NeuRAD: Neural Rendering for Autonomous Driving

Adam Tonderski*, Carl Lindström*, Georg Hess*, William Ljungbergh, Lennart Svensson, Christoffer Petersson

zenseact · Lund University · Chalmers University of Technology · Linköping University · WASP Wallenberg AI, Autonomous Systems and Software Program

Contributions

- ★ **Full rendering capabilities:** Neural rendering for full 360° camera and lidar on dynamic AD data.
- ★ **AD data modeling:** Novel strategies to model camera and lidar data in unbounded AD scenes.
- ★ **State-of-the-art performance:** Improved metrics on five popular AD datasets.
- ★ **Open-source:** Built on top of Nerfstudio, code released at https://github.com/georghess/neurad-studio.

Motivation: Driving data is boring

Collected data is mostly **uneventful!**

Safe autonomous driving requires handling of **corner cases.**

Manual collections and game engine simulations are **expensive**, **time-consuming** and **scales poorly.**

NeuRAD: Neural Rendering for AD data

1. Collect sensor data
2. Learn scene representation
3. Render sensor data from new views

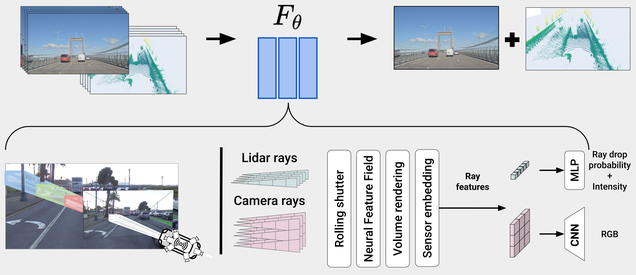

Method: Automotive data modeling

Dynamic and unbounded scenes:

- Our scene representation is **decomposed** into static and dynamic parts for increased controllability.
- **Scale-dependent downweighting** and **contraction** enables learning features at multiple scales, without anti-aliasing.

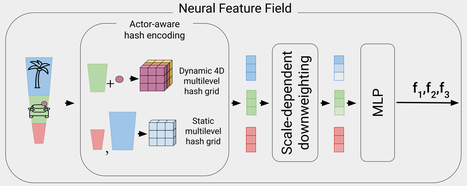

Rolling shutter: Modeling rolling shutter for both camera and lidar to account for a fast moving sensor rig.

Camera-specific embeddings: Accounting for varying camera settings.

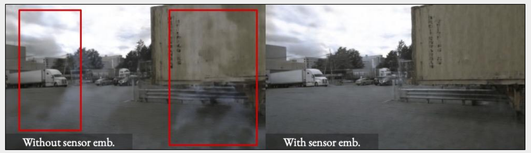

Learnable ray drop probability: Modeling the phenomenon that lidar rays can travel far without hitting a surface, or hit surfaces from which the beam bounces off.

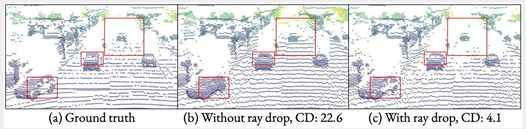

Proposal sampling vs. occupancy:

1. Sampling surfaces far from any lidar points.
2. Recovering thin structures or fine details of close-up surfaces.

Qualitative Results

Realistic rendering: High-quality renders of both camera and lidar data, across multiple datasets.

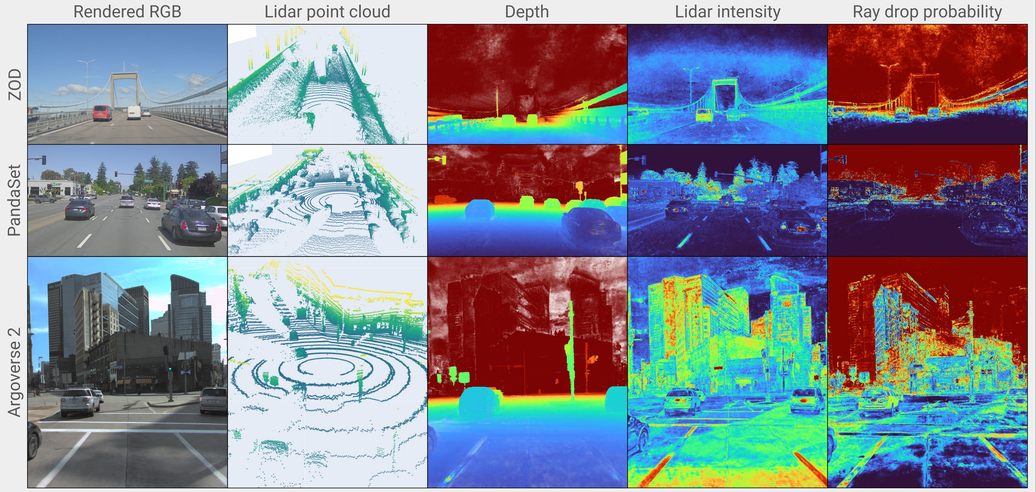

Scalable sensor-realistic simulations: The learned scene can be easily manipulated by controlling the self-driving vehicle, changing its sensor placement or moving other actors.

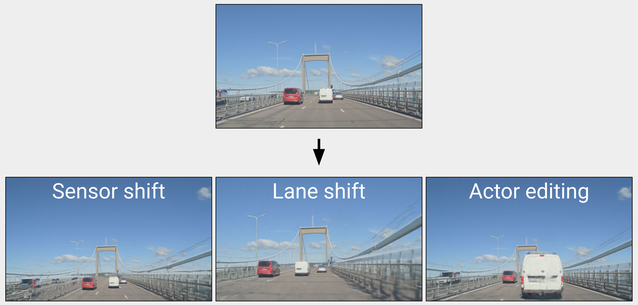

Quantitative Results

Dataset	Model	PSNR ↑	SSIM ↑	LPIPS ↓
Panda FC	UniSim	25.63	0.745	0.288
	NeuRAD	**26.58**	**0.778**	**0.190**
Panda 360°	UniSim*	23.50	0.692	0.330
	NeuRAD	**25.97**	**0.758**	**0.242**
nuScenes	S-NeRF	26.21	**0.831**	0.228
	NeuRAD	**26.99**	0.815	**0.225**
KITTI MOT	MARS	24.00	0.801	0.164
	NeuRAD	**27.00**	**0.795**	**0.082**
Argo2	UniSim*	23.22§	0.661§	0.412§
	NeuRAD	**26.22**	**0.717**	**0.315**
Zenseact Open Dataset	UniSim*	27.97	0.777	0.239
	NeuRAD	**29.49**	**0.809**	**0.226**

CVPR SEATTLE, WA JUNE 17-21, 2024