LiDAR Snowfall Simulation for Robust 3D Object Detection

ETH zürich · CVL Computer Vision Lab · TRACE
Toyota Research on Automated Cars in Europe

Martin Hahner[1], Christos Sakaridis[1], Mario Bijelic[2], Felix Heide[2], Fisher Yu[1], Dengxin Dai[1,3], and Luc Van Gool[1,4]

[1] ETH Zurich [2] Princeton University [3] MPI for Informatics [4] KU Leuven

KU LEUVEN · PRINCETON UNIVERSITY

CVPR JUNE 19-24 2022 NEW ORLEANS • LOUISIANA

What is a LiDAR sensor?

- a LiDAR sensor is an *active illumination* sensor which can measure precise distances to objects in a scene via individual light beams
- in contrast to *passive illumination* sensors, such as RGB cameras, a LiDAR sensor operates equally well at day and night-time
- a LiDAR sensor typically consists of a light transmitter Tx and receiver Rx which have (see Figure below)
- the transmitter transmits light beams and the receiver measures the time it takes for the light beams to return to the sensor
- if a transmitted light beam is received by the receiver, the beam must have hit a reflective surface or object ("target")
- leveraging the constant speed of light, the distance to that target can then typically be determined with an accuracy of +/- 3 cm

ToF LiDAR example © Li et al. - What happens to a ToF LiDAR in fog?

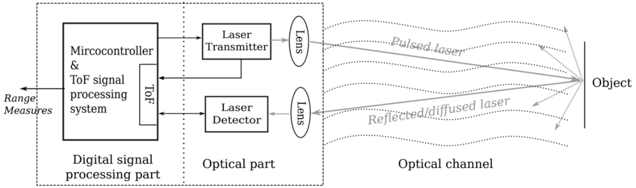

What are the problems?

- LiDAR sensors operate with light beams in the near infrared light (NIR) spectrum, typically in the range of 900-1500 nm
- light beams at such a wavelength do **not** penetrate water particles of any aggregation state
- in snowfall, this leads to two things:
  - backscatter : distance measurements originating from snowflakes in the atmosphere
  - attenuation: reduced intensity measurements of valid targets due to intereference with snowflakes
- this **domain shift** from clear to snowy conditions (due to the aforementioned attenuation and backscatter of snowflakes) negatively impacts the performance of downstream tasks
- most of today's autonomous driving datasets are captured in clear weather and do contain **no snowy (training) data**

What is our proposed solution?

- we propose **a snowfall and wet ground simulation pipeline** that takes a LiDAR point cloud captured in clear weather as input and outputs a point cloud of the same scene as if it was captured in wintery snowfall conditions
- with this pipeline we can generate **lots of snowy LiDAR data with various snowfall intensities and ground wetness levels** (since there are so many datasets available that were captured in clear conditions and therefor fit our input criteria)
- with this amount of data, **for the first time**, now we have enough snowy data to **train downstream machine learning tasks**

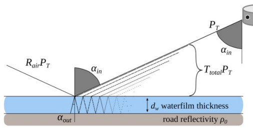

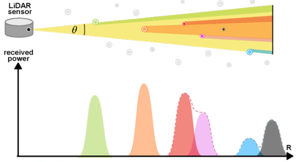

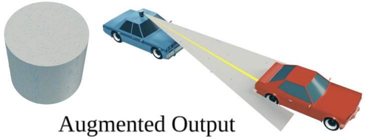

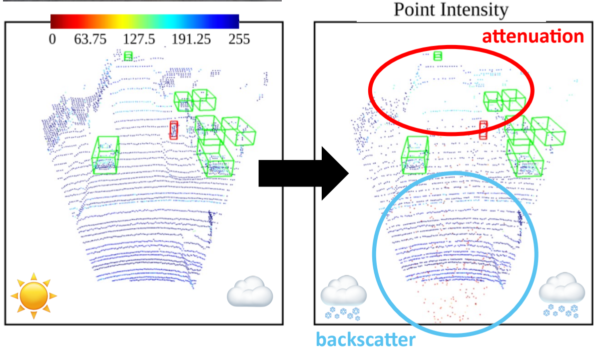

What are our results?

- we show **consistent performance improvements** of several recent 3D object detection methods in scenes captured in real-world snowfall without sacrificing performance on clear scenes

3D object detection in real-world clear weather / heavy snowfall

moderate cars	PV-RCNN		VoxelRCNN		CenterPoint	
trained on clear	45.36	39.69	45.19	39.47	44.11	38.68
trained on ours	**45.71**	**41.79**	**45.20**	**40.76**	**44.33**	**40.14**
3D AP Δ	+0.35	+2.10	+0.01	+1.29	+0.22	+1.46

Conclusion

- we propose **an algorithm to convert LiDAR point clouds** captured in clear weather into LiDAR point clouds as if they were captured in wintery snowfall conditions

  with this algorithm
  - we address the lag of **snowy LiDAR training data**
  - we can process not only one specific, but **any LiDAR dataset**
  - we can generate not only one snowfall intensity and one ground wetness level of an individual scene, but really **any desired snowfall intensity and ground wetness level**
- we show that training on this augmented data leads to improvements on 3D object detection, a **major downstream task** for LiDAR point clouds
- for more information, please refer to the paper, the supplementary materials and/or our published code, all available at the link below

trace.ethz.ch/lidar_snow_sim