LiDAR Snowfall Simulation for Robust 3D Object Detection

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

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

Abstract

3D object detection is a central task for applications such as autonomous driving, in which the system needs to localize and classify surrounding traf c agents, even in the presence of adverse weather. In this paper, we address the problem of LiDAR-based 3D object detection under snowfall. Due to the dif culty of collecting and annotating training data in this setting, we propose a physically based method to simulate the effect of snowfall on real clear-weather LiDAR point clouds. Our method samples snow particles in 2D space for each LiDAR line and uses the induced geometry to modify the measurement for each LiDAR beam accordingly. Moreover, as snowfall often causes wetness on the ground, we also simulate ground wetness on LiDAR point clouds. We use our simulation to generate partially synthetic snowy LiDAR data and leverage these data for training 3D object detection models that are robust to snowfall. We conduct an extensive evaluation using several state-of-the-art 3D object detection methods and show that our simulation consistently yields signi cant performance gains on the real snowy STF dataset compared to clear-weather baselines and competing simulation approaches, while not sacri cing performance in clear weather. Our code is available at github.com/SysCV/LiDAR_snow_sim.

1. Introduction

A light detection and ranging (LiDAR) sensor is an active range sensor useful for several applications [10, 27, 39, 61]. Its high-quality 3D output renders LiDAR the modality of choice for several tasks that require 3D reasoning, such as 3D object detection [25, 59]. As LiDAR sensors are becoming increasingly cheaper [53], their integration into autonomous cars becomes increasingly feasible as well.

Nonetheless, previous sensor tests have revealed that such active pulsed systems are vulnerable in scattering media, leading to decreases of perception distances in various weather conditions such as rain [4, 14, 52], fog [2, 4, 18, 19, 52], and snow [19, 24, 28], as shown in Fig. 1. In these conditions, the optical medium contains particles of water or snow which interact with the laser beam and absorb, re ect or refract its photons. This results in two effects: (i) attenuation of the received power that corresponds to the target at the line of sight, and (ii) backscattering from particles leading to spurious maxima in the received power and thus to spurious returns at ranges different from the true range of the target. Consequently, there is a severe degradation of measurement quality due to intense noise, a large domain shift relative to point clouds captured in clear weather, and hence a detrimental effect on performance of high-level tasks such as 3D object detection [1, 16]. Yet, achieving robust perception in adverse weather is a desirable goal as fatality rates for human drivers are notably higher in adverse weather, as reported by the US Department of Transportation [49] and the European Commission [8].

Figure 1. 3D object detection results in heavy snowfall with prior training on the proposed data augmentation scheme (top right) in comparison to no augmentation (top left). The bottom row shows the RGB image as reference.

Since adverse-weather data are hard to collect [1], previous works have investigated simulation methods to close the domain gap for camera data in fog [36] and rain [48].

Figure 2. Sketch of a LiDAR sensor where the transmitter Tx and the receiver Rx do not have coaxial optics, but have parallel axes (called a bistatic beam configuration).

The number of detection stages is most often classified into single [25] and two-stage approaches [39–41, 57, 59, 62], where single-stage approaches directly discretize the input space and predict objects for each individual cell following [26]. Two-stage approaches first predict proposals and refine them in a subsequent pooled feature space following the general idea of [34].

3. Snowfall Simulation on LiDAR Point Clouds

Pulse propagation in free space can be modeled with geometrical optics for cost-effective LiDAR systems. Such systems apply an array of synchronized near-field infrared pulse emitters Tx and avalanche photodiodes (APDs) as receivers Rx depicted in Fig. 2 and described in [55]. The sent-out laser pulse P_0 is reflected by a solid scene object, often referred to as target, with reflectivity ρ_0, and it is captured by the receiver, providing the time delay t of the captured echo and its corresponding power P_R. The object distance R is calculated by applying $R = \frac{ct}{2}$, where c is the speed of light. The 3D position $[x, y, z]$ of the object is obtained by using the direction in which the pulse was emitted. For extended objects, geometric optics [11] can be applied to model the received power P_R following

$$P_R(R) = C_A P_0 \rho_0 \frac{\cos(\alpha_{in})}{R^2}; \quad (1)$$

which holds for objects with a diameter larger than the beam diameter at distance R and requires additional information about (i) the incident angle α_{in} and (ii) the system constant C_A independent of range and time. However, the received laser power is typically corrected [11], as C_A differs for each scanning layer due to different optics and beam divergences. Four different levels of intensity calibration can be reported according to [21]. For the Velodyne HDL-S3D sensor used in our experiments, a beam divergence correction is applied following the sensor manual [50]. This correction is defined as

$$i = P_R + f_s \left(f_o - \left(1 - \frac{R}{R_{max}}\right)^2\right); \quad (2)$$

where f_s is the focal slope and f_o is the focal offset. The parameters for each laser are retrieved from the factory side calibration. Before applying the proposed simulation methods, we first retrieve the raw intensities by inverting this intensity calibration. In snowfall, the optical medium contains particles which are smaller than the beam diameter, so Mie scattering and the exact spatial distribution of the particles must be taken into account [33].

Pulse propagation in the presence of scattering particles is described by a linear model introduced in [33], which is valid for non-elastic scattering. This model expresses the range-dependent received power P_R as a time-wise convolution between the time-dependent transmitted signal power P_T and the impulse response H of the optical system:

$$P_R(R) = C_A \int_0^{2R/c} P_T(t)\, H\left(R - \frac{ct}{2}\right) dt; \quad (3)$$

with the time signature of the transmitted pulse given by

$$P_T(t) = \begin{cases} P_0 \sin^2\left(\frac{\pi}{2\tau_H} t\right); & 0 \le t \le 2\tau_H \\ 0 & \text{otherwise} \end{cases}; \quad (4)$$

τ_H is the half-power pulse width, set to 10 ns for the Velodyne HDL-S3D sensor. The impulse response H can be factored into the impulse responses of the optical channel, H_C, and the target, H_T:

$$H(R) = H_C(R)\, H_T(R). \quad (5)$$

H_C depends on the beam divergence, the overlap of transmitter and receiver described by $\eta(R)$ as well as the transmittance $T(R)$ of the medium through

$$H_C(R) = \frac{T^2(R)}{R^2}\, \eta(R). \quad (6)$$

The transmittance $T(R)$ is equal to 1 in the part of the medium that is not occupied by snow particles, assuming absence of other scatterers. The overlap $\eta(R)$ can be geometrically derived from Fig. 2 as

$$\eta(R) = \begin{cases} 0; & R \le R_1 \\ \frac{R - R_1}{R_2 - R_1}; & R_1 < R < R_2 \\ 1; & R_2 \le R. \end{cases} \quad (7)$$

The impulse response of the target, H_T, allows us to model snow particles as we detail in the following.

Scene reflection defines the particle interaction with the laser pulse through H_T. For an extended solid target object we can write

$$H_T(R) = \rho_0\, \delta(R - R_0); \quad (8)$$

with ρ_0 being the reflectivity of the object and δ the Dirac delta function. However, in snowfall, apart from the solid target object, the laser beam is also partially reflected by snow particles.

Figure 5. A real-world capture on a dry highway (top), a real-world capture with a water height of $d_w = 0.53$ mm (middle) and the synthesized road wetness from the clear reference (bottom).

Algorithm 1 LiDAR snowfall simulation

```
procedure SNOWFALL(pc, n_l, ..., R_max, ..., r_s, ..., 0)
    for l in n_l do                                   . for each layer
        pc_l ← pc.SELECT(layer = l)
        f_d, f_s, i_max ← LOAD_CALIB(l)
        f_o ← (1 − f_d/13100)^2                       . focal offset [50]
        s ← SAMPLE_SNOWFLAKES(R_max, r_s)             . in 2D [15]
        for p in pc_l do                              . for each point in layer
            x, y, z, i ← p
            R_0 ← k_p k_2
            t ← GET_PARTICLES_IN_BEAM(s, x, y, R_0, ...)   . in 2D
            if len(t) > 1 then                        . otherwise no interference
                P_{R,snow} ← 0                        . initialize with zeros
                for i, R, ... in t do                 . for each target
                    if R = R_0 then                   . original target
                        P_R ← i · f_s · f_o · (1 − R/R_max)^2
                        C_A ← P_0 · P_R/P_0 · R_0^2   . follows from Eq. (1)
                    else                              . snow flake
                        P_R ← s · i_max
                        C_A ← P_0 · P_R/P_0
                    end if
                    P_{R,snow} += Eq.(11)(C_A, P_0, R, ..., ..., ...)
                end for
                P_R ← max(P_{R,snow})
                R ← arg max(P_{R,snow}) − c · H/2
                i ← P_R + i_max · f_s · f_o · (1 − R/R_max)^2
                (x, y, z) ← R/R_0 · (x, y, z)
                p ← x, y, z, i
            end if
        end for
    end for
    return pc
end procedure
```

Otherwise the target object intensity is attenuated according to its occlusion percentage. Our complete snowfall simulation is presented in Algorithm 1. In Fig. 3 we show a winterly example scene, once augmented with our snowfall simulation and once with the one proposed in LISA [22].

3.1. Wet Ground Model

LiDAR readings are affected by the wetness of surrounding surfaces. Emitted light pulses are re ected specularly from wet ground leading to signi cantly attenuated laser echoes depending on the water height [5]. Analysing the road wetness statistics of STF [1] (given in the supplementary materials), it becomes apparent that wet roads occur together with adverse weather such as snowfall and are the main cause for lost points on road surfaces (see Fig. 5).

Figure 6. Visualization showing the geometrical optical model which describes the re ection on a wet road surface.

To model the attenuation caused by road wetness, we apply geometric optics modeling single rays and their refraction on a thin layer of water illustrated in Fig. 6. A qualitative example is shown in Fig. 5 (bottom). We use the refractive indices (n_{in}, n_{out}) and angles $(\theta_{in}, \theta_{out})$ at the transition point. The angle θ_{out} can be calculated based on Snell's law:

$$\sin(\theta_{out}) = \frac{n_{in}}{n_{out}} \sin(\theta_{in}). \tag{12}$$

The change in pulse amplitude is modeled by the Fresnel equations described by the perpendicular and parallel transmissionst$^?, t^=$ and re ectionsr$^?, r^=$ with respect to the ground, respectively (details in the supplementary materials). Going from amplitude to transmitted power, we can deduce the power re ectionR_k^p and transmissionT_k^p:

$$R_k^p = (r_k^p)^2; \tag{13}$$

$$T_k^p = \frac{n_2 \cos(\theta_{in})}{n_1 \cos(\theta_{out})} t_k^p; \tag{14}$$

where$p \in (?, =)$ refers to the polarization. We sum up all individual light rays which traverse back to the sensor, as shown in Fig. 6, leading to the geometric series

$$T_{total}^p = T_{air}^p \, {}_0T_{water}^p \sum_{k=0}^{\infty} ({}_0R_{water}^p)^k; \tag{15}$$

which can be simpli ed for$({}_0R_{water}) < 1$ to

$$T_{total}^p = T_{air}^p \, {}_0T_{water}^p / (1 - {}_0R_{water}^p). \tag{16}$$

As the actual polarization is proprietary, we assume that the manufacturer optimized the polarization depending on the distance for best possible performance, implying

$$T_{total} = \max\left(T_{total}^=, T_{total}^?\right). \tag{17}$$

Detection method	Simulation method	heavy snowfall ↑				light snowfall ↑				clear weather ↑			
		0-80m	0-30m	30-50m	50-80m	**0-80m**	0-30m	30-50m	50-80m	**0-80m**	0-30m	30-50m	50-80m
PV-RCNN [39]	*None*	39.69	65.05	36.14	8.03	41.13	69.24	39.72	**11.68**	45.36	**72.34**	42.48	10.53
	Fog [16]	38.19	64.72	33.38	7.49	39.82	68.41	38.68	9.65	43.37	71.05	40.03	9.90
	DROR [6]	38.57	64.27	35.40	8.07	39.33	66.73	38.14	10.51	41.44	67.76	38.48	9.44
	LISA [22]	39.21	64.21	35.34	8.64	41.60	69.15	**41.08**	11.15	45.30	71.06	42.86	11.45
	Ours-wet	40.03	65.34	35.82	**9.31**	41.07	68.49	40.03	11.02	44.81	71.60	42.71	10.63
	Ours-snow	41.61	67.44	**37.47**	8.84	41.20	68.79	40.20	11.13	45.61	72.14	**43.40**	11.21
	Ours-snow+wet	**41.79**	**68.39**	37.14	8.85	**41.79**	**70.30**	41.01	11.28	**45.71**	71.88	43.31	**11.69**
VoxelRCNN-Car [9] (single class method)	*None*	39.47	65.14	36.29	6.83	41.25	69.12	39.86	**11.81**	45.19	72.33	**43.20**	10.69
	Fog [16]	40.06	65.58	36.78	7.33	41.10	68.93	39.25	10.98	44.46	71.67	41.78	10.84
	DROR [6]	38.16	64.97	33.23	6.83	38.48	66.93	35.68	9.97	40.65	67.94	36.85	8.45
	LISA [22]	39.06	66.61	33.56	6.93	40.68	68.80	38.78	10.75	45.03	72.05	41.96	10.59
	Ours-wet	40.28	67.37	36.03	6.95	40.89	69.41	39.28	10.37	43.98	72.08	40.23	9.23
	Ours-snow	**41.20**	**68.27**	37.18	**7.90**	**41.75**	**70.22**	**40.95**	11.78	44.52	**72.40**	42.23	10.39
	Ours-snow+wet	40.76	66.58	**37.74**	7.78	41.57	68.94	40.82	11.66	**45.20**	71.30	42.84	**11.23**
CenterPoint [59]	*None*	38.68	63.50	34.20	8.23	40.91	68.36	39.96	10.98	44.11	71.66	40.97	9.76
	Fog [16]	38.82	64.43	33.70	8.69	40.82	68.71	40.05	10.48	43.79	70.34	41.22	10.77
	DROR [6]	38.42	64.47	34.31	8.27	38.69	65.62	37.59	10.26	40.80	67.62	36.61	8.69
	LISA [22]	38.11	63.74	33.58	7.67	40.26	68.16	39.25	11.20	44.70	71.54	41.46	**11.23**
	Ours-wet	39.03	63.17	36.08	7.69	40.65	68.59	39.58	10.87	43.85	70.20	41.66	10.82
	Ours-snow	39.81	64.75	35.77	**9.06**	40.68	68.71	39.65	10.40	**44.76**	**71.90**	**42.91**	10.13
	Ours-snow+wet	**40.14**	**65.54**	**37.20**	8.41	**41.23**	**69.19**	**40.07**	**11.84**	44.33	71.24	41.49	10.02
Part-A² [42]	*None*	36.59	63.50	30.17	**6.86**	38.03	65.83	35.95	**9.39**	42.81	70.27	39.90	9.18
	Fog [16]	35.98	62.23	30.94	6.39	38.17	66.07	37.12	8.63	41.82	67.44	39.73	9.39
	DROR [6]	35.85	65.36	27.99	6.13	35.43	63.50	32.87	7.95	39.48	66.92	35.18	8.61
	LISA [22]	37.12	65.57	30.05	5.96	38.04	66.62	36.72	8.25	41.92	70.02	38.84	7.79
	Ours-wet	37.13	65.02	30.46	6.14	38.29	65.80	38.12	8.85	41.92	69.12	39.37	9.00
	Ours-snow	**37.73**	**66.44**	29.54	6.46	38.23	**67.72**	36.02	7.41	**43.41**	**71.00**	**40.63**	**9.74**
	Ours-snow+wet	37.52	64.97	**31.65**	6.48	**38.66**	67.00	**38.18**	8.43	42.45	70.27	39.19	8.83
PointRCNN [41]	*None*	36.68	61.74	33.25	6.14	39.04	66.90	39.72	9.28	41.79	**68.58**	**40.34**	7.96
	Fog [16]	36.56	62.93	33.03	5.52	38.37	65.71	39.60	8.50	41.28	67.79	38.82	7.88
	DROR [6]	36.14	62.64	31.64	5.28	36.31	63.52	36.62	7.77	39.08	64.96	36.54	7.70
	LISA [22]	36.68	62.85	31.80	5.78	38.08	65.20	38.96	8.96	**41.80**	68.17	39.79	**8.35**
	Ours-wet	37.07	62.92	**33.76**	**6.38**	**39.46**	66.33	**41.51**	**9.69**	41.70	67.41	39.59	8.05
	Ours-snow	**37.59**	**63.99**	33.63	6.15	38.60	65.18	39.20	9.47	41.43	67.61	39.89	8.15
	Ours-snow+wet	37.51	63.77	33.01	5.90	39.15	**67.10**	39.89	9.43	41.34	67.76	39.58	7.87
SECOND [57]	*None*	36.08	61.53	30.92	6.60	37.77	**65.68**	36.06	9.80	42.10	67.82	39.52	10.81
	Fog [16]	36.08	61.65	31.25	7.65	37.31	64.27	35.43	**10.55**	42.34	**69.85**	39.20	10.13
	DROR [6]	35.04	60.72	28.79	7.88	35.09	62.24	32.09	8.85	38.96	64.74	35.50	9.76
	LISA [22]	35.90	59.31	**32.81**	7.44	**38.07**	64.38	**38.47**	10.32	41.75	67.01	38.24	**11.47**
	Ours-wet	36.79	61.18	32.48	**8.86**	38.04	64.73	37.31	9.95	42.48	69.32	39.28	10.58
	Ours-snow	36.83	61.94	31.40	8.50	37.99	65.40	36.70	9.59	**42.72**	69.16	**40.13**	10.86
	Ours-snow+wet	**37.03**	**64.00**	31.06	7.05	37.11	65.42	35.75	8.69	42.31	69.80	38.25	10.18
PointPillars [25]	*None*	30.85	52.45	27.31	5.59	34.09	59.88	32.80	8.52	38.24	64.02	35.76	8.01
	Fog [16]	30.39	52.13	26.79	5.71	35.38	60.81	35.15	9.60	37.74	64.56	34.48	7.30
	DROR [6]	29.32	**54.52**	21.88	4.82	30.99	57.17	28.43	6.95	34.72	60.59	30.34	6.72
	LISA [22]	28.70	49.78	24.98	5.63	33.87	60.93	31.38	8.70	37.92	63.98	34.61	7.94
	Ours-wet	31.58	52.81	28.64	6.93	34.58	60.57	34.00	8.66	38.10	**65.12**	34.31	7.75
	Ours-snow	**32.94**	54.21	**29.79**	**7.81**	**35.96**	**61.50**	**35.67**	**10.13**	**39.25**	64.37	**36.65**	**8.80**
	Ours-snow+wet	31.38	54.09	27.67	6.14	34.18	60.27	33.12	8.08	38.17	64.40	34.73	7.95

Table 1. Comparison of simulation methods for 3D object detection in snowfall on STF [1]. We report 3D average precision (AP) of moderate cars on three STF splits: the heavy snowfall test split with 1404 samples, the light snowfall test split with 2512 samples and the clear-weather test split with 1816 samples. "Ours-wet": our wet ground simulation, "Ours-snow": our snowfall simulation, "Ours-snow+wet": cascaded application of our snowfall and wet ground simulation.

model (Ours-snow+wet) consistently improves the performance on the most challenging test case, i.e. heavy snowfall, for all methods by a significant margin compared to both the baseline approach as well as all competing simulation [16, 22] and denoising [6] methods. This improvement on heavy snowfall is particularly pronounced for the best-performing detection method, i.e. PV-RCNN [39], for which our full simulation beats the clear-weather baseline by a notable 2.1% in AP.

For PV-RCNN [39], CenterPoint [59] and Part-A² [42], our full simulation also delivers the best performance among all methods on light snowfall, showing that the benefit of our simulation extends to all snowfall intensities. Moreover, on the clear-weather test split, our snowfall simulation without wet ground modeling demonstrates the best performance among all competing approaches for six out of the seven detection methods, consistently improving upon the clear-weather baseline. This finding shows that using

Figure 7. Qualitative comparison of PV-RCNN [39] on samples from STF [1] containing heavy snowfall. The leftmost column shows the corresponding RGB images. The rest of the columns show the LiDAR point clouds with ground-truth boxes and predictions using the clear-weather baseline ("no augmentation"), DROR [6], LISA [22], and our fully- edged simulation ("our snow+wet augmentation").

our snowfall simulation for training increases detection performance on severe snowy conditions, while not sacri cing but rather improving performance on clear weather as well.

Using simulation methods designed for different adverse conditions, such as the fog simulation in [16], does not transfer well to snowfall as the respective physical models differ; performance of [16] is slightly lower than the clear-weather baseline on both snowfall splits for most detection methods. The application of DROR [6] as an enhancement step removing clutter points achieves among the lowest results, because it also removes several valid points, which do not belong to the snowfall clutter.

4.2. Qualitative Results

Qualitative results showing the proposed data augmentation scheme are presented in Fig. 7. Here, PV-RCNN [39] is compared to the clear-weather baseline with no augmentation, DROR [6] and LISA [22]. In the rst row, we see that the pedestrian inside the snowfall clutter can only be detected when our proposed data augmentation is applied during training. In the second row, additional false positives appear for all competing approaches. The bottom row shows a dif cult highway scene with whirled-up snow dust. Our data augmentation approach generalizes well to this example, being the only method that detects the lead vehicle. Note also that in such a scenario with whirled-up snow dust, DROR [6] cannot remove the clutter completely.

5. Conclusion

In this work, we have introduced a novel method for realistic synthesis of winter scenes from clear LiDAR captures modeling snowfall and wet surfaces in a physically accurate way. Further, we have proven the effectiveness of the proposed algorithm, testing the augmentation with seven different 3D object detection methods and achieving consistent improvements of up to 2:1% in AP in heavy snowfall. As future work, we envision the exploration of temporal cues for robust LiDAR-based 3D object detection.

Acknowledgements. This work is funded by Toyota Motor Europe via the research project TRACE-Z rich. The work also received funding by the AI-SEE project with national funding from the Austrian Research Promotion Agency (FFG), Business Finland, Federal Ministry of Education and Research (BMBF) and National Research Council of Canada Industrial Research Assistance Program (NRC-IRAP). We also thank the Federal Ministry for Economic Affairs and Energy for support within "VVM-Veri cation and Validation Methods for Automated Vehicles Level 4 and 5", a PEGASUS family project. Felix Heide was supported by an NSF CAREER Award (2047359), a Sony Young Faculty Award, and a Project X Innovation Award. We thank Emmanouil Sakaridis for verifying our derivation of occlusion angles in our snowfall simulation.

[27] Weixin Lu, Yao Zhou, Guowei Wan, Shenhua Hou, and Shiyu Song. L3-Net: Towards learning based LiDAR localization for autonomous driving. In IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2019. 1

[28] Sebastien Michaud, Jean-François Lalonde, and Philippe Giguere. Towards characterizing the behavior of LiDARs in snowy conditions. In Workshop on Planning, Perception and Navigation for Intelligent Vehicles, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2015. 1, 2

[29] Matthew Pitropov, Danson Evan Garcia, Jason Rebello, Michael Smart, Carlos Wang, Krzysztof Czarnecki, and Steven Waslander. Canadian adverse driving conditions dataset. The International Journal of Robotics Research, 40:681 – 690, 2021. 2, 6

[30] Charles Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas Guibas. Frustum PointNets for 3D object detection from RGB-D data. In IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2018. 2

[31] Charles Qi, Hao Su, Kaichun Mo, and Leonidas Guibas. PointNet: Deep learning on point sets for 3D classification and segmentation. In IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2017. 2

[32] Roy Rasmussen, Jothiram Vivekanandan, Jeffrey Cole, Barry Myers, and Charles Masters. The estimation of snowfall rate using visibility. Applied Meteorology, 38(10):1542 – 1563, 1999. 2

[33] Ralph Rasshofer, Martin Spies, and Hans Spies. Influences of weather phenomena on automotive laser radar systems. Advances in Radio Science, 9, 2011. 2, 3

[34] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In Neural Information Processing Systems (NeurIPS), 2015. 3

[35] Christos Sakaridis, Dengxin Dai, Simon Hecker, and Luc Van Gool. Model adaptation with synthetic and real data for semantic dense foggy scene understanding. In European Conference on Computer Vision (ECCV), 2018. 2

[36] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Semantic foggy scene understanding with synthetic data. International Journal of Computer Vision, 126(9):973–992, 2018. 1, 2

[37] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. ACDC: The Adverse Conditions Dataset with Correspondences for semantic driving scene understanding. IEEE International Conference on Computer Vision (ICCV), 2021. 2

[38] R. S. Sekhon and R. C. Srivastava. Snow size spectra and radar reflectivity. Atmospheric Science, 27(2):299 – 307, 1970. 2

[39] Shaoshuai Shi, Chaoxu Guo, Li Jiang, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. PV-RCNN: Point-voxel feature set abstraction for 3D object detection. In IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2020. 1, 2, 3, 6, 7, 8

[40] Shaoshuai Shi, Li Jiang, Jiajun Deng, Zhe Wang, Chaoxu Guo, Jianping Shi, Xiaogang Wang, and Hongsheng Li. PV-RCNN++: point-voxel feature set abstraction with local vector representation for 3D object detection. arXiv preprint 2102.00463, 2021. 2, 3

[41] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. PointRCNN: 3D object proposal generation and detection from point cloud. In IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2019. 2, 3, 6, 7

[42] Shaoshuai Shi, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. From points to parts: 3D object detection from point cloud with part-aware and part-aggregation network. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), 2020. 2, 6, 7

[43] Zheng Shi, Ethan Tseng, Mario Bijelic, Werner Ritter, and Felix Heide. Zeroscatter: Domain transfer for long distance imaging and vision through scattering media. In IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2021. 2

[44] Jungil Shin, Hyunsuk Park, and Taejung Kim. Characteristics of laser backscattering intensity to detect frozen and wet surfaces on roads. Sensors, 2019:8973248, 2019. 2

[45] Andrea Simonelli, Samuel Rota Bulo, Lorenzo Porzi, Manuel Lopez-Antequera, and Peter Kontschieder. Disentangling monocular 3D object detection. In IEEE International Conference on Computer Vision (ICCV), 2019. 2, 6

[46] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, Vijay Vasudevan, Wei Han, Jiquan Ngiam, Hang Zhao, Aleksei Timofeev, Scott Ettinger, Maxim Krivokon, Amy Gao, Aditya Joshi, Yu Zhang, Jonathon Shlens, Zhifeng Chen, and Dragomir Anguelov. Scalability in perception for autonomous driving: Waymo open dataset. In IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2020. 2

[47] OpenPCDet Development Team. OpenPCDet: An open-source toolbox for 3D object detection from point clouds. https://github.com/open-mmlab/OpenPCDet, 2020. Accessed: 15 November 2021. 6

[48] Maxime Tremblay, Shirsendu S. Halder, Raoul de Charette, and Jean-François Lalonde. Rain rendering for evaluating and improving robustness to bad weather. International Journal of Computer Vision (IJCV), 126, 2021. 1, 2

[49] U.S. Department of Transportation: Federal Highway Administration. How do weather events impact roads? https://ops.fhwa.dot.gov/weather/q1_roadimpact.htm, 2019. Accessed: 15 November 2021. 1

[50] Velodyne. Hdl-64e user's manual, 2007. 3, 5

[51] Sourabh Vora, Alex Lang, Bassam Helou, and Oscar Beijbom. PointPainting: Sequential fusion for 3D object detection. In IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2020. 2

[52] Andrew Wallace, Abderrahim Halimi, and Gerald Buller. Full waveform LiDAR for adverse weather conditions. IEEE Transactions on Vehicular Technology, 69, 2020. 1, 2