Fog Simulation on Real LiDAR Point Clouds for 3D Object Detection in Adverse Weather

ICCV 2021 VIRTUAL OCTOBER 11-17

Martin Hahner [1] mhahner@vision.ee.ethz.ch
Christos Sakaridis [1] csakaridis@vision.ee.ethz.ch
Dengxin Dai [1,2] ddai@mpi-inf.mpg.de
Luc van Gool [1,3] luc.vangool@kuleuven.be

[1] ETH Zürich [2] MPI for Informatics [3] KU Leuven

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 Tx transmits light beams and the receiver Rx 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 *c*, the distance to that target can then typically be determined with an accuracy of +/- 3 cm

ToF LiDAR example

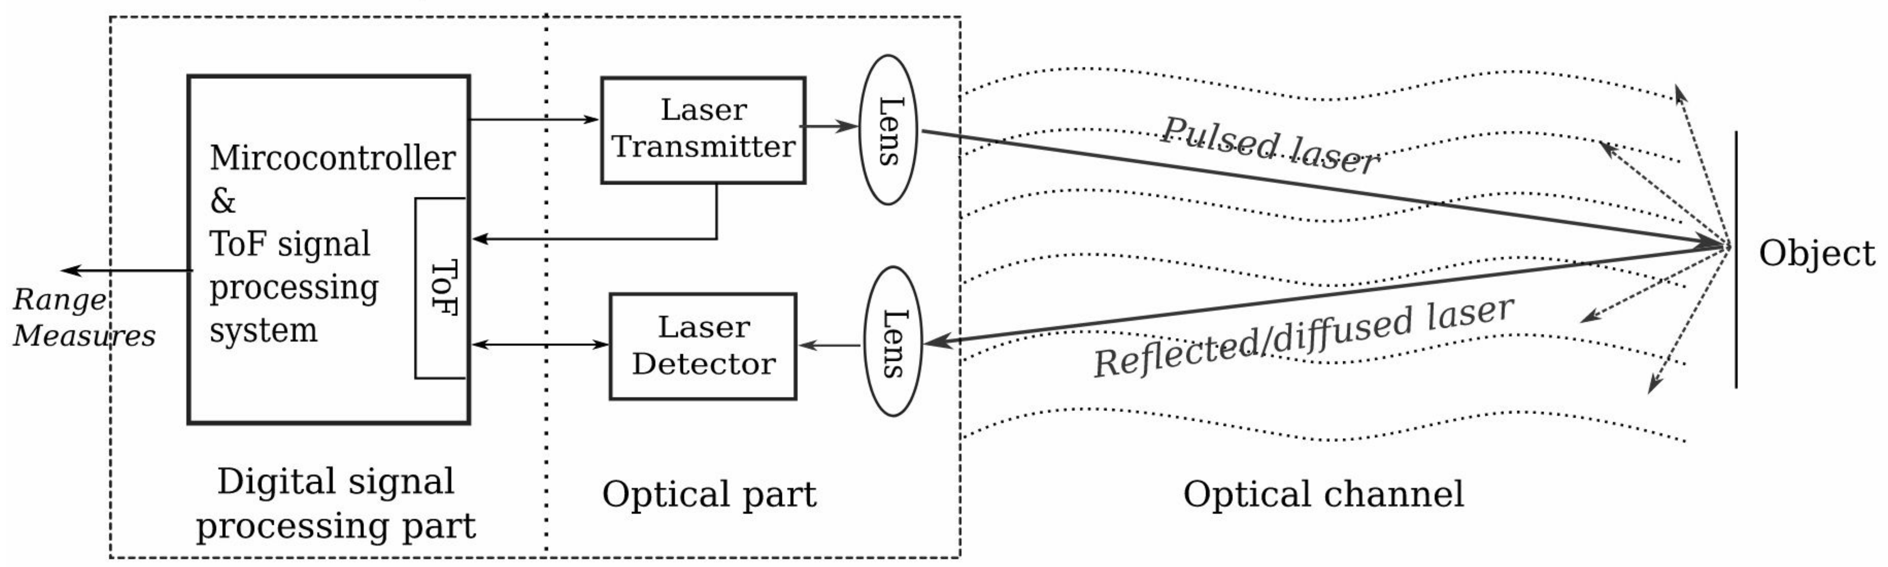

© Li et al. - What happens to a ToF LiDAR in fog?

What are the problems?

- LiDAR sensors operate with light beams in the near infrared light (NIR) spectrum, typically ~900-1500 nm wavelength
- light beams at such a wavelength do **not** penetrate water particles
- in fog, this leads to two things:
  - backscatter: distance measurements originating from the accumulation of water particles in the foggy atmosphere
  - attenuation: objects of interest in the fog are not reached/detected by the emitted light beams due to attenuation
- this **domain shift** from clear to foggy conditions (due to the aforementioned attenuation and backscatter of water particles) negatively impacts the performance of downstream tasks
- most of today's autonomous driving datasets are captured in clear weather and do contain **no foggy (training) data**

What is our proposed solution?

- we propose **a fog 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 foggy conditions
- with this pipeline we can generate physically realistic foggy LiDAR data
- in fact, **lots of foggy LiDAR data with various fog densities** and scene complexity (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 foggy data to **train downstream machine learning tasks**

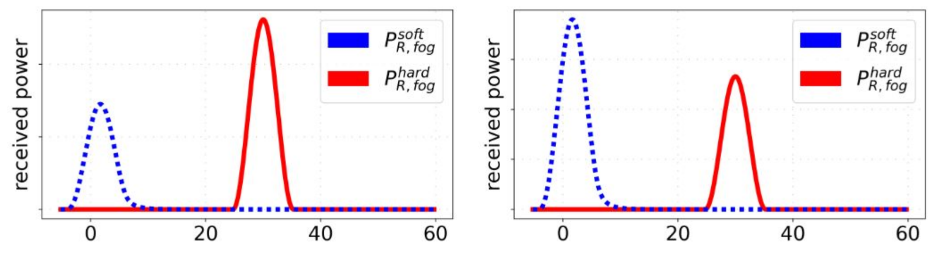

(a) $P^{soft}_{R,fog}(R) < P^{hard}_{R,fog}$ (b) $P^{soft}_{R,fog} > P^{hard}_{R,fog}$

Figure 4: The two terms of the received signal power $P_{R,fog}$ from a single LiDAR pulse, associated to the solid object that reflects the pulse ($P^{hard}_{R,fog}$) and the soft fog target ($P^{soft}_{R,fog}$), plotted across the range domain. While in (a) the fog is not thick enough to yield a return, in (b) it is thick enough to yield a return that overshadows the solid object at $R_0 = 30$m.

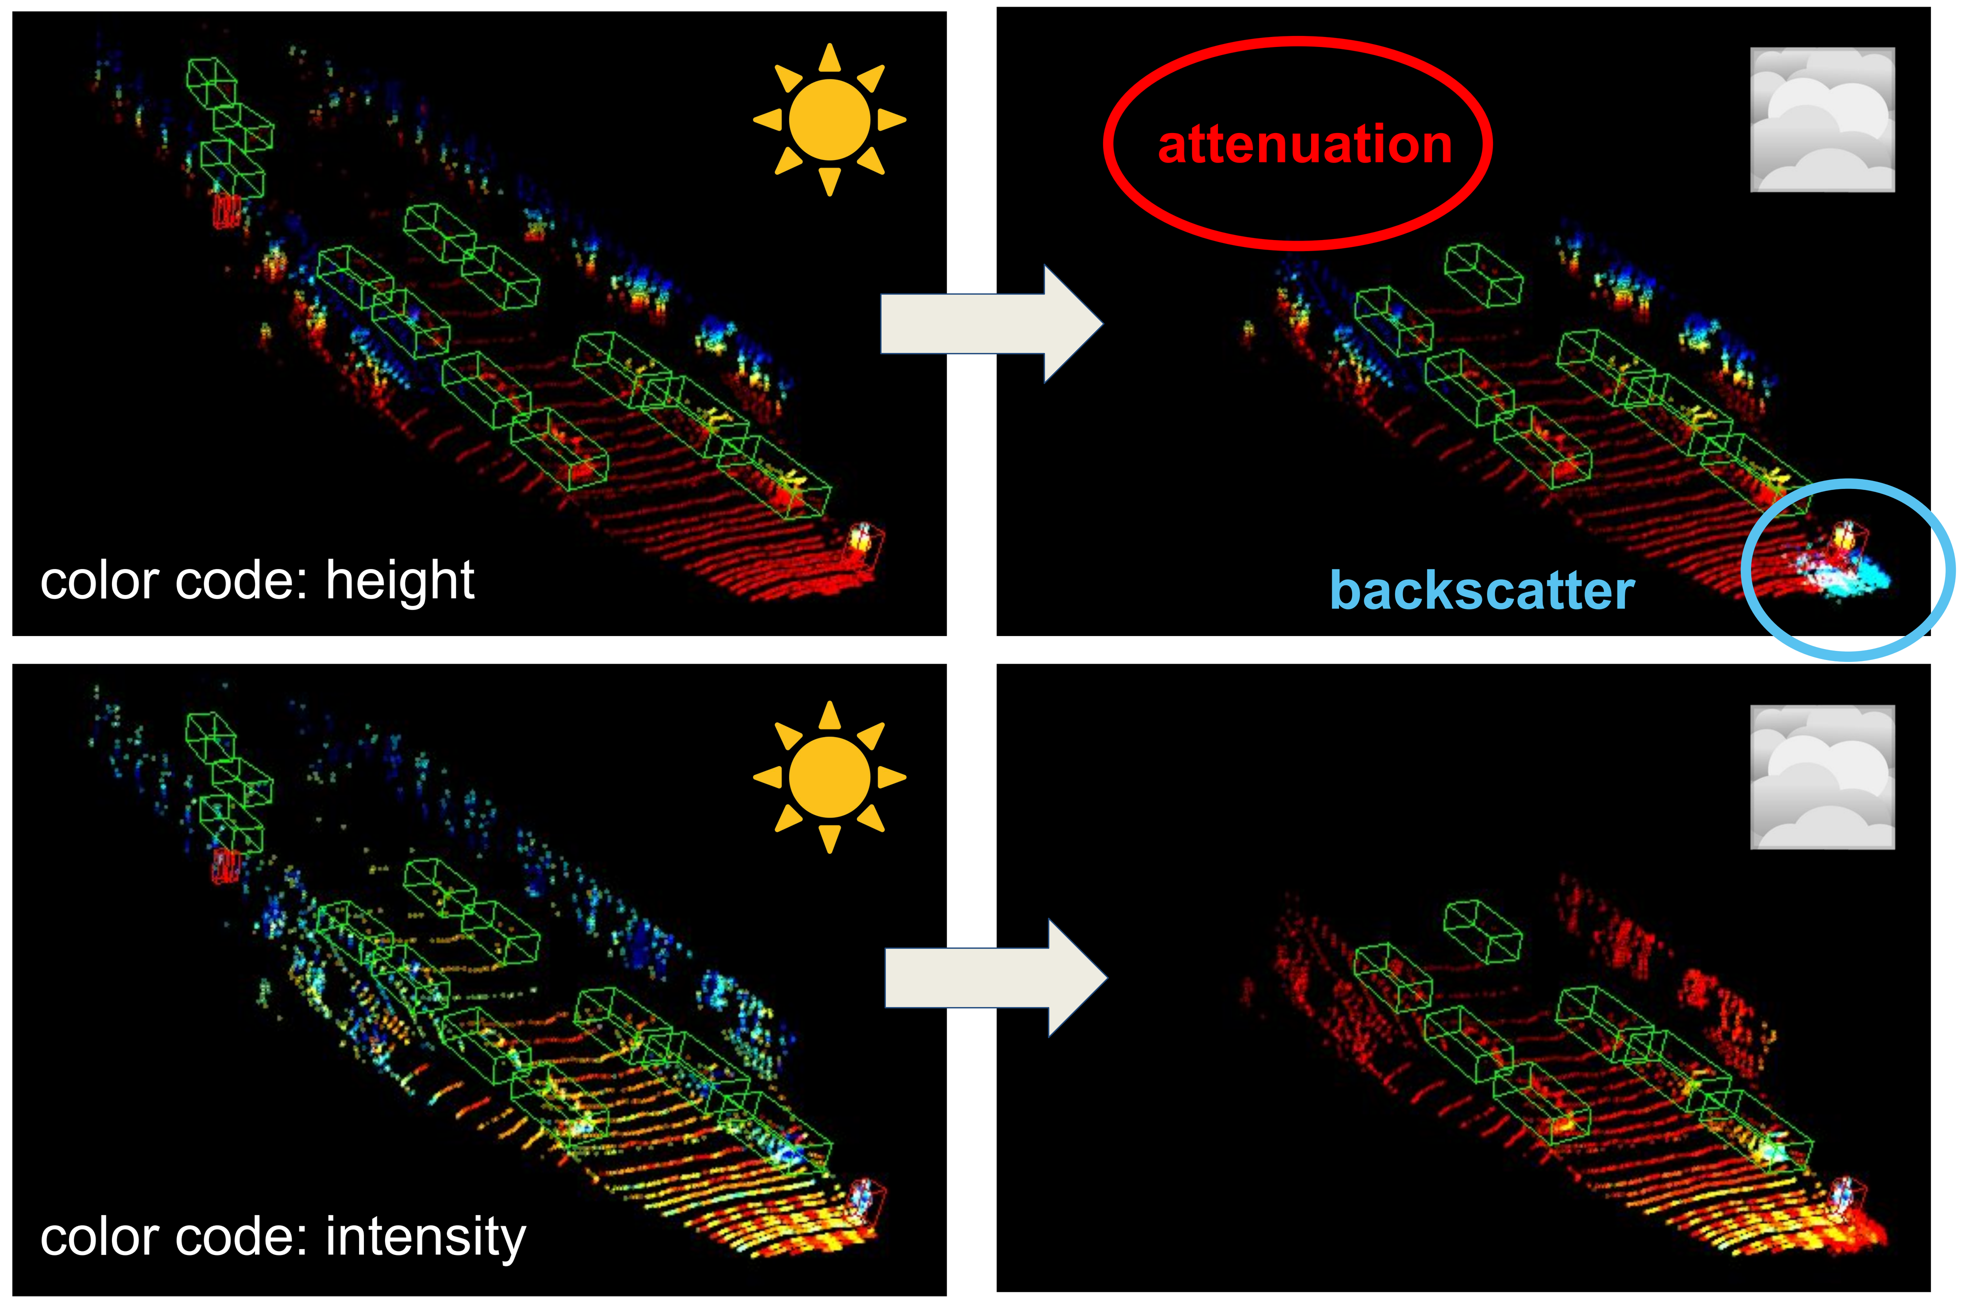

What are our result?

- we show **consistent performance improvements** of several state-of-the-art 3D Object Detection methods in scenes captured in real-world foggy conditions without sacrificing performance on clear scenes

3D Object Detection in real-world clear weather / dense fog

mod. Car AP@.5IoU	PointPillars	Part-A²	SECOND	PointRCNN	PV-RCNN
trained on clear scenes	74.64 / 35.23	73.79 / 38.15	75.20 / 42.99	76.52 / 45.03	77.05 / 46.00
trained on ours	74.56 / **39.14**	73.86 / **42.75**	75.95 / **44.01**	76.58 / **47.99**	76.75 / **47.38**
Δ	-0.08 / **+3.91**	+0.07 / **+4.60**	+0.75 / **+1.02**	+0.06 / **+2.96**	-0.30 / **+1.38**

Conclusion

- we propose **an algorithm to convert LiDAR point clouds** captured in clear weather into LiDAR point clouds as if they were captured in foggy conditions
- with this algorithm
  - we address the lag of **foggy LiDAR training data**
  - we can process not only one specific, but **any LiDAR dataset**
  - we can generate not only one fog density level of an individual scene, but really **any desired fog density 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
- for a **live demo** of the published code, join our **Q&A session**

Paper, Code, Supplementary Materials, all available at → trace.ethz.ch/lidar_fog_sim