DRUM: Diffusion-based Raydrop-aware Unpaired Mapping for Sim2Real LiDAR Segmentation

ICRA IEEE INTERNATIONAL CONFERENCE ON ROBOTICS AND AUTOMATION 2026 VIENNA

Tomoya Miyawaki[1] **Kazuto Nakashima**[1] **Yumi Iwashita**[2] **Ryo Kurazume**[1]
[1]Kyushu University, Japan [2]NASA/Caltech JPL, USA

TL;DR: DRUM uses a real-data diffusion prior to turn labeled synthetic LiDAR scans into pseudo-real scans, improving Sim2Real LiDAR segmentation.

Motivation

Problem: Manual annotation for LiDAR segmentation is costly.

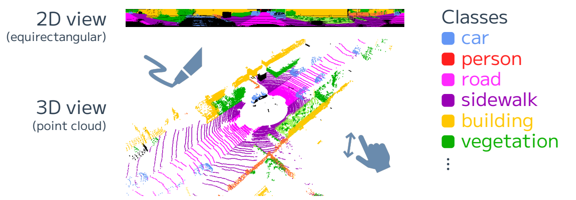

Simulators can automatically generate labeled LiDAR samples, but the samples lack real-world LiDAR effects: **reflectance** & **raydrop**.

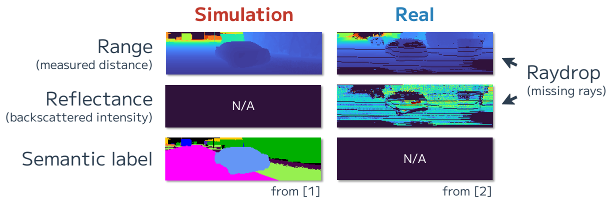

Goal

Given **labeled simulation** scans and **unlabeled real** scans, generate a **labeled pseudo-real** dataset for LiDAR segmentation.

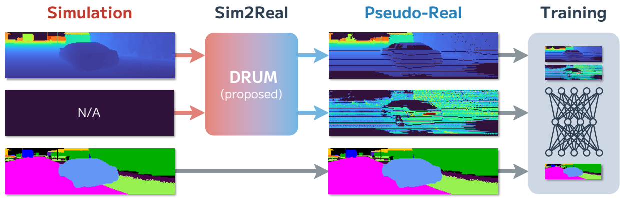

Our approach:

- Reproduce **reflectance** and **raydrop** in simulation scans.

Challenges:

- Simulation and real datasets are **unpaired**.
- Reflectance requires **completion**, while raydrop requires **corruption**.

Method

1 Model the real-data distribution using a diffusion model.

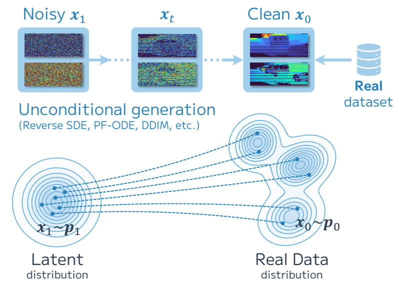

- An unconditional diffusion model is trained on unlabeled real LiDAR scans.
- Range and reflectance are represented as a two-channel equirectangular image, with range values on a logarithmic scale.

2 Map a simulation sample into the real domain using raydrop-aware masked guidance.

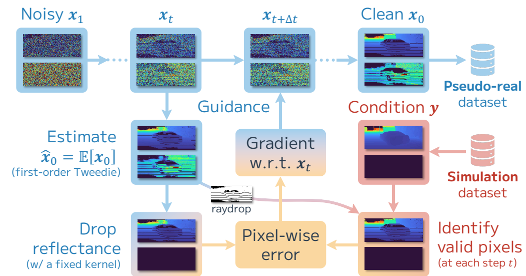

- **Problem:** The missing reflectance in the simulation input y can be inpainted using guidance techniques [4], but naïve guidance suppresses raydrop because y is "too clean."
- **Key idea:** We identify valid pixels from the tentative estimate $\hat{x}_0$ and apply masked guidance to preserve the generated raydrop.

3 Train segmentation models on the pseudo-real dataset.

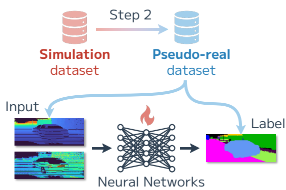

- Any type of model can be trained with the pseudo-real input and the labels provided by the sim.

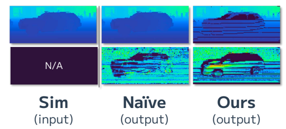

Implementation: We employ R2DM [3], a pixel-space diffusion model of range & reflectance images. The guidance framework is built upon ΠGDM [4].

Evaluation

Q1: Does DRUM improve Sim2Real LiDAR segmentation?

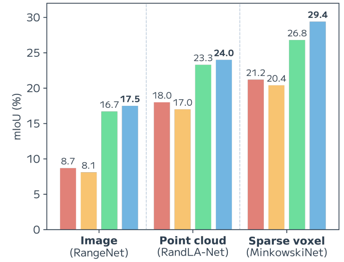

Qualitative results of Image-based segmentation

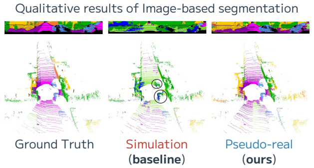

Settings:

- A diffusion model is pre-trained on SemanticKITTI "train" [2] w/o labels
- Segmentation models (3 types) are trained on SynLiDAR "all" [1] and evaluated on SemanticKITTI "val" [2]

Comparison methods:

- Range (sim)
- Range (sim), Reflectance (estimation by [1])
- Range (sim), Raydrop (**ours**)
- Range (sim), Raydrop (**ours**), Reflectance (**ours**)

Results:

- Our pseudo-real data improves mIoU across image, point, and voxel representations.
- Large gains in the vehicle classes: e.g., RangeNet improves car IoU by +51 pts.

Q2: What if some real scans are labeled?

Settings:

- Use {1%, 20%, 50%, 100%} of labeled real data with/without our labeled pseudo-real scans.
- Model: RangeNet (Image-based segmentation)

Results:

- Consistent improvement over multiple ratios.
- Large gains in the minority classes: +15 pts for person IoU, +22 pts for bicyclist IoU.

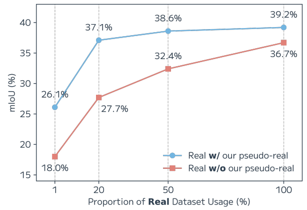

Project page
w/ code & demo

[1] **SynLiDAR** (Xiao et al. AAAI'22), [2] **SemanticKITTI** (Behley et al. ICCV'19), [3] **R2DM** (Nakashima et al. ICRA'24), [4] **ΠGDM** (Song et al. ICLR'23) [5] **RangeNet** (Milioto et al. IROS'19) [6] **RandLA-Net** (Hu et al. CVPR'20) [7] **MinkowskiNet** (Choy et al. CVPR'19)