Neural Simulation for Autonomous Driving

Adam Tonderski | Zenseact + Lund University | C. Petersson + K. Åström

Paper I: NeuRAD - Neural Rendering for Autonomous Driving (CVPR 24)

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

Paper II: NeuroNCAP - Photorealistic Closed-Loop Safety Testing for Autonomous Driving (ECCV 24)

William Ljungbergh*, Adam Tonderski*, Joakim Johnander, Holger Caesar, Kalle Åström, Mikael Felsberg, Christoffer Petersson

NeuRAD: Neural Rendering for AD

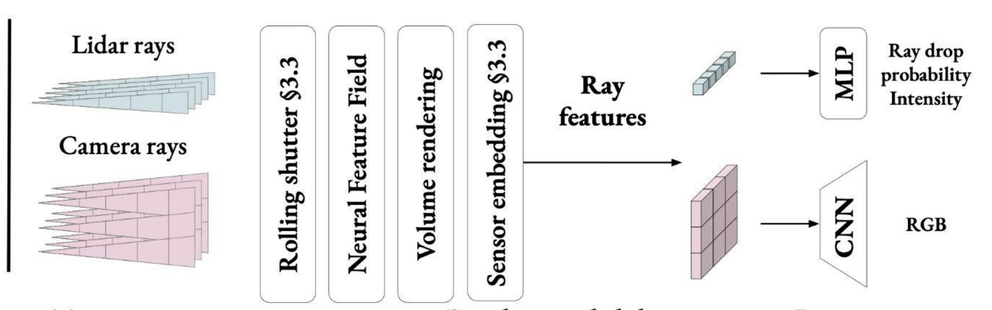

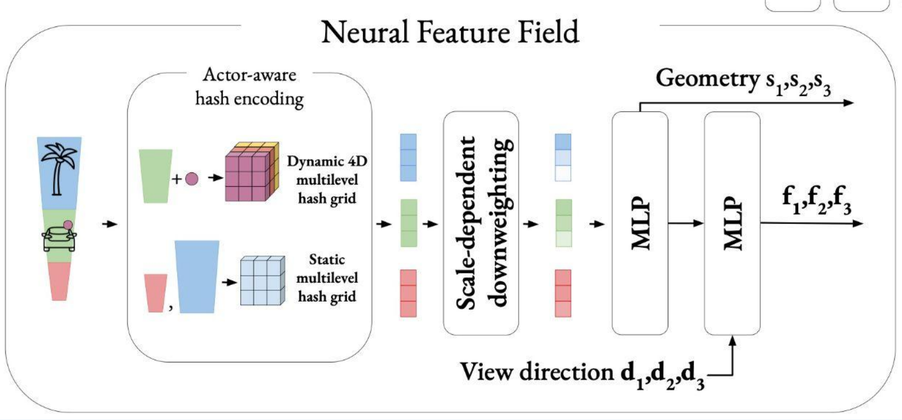

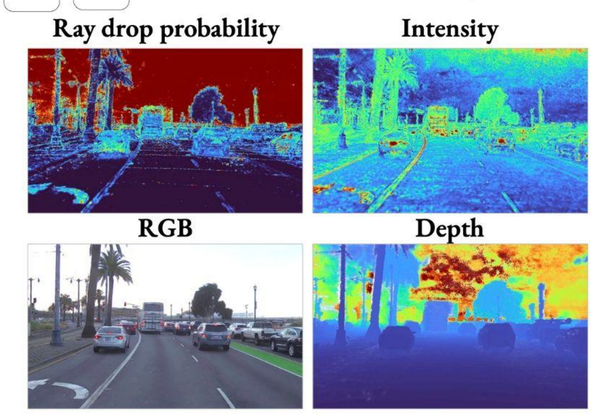

NeuroNCAP: Closed-loop safety testing

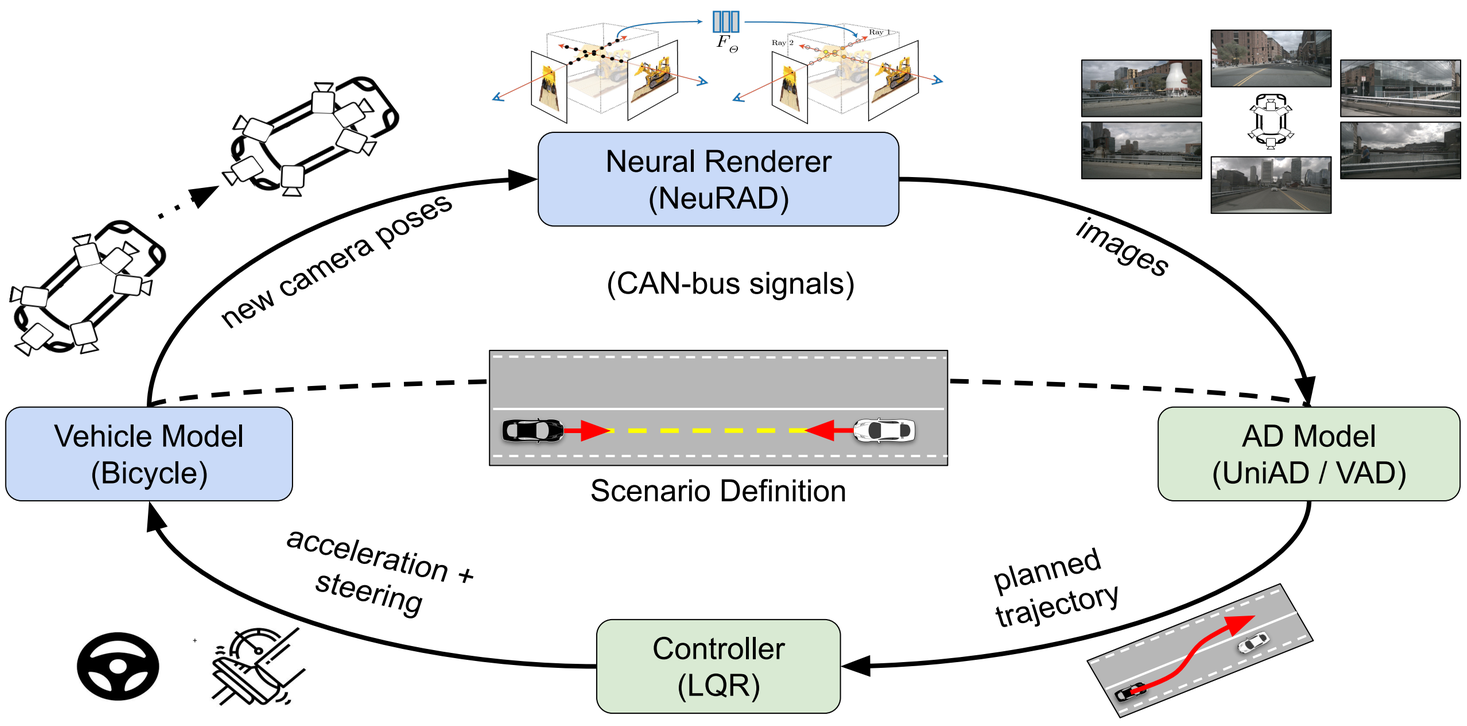

SoTA methods crash a lot!

Model	Post-proc.	NeuroNCAP Score ↑ Avg.	Stat.	Frontal	Side	Collision rate (%) ↓ Avg.	Stat.	Frontal	Side
Base-U	-	2.65	4.72	1.80	1.43	69.9	9.6	100.0	100.0
Base-V	-	2.67	4.82	1.85	1.32	68.7	6.0	100.0	100.0
UniAD	x	0.73	0.84	0.10	1.26	88.6	87.8	98.4	79.6
VAD†	x	0.66	0.47	0.04	1.45	92.5	96.2	99.6	81.6
UniAD†	✓	1.84	3.54	0.66	1.33	68.7	34.8	92.4	78.8
VAD	✓	2.75	3.77	1.44	3.05	50.7	28.7	73.6	49.8

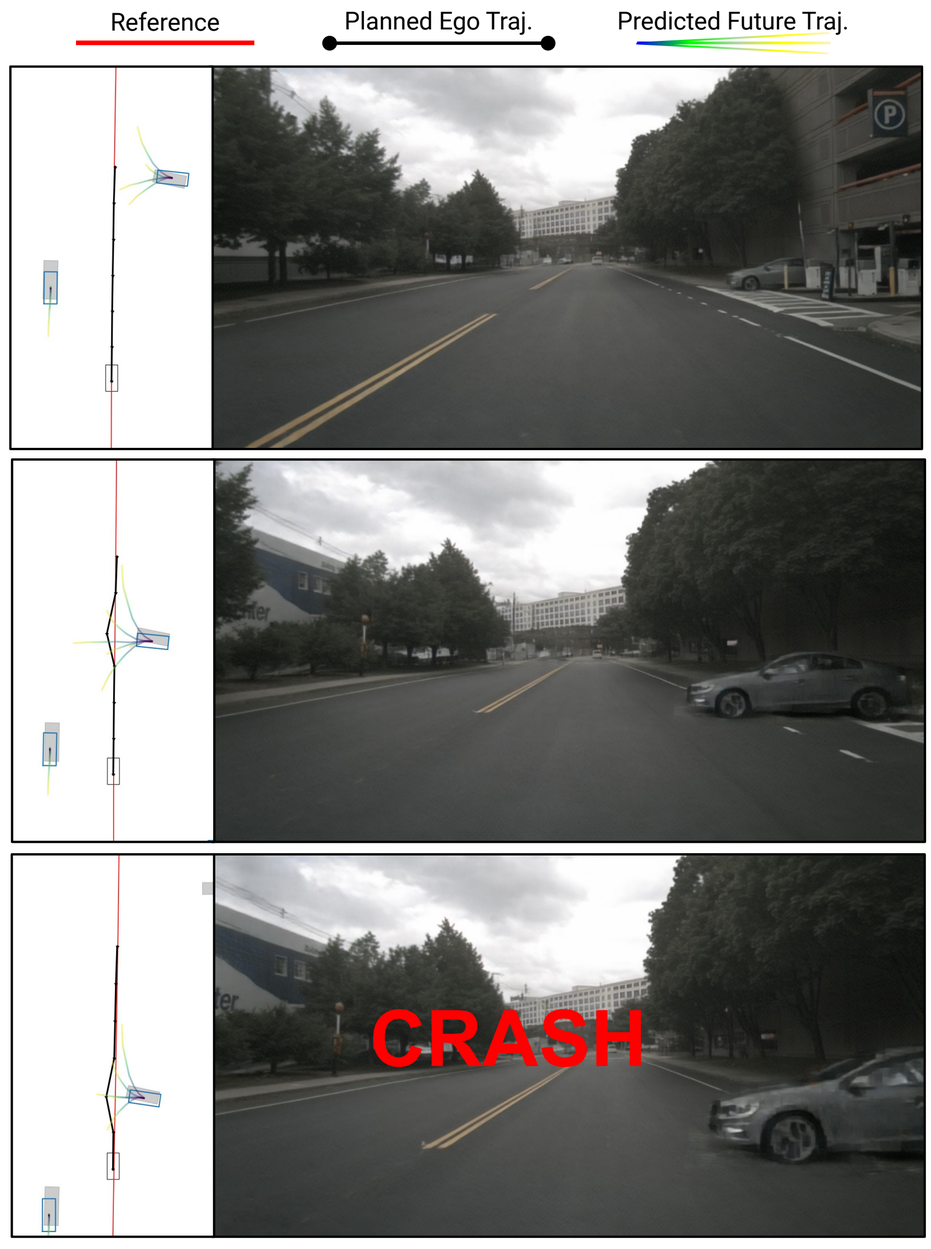

Construct safety-critical scenarios in the wild

With inspiration from the industry standard EuroNCAP, we define three safety-critical scenario types...

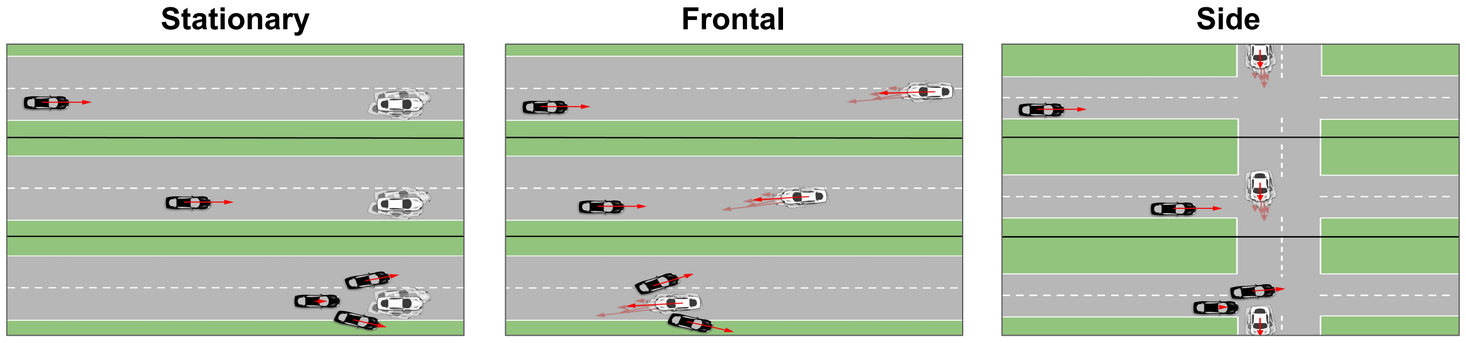

...and then we use neural rendering to turn a nominal driving scene into these types of scenarios

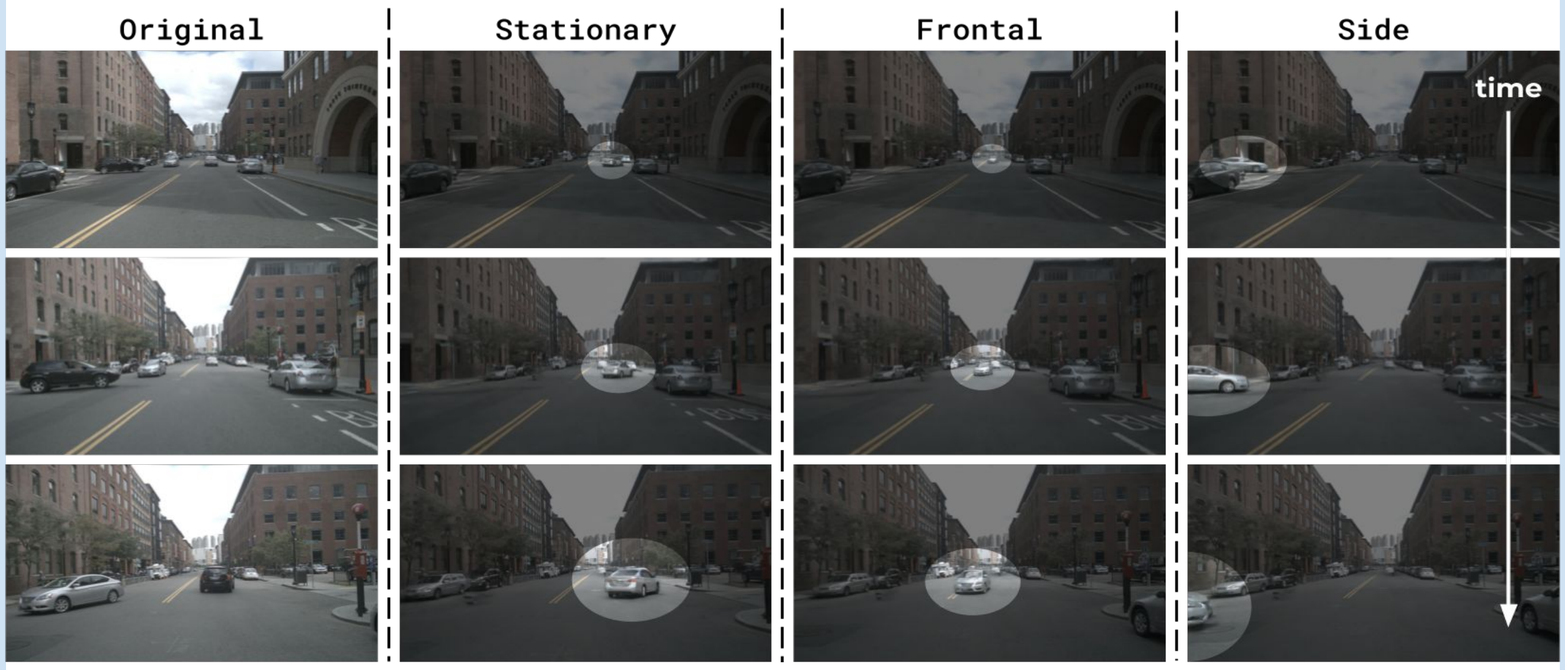

Generative actor refinement

Problem: NeRF and 3DGS cannot reliably extrapolate to unseen regions.
Solution:

- Initial training on seen regions
- Render unseen regions
- Improve rendered images w/ generalist diffusion models (e.g. SD3, FLUX)
- Use refined images to finetune NeRF, 3DGS reconstructions