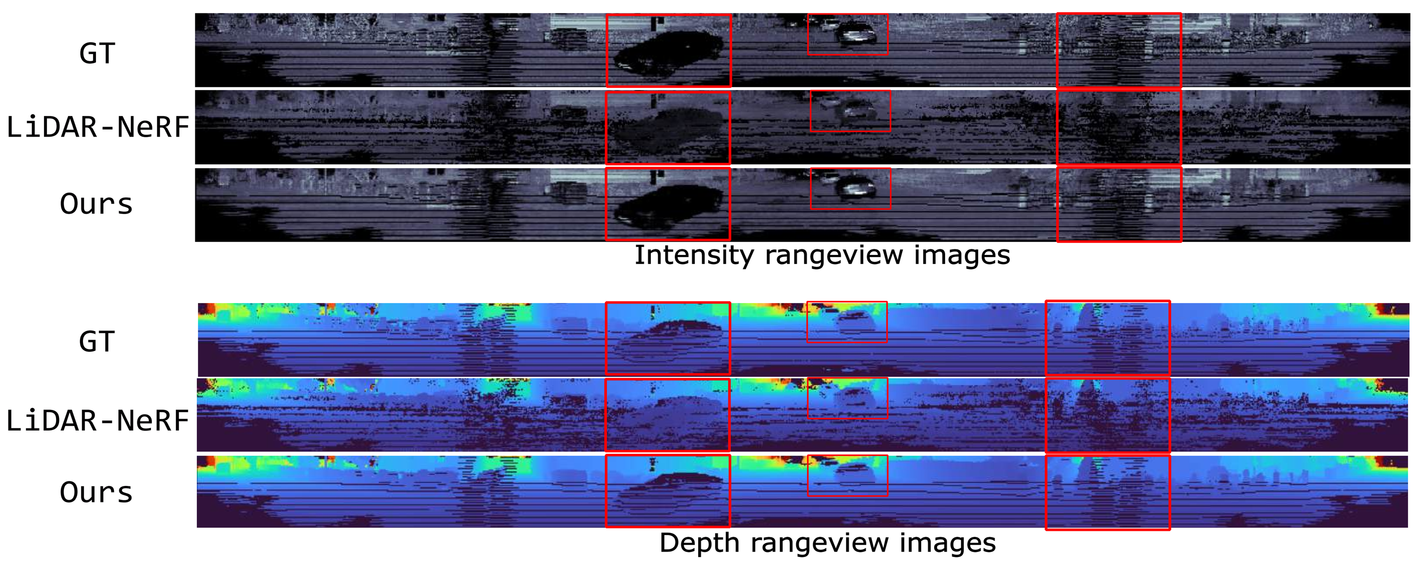

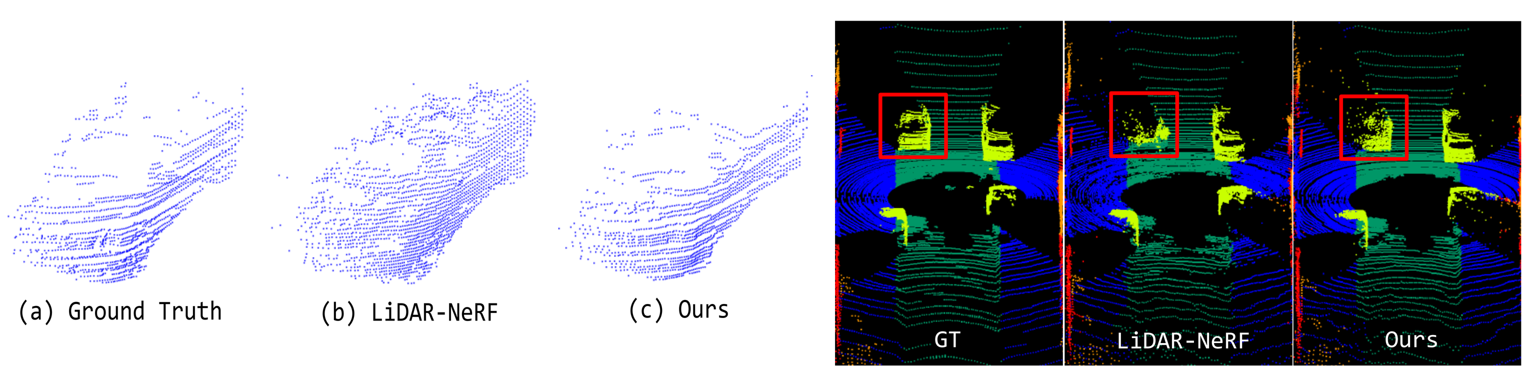

3 Quantitative Results

Table 1: Quantitative comparison on KITTI-360 dataset. We compare our method to different previous approaches and highlight the best results in **bold** and the second-best results in *italic*.

Method	Point Cloud		Depth					Intensity				
	CD↓	F-score↑	RMSE↓	MedAE↓	LPIPS↓	SSIM↑	PSNR↑	RMSE↓	MedAE↓	LPIPS↓	SSIM↑	PSNR↑
LiDARsim	3.2228	0.7157	6.9153	0.1279	0.2926	0.6342	21.4608	0.1666	0.0569	0.3276	0.3502	15.5853
NKSR	1.8982	0.6855	5.8403	0.0996	*0.2752*	0.6409	23.0368	0.1742	0.0590	0.3337	0.3517	15.2081
PCGen	0.4636	0.8023	5.6583	0.2040	0.5391	0.4903	23.1675	0.1970	0.0763	0.5926	0.1351	14.1181
D-NeRF	0.1442	*0.9128*	*4.0194*	0.0508	0.3061	*0.6634*	*26.2344*	0.1369	*0.0440*	0.3409	0.3748	*17.3554*
TiNeuVox-B	0.1748	0.9059	4.1284	*0.0502*	0.3427	0.6514	26.0267	*0.1363*	0.0453	0.4365	0.3457	17.3535
K-Planes	*0.1302*	0.9123	4.1322	0.0539	0.3457	0.6385	26.0236	0.1415	0.0498	0.4081	0.3008	17.0167
LiDAR-NeRF	0.1438	0.9091	4.1753	0.0566	0.2797	0.6568	25.9878	0.1404	0.0443	*0.3135*	*0.3831*	17.1549
Lid-Lab-Nerf (ours)	**0.1183**	**0.9154**	**2.9504**	**0.0317**	**0.0768**	**0.8629**	**28.6802**	**0.1056**	**0.0263**	**0.1466**	**0.6217**	**19.5330**

Table 2: Segmentation performance of MinkUNet models trained on real KITTI-360 data and on our synthetic dataset, evaluated on real KITTI-360 scans. The model trained on real data achieves a slightly higher mean IoU overall, while the synthetic-data model remains competitive and outperforms in certain classes.

Model	Sidewalk IoU	Road IoU	Building IoU	Vehicle IoU	Vegetation IoU	Others IoU	Mean IoU
MinkUNet (trained on real KITTI-360)	35.45	36.55	33.89	40.25	14.92	3.84	27.49
MinkUNet (trained on our synthetic data)	34.85	33.67	30.41	40.39	16.70	5.93	27.00

4 Other Applications

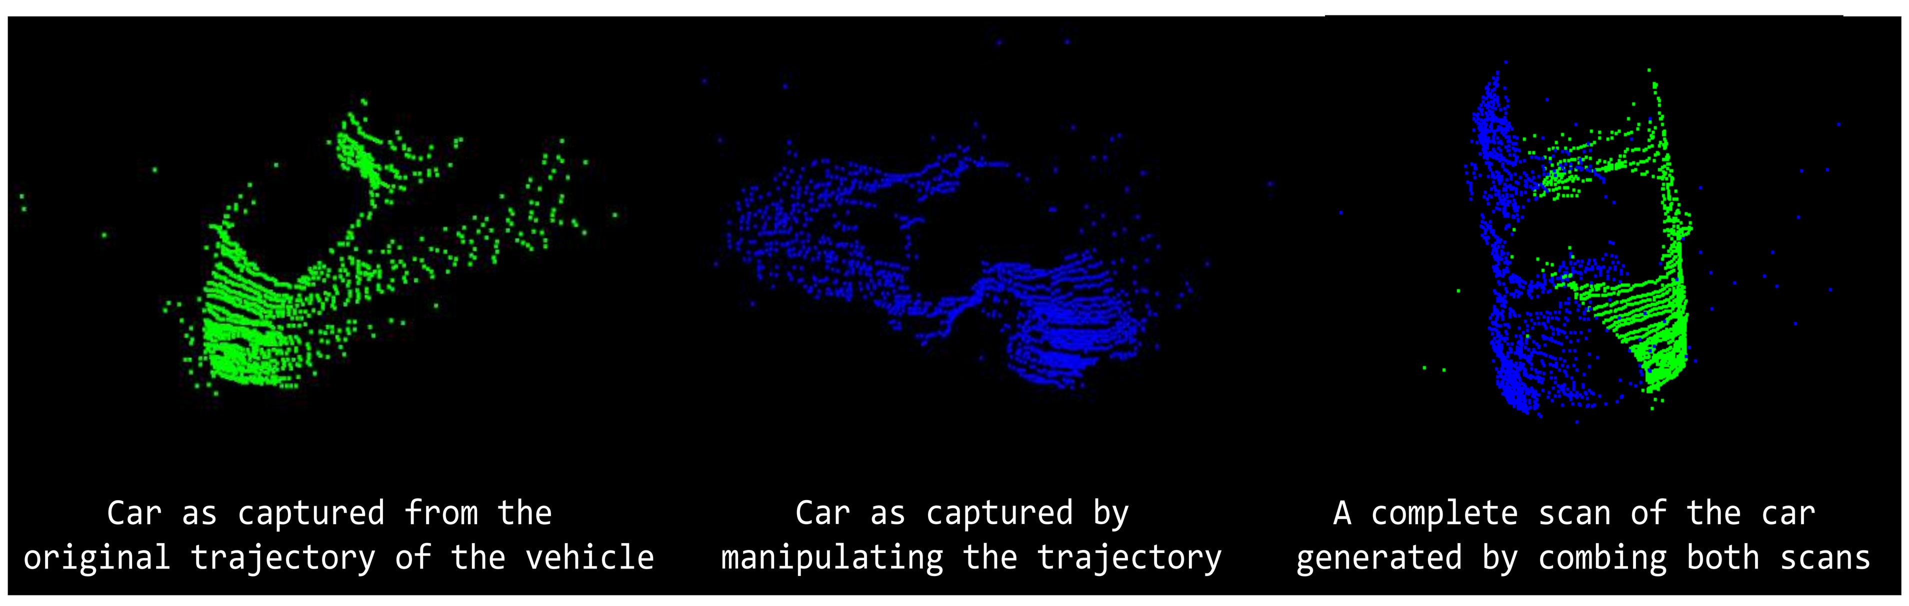

Object/Scene Completion

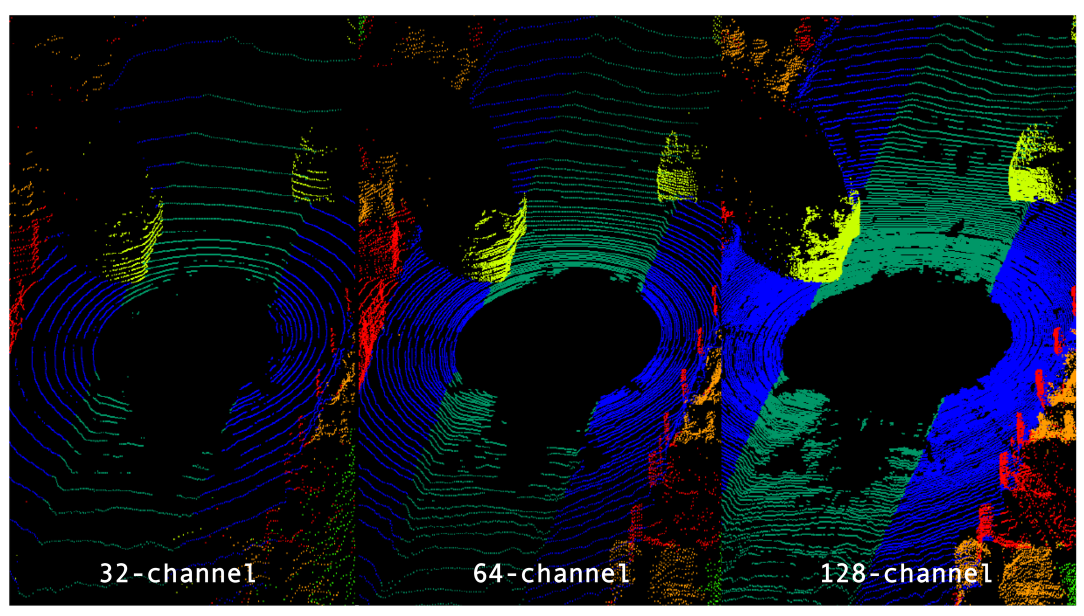

Increase or decrease the number of channels during simulations