MAC-NeRF: Motion-Aware Curriculum Learning for Dynamic LiDAR NeRFs

Shangshu Yu[1] Xiaotian Sun[2] Wen Li[3†] Rui She[4]
Hanyun Wang[5] Sheng Ao[2] Chenglu Wen[2] Cheng Wang[2†]

[1]Northeastern University [2]Xiamen University [3]University of Bristol
[4]Beihang University [5]Sun Yat-sen University

ICML
International Conference
On Machine Learning

Challenges & Motivation

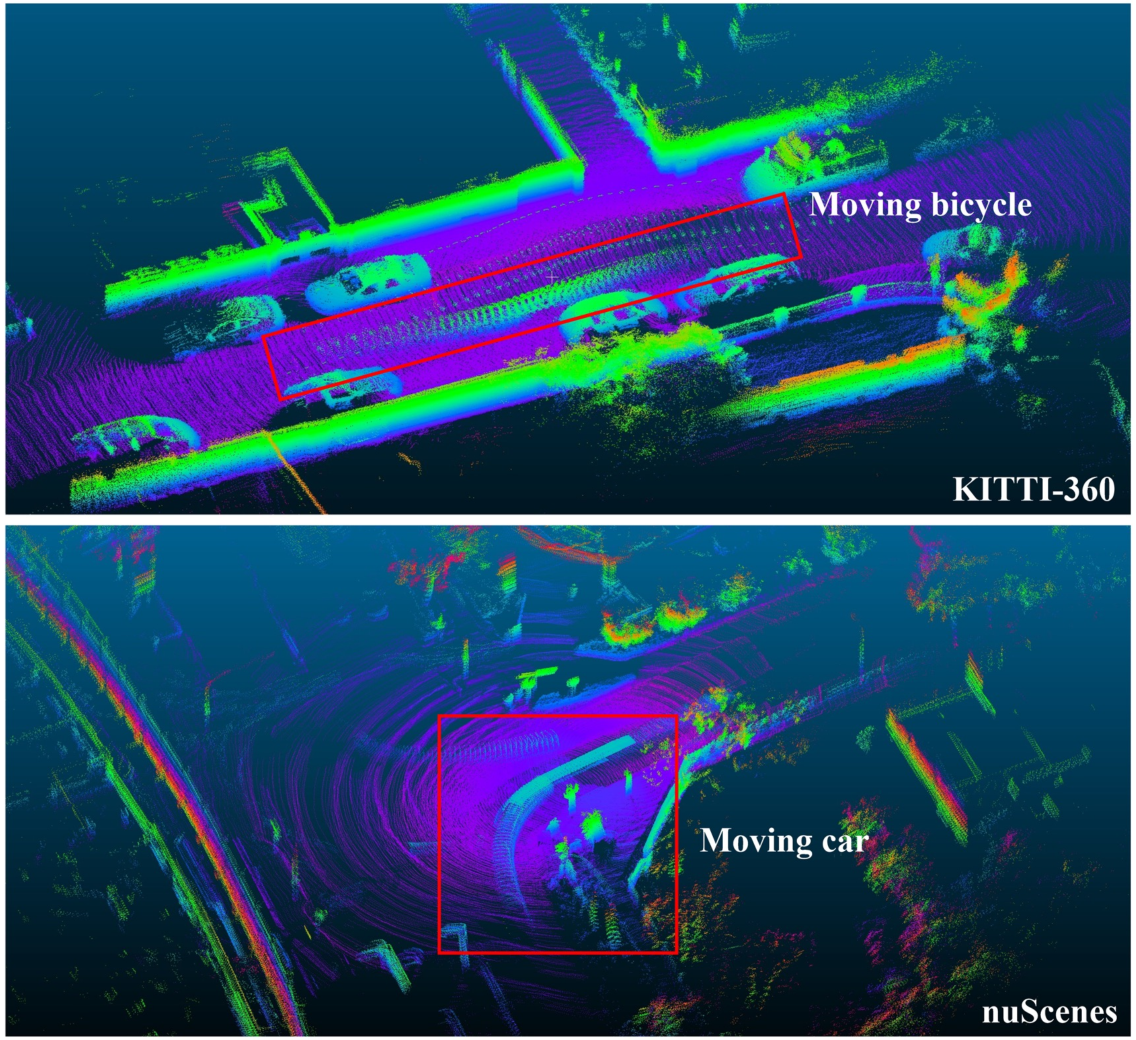

- Despite the potential of LiDAR NeRFs, learning accurate dynamic representations remains a critical challenge, as moving objects introduce **motion-induced supervision conflicts** and **geometric ambiguities**.

- Specifically, moving objects violate multi-view consistency, generating conflicting supervision driving the former, while inducing ambiguities between genuine structures and ghosting artifacts resulting in the latter.

Contribution

- We propose MAC-NeRF, a novel LiDAR NeRF framework that introduces motion-aware curriculum learning to address motion-induced supervision conflicts and geometric ambiguities, enabling high-fidelity NVS.
- We propose a Rectified Temporal Consistency module, resolving motion-induced supervision conflicts by filtering out erroneous motion and prioritizing trustworthy temporal correspondence learning.
- We devise the Confidence-Modulated Frequency Regularization mechanism, which adaptively anneals the regularization bandwidth to progressively suppress ghosting artifacts while preserving structure details.

Framework

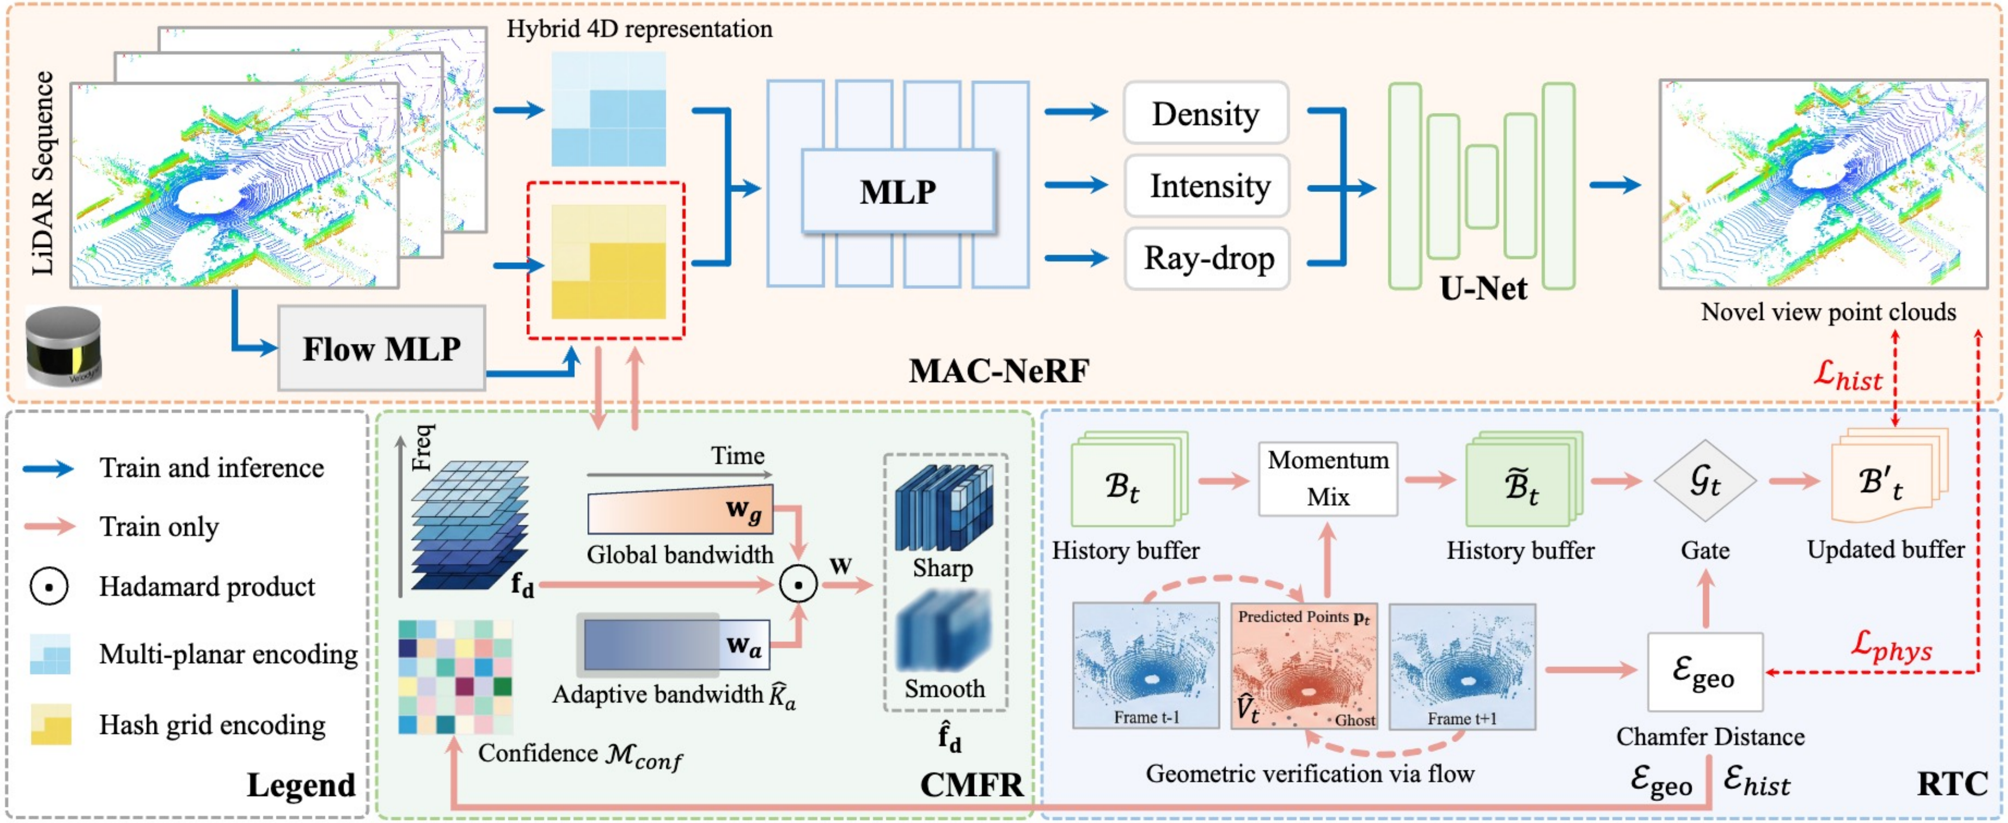

Figure 2. Overview of MAC-NeRF. The pipeline consists of a hybrid 4D representation with a scene flow MLP to render density, intensity, and ray-drop for LiDAR NVS. The proposed **RTC** module (blue box) maintains a history buffer $\mathcal{B}_t$ as stable supervision to first anchor predictions, while progressively employing geometric verification (via $\mathcal{E}_{geo}$) to filter out erroneous motion. The proposed **CMFR** strategy (green box) utilizes the geometric confidence $\mathcal{M}_{conf}$ derived from RTC to adaptively modulate the frequency regularization bandwidth, suppressing ghosting artifacts before recovering fine details. RTC and CMFR are only employed during training.

Experiment

Table 1. Quantitative results on the **KITTI-360** dataset. We group baselines into Static and Dynamic NVS categories. The best result is in **bold**, and the second best is underlined.

	Method	Point Cloud		Depth					Intensity				
		CD↓	F-score↑	RMSE↓	MedAE↓	LPIPS↓	SSIM↑	PSNR↑	RMSE↓	MedAE↓	LPIPS↓	SSIM↑	PSNR↑
Static	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
	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
	LiDAR-GS	0.1288	0.9156	3.9333	0.0530	0.2475	0.6276	26.1326	0.1336	0.0453	0.3411	0.3737	17.4554
Dynamic	LiDAR4D	0.1089	0.9272	3.5256	0.0404	0.1051	0.7647	27.4767	0.1195	0.0327	0.1845	0.5304	18.5561
	LiDAR-RT	0.1077	0.9255	3.4671	0.0512	0.1016	0.8406	27.6755	0.1115	0.0271	0.1812	0.6077	19.0862
	STGC-NeRF	0.0997	0.9325	3.0794	0.0277	0.0681	0.8774	28.6796	0.0995	0.0262	0.1479	0.6563	20.0825
	MAC-NeRF	**0.0961**	**0.9367**	**3.0262**	**0.0272**	**0.0675**	**0.8806**	**28.7475**	**0.0984**	**0.0258**	**0.1489**	**0.6623**	**20.1862**

Table 2. Quantitative results on the **nuScenes** dataset. The notations are consistent with Tab. 1.

	Method	Point Cloud		Depth					Intensity				
		CD↓	F-score↑	RMSE↓	MedAE↓	LPIPS↓	SSIM↑	PSNR↑	RMSE↓	MedAE↓	LPIPS↓	SSIM↑	PSNR↑
Static	LiDARsim	12.1383	0.6512	10.5539	0.3572	0.1871	0.5653	17.7841	0.0659	0.0115	0.1160	0.5170	23.7791
	NKSR	11.4910	0.6178	9.3731	0.5763	0.2111	0.5637	18.7774	0.0680	0.0119	0.1290	0.5031	23.4905
	PCGen	2.1998	0.6341	8.8364	0.4011	0.1792	0.5440	19.2799	0.0768	0.0147	0.1308	0.4410	22.4428
	LiDAR-NeRF	0.3225	0.8576	7.1566	0.0338	0.0702	0.7188	21.2129	0.0467	0.0076	0.0483	0.7264	26.9927
	LiDAR-GS	0.2829	0.8987	6.8806	0.0290	0.1108	0.7286	21.7803	0.0444	0.0077	0.0527	0.7346	27.2427
Dynamic	LiDAR4D	0.2443	0.8915	6.7831	0.0258	0.0569	0.7396	21.7189	0.0426	0.0071	0.0459	0.7498	27.7977
	LiDAR-RT	0.2300	0.9000	6.6195	0.0245	0.0493	0.7655	21.8720	0.0422	0.0081	0.0483	0.7542	27.8105
	STGC-NeRF	0.2204	0.9070	6.5361	0.0240	0.0486	0.7741	22.0044	0.0417	0.0082	0.0418	0.7566	27.9989
	MAC-NeRF	**0.2111**	**0.9087**	**6.4559**	**0.0230**	**0.0484**	**0.7758**	**22.0918**	**0.0410**	**0.0061**	**0.0399**	**0.7597**	**28.0297**

Experiment

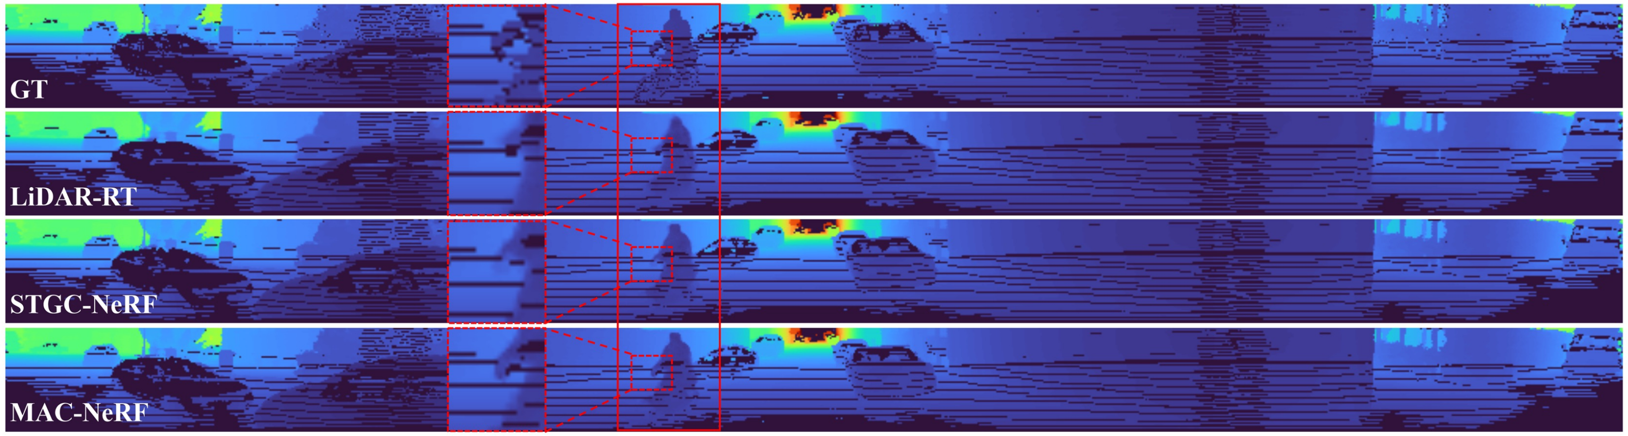

(a) Rendered depth map

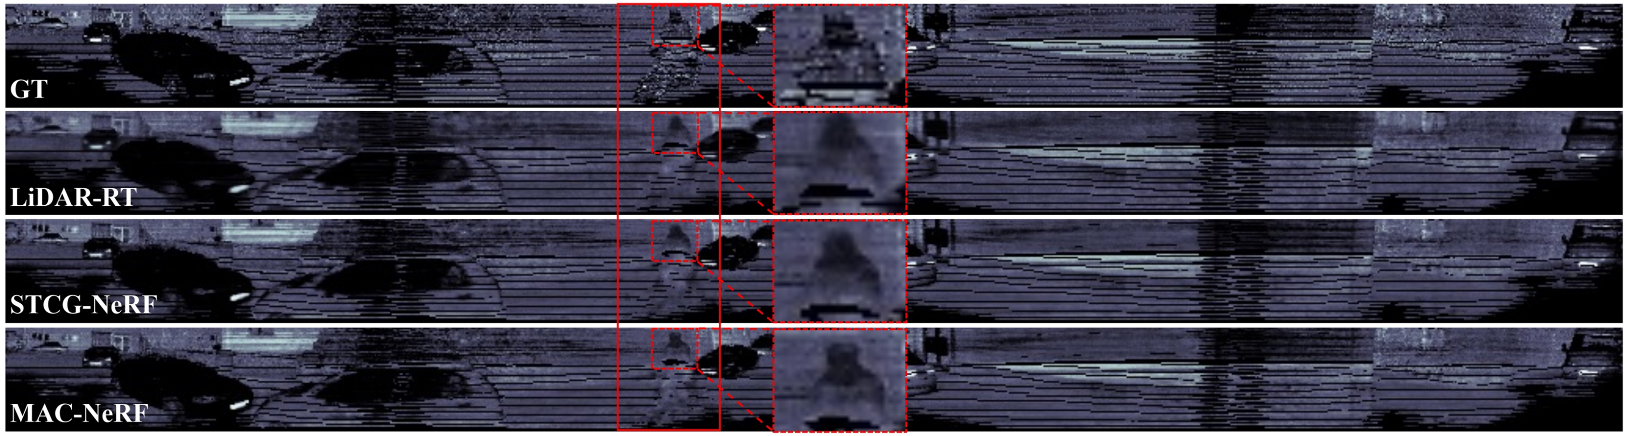

(b) Rendered intensity map

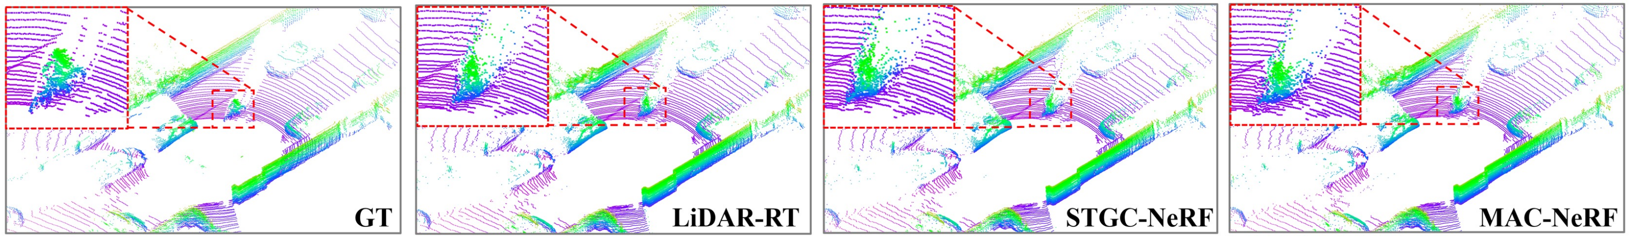

(c) Rendered LiDAR point clouds

T h a n k s !

- **Homepage：https://psyz1234.github.io/**