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 traffic 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 difficulty 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 significant performance gains on the real snowy STF dataset compared to clear-weather baselines and competing simulation approaches, while not sacrificing 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, reflect

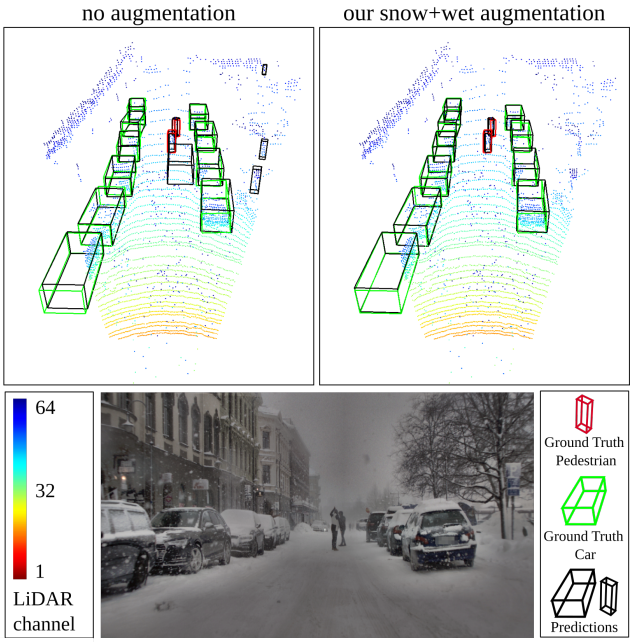


Figure 1. 3D object detection results in heavy snowfall with prior training on the proposed data augmentation scheme (top right) in comparison to no augmentation (top left). The bottom row shows the RGB image as reference.

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].

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].

More recently, simulation methods for LiDAR sensors in fog [1, 16] and rain [14, 22] have also been proposed. Motivated by this line of work, we introduce a physically based method to simulate snowfall on real clear-weather LiDAR point clouds. In particular, we use the linear system introduced in [33] to model the transmission of LiDAR pulses and the associated received power at the sensor. We simulate snowfall by explicitly sampling snow particles and modeling them as opaque spheres, the size of which is controlled by the snowfall rate [15, 32]. In our sampling, we obey the exclusion principle that no two particles intersect with each other. Given the sample of snow particles, we compute for each LiDAR beam the set of particles that intersect with it and derive the angle of the beam cross-section that is reflected by each particle, taking potential occlusions into account. This derivation directly delivers the modified impulse response of the linear system in the presence of snowfall, which allows the analytical calculation of the received power at the sensor.

Another condition associated with snowfall is wetness on the ground. This emerging thin water layer increases the specular component of reflection by the ground surface [44]. To model the ground reflection, we introduce an optical model using the Fresnel equations and the reflection on thin surfaces, which provides adapted reflectance values for wet surfaces.

The generated partially synthetic point clouds with our snowfall and wet ground simulation are used as training data for optimizing state-of-the-art 3D object detection methods, so that the learned models are more robust under snowfall. The hope is that our physically based simulation is realistic enough to relieve us from the need for real snowy training samples. We benchmark the models trained in this regime on the challenging real snowy subset of the STF dataset [1] and find that the models trained on our simulated snow consistently achieve significant performance gains over baseline models trained only on clear weather and competing simulation methods.

2. Related Work

Adverse weather research can be subdivided in meteorological publications providing fundamental knowledge for computer vision approaches [15, 32, 33, 38], phenomenological reasoning of introduced disturbance patterns of different weather conditions [2, 4, 24] and the application of computer vision algorithms to such challenging conditions [1, 29, 37, 48]. In [33] a general theoretical framework predicting the influence of various weather types on LiDAR has been studied, including rain, fog and snow. The authors study these weather effects following the statistical distribution of the scattering particles, which are introduced e.g. for snow in [15] and [38]. Further implications on visibility in snowfall are presented in [32].

Those resulting disturbance patterns and their strength are phenomenologically investigated for rain in [4, 14, 52], fog in [2, 4, 18, 19, 52], snow in [19, 24, 28] and wet surfaces in [5, 44, 56]. Algorithmically, authors have tried to tackle those conditions by robust fusion algorithms [1], developed simulation techniques as data augmentation for camera data in [36, 48] and LiDAR data in [1, 14, 16, 22]. Authors also investigated enhancement technologies to remove adverse weather effects in [6, 17, 43] or applied domain adaptation methods adjusting clear-weather algorithms to adverse weather in [37]. Underlying deep learning data sets containing adverse weather samples were introduced in [1, 3, 29, 37, 46]. Yet classical data loop approaches are difficult to apply as adverse weather samples are rare and well underrepresented [1].

Simulation of adverse weather allows to mitigate the rarity of adverse weather effects and difficulties in data collection campaigns as for example shown in [35, 36, 48] through data synthesis. Additionally, it enables to generate reproducible conditions with clear ground truth necessary to learn image enhancement techniques in [17, 43] or to investigate adverse weather noise-dependent performance decrease reproducible for different weather effects in [48]. LiDAR simulation methods were explicitly studied in [1, 14, 16, 22]. Developments started in [1] with a data driven approach in fog which was extended by [16], introducing a physically based model and achieving higher 3D object detection performances. A simulation method for rain is introduced in [14] and a general approach for snow, rain and fog in [22]. Contrary to [22], our snowfall simulation involves a continuous formulation in the power signal domain, which allows us to superimpose reflections by different particles and to reason about occlusions between particles and the target, thereby adhering better to the physics of the laser transmission. Additionally, we take into account the effects of wet roads, allowing us to estimate laser hardware parameters as noise floor and sent intensity from the dry road intensities in a data-driven way.

3D object detection has seen tremendous progress in recent years. Several methods have been presented for RGB camera [45, 54, 60], LiDAR [25, 31, 39–42, 62], gated imagers [20] or the fusion of multiple modalities in [7, 23, 30, 51]. Across many dataset leaderboards, however, the top performance positions are typically all sorted out among LiDAR based methods [3, 13, 46]. In our work, we utilize the methods PV-RCCN [39], VoxelRCNN-Car [9], CenterPoint [59], Part-A^2 [42], PointRCNN [41], SECOND [57], and PointPillars [25]. The methods differ in point cloud representations, the used feature extraction backbones and the number of detection stages. As input modalities, point clouds are treated *e.g.* in voxel space [39, 40, 42, 57, 62], inferring the raw point clouds [31, 41] or using abstract representations such as pillars in [25].

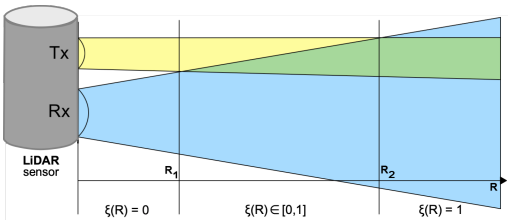


Figure 2. Sketch of a LiDAR sensor where the transmitter Tx and the receiver Rx do not have coaxial optics, but have parallel axes (called a bistatic beam configuration).

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 τ of the captured echo and its corresponding power P_R. The object distance R is calculated by applying $R = c\tau$, 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

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

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

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

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:

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

with the time signature of the transmitted pulse given by

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

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:

$$H(R) = H_C(R)\, H_T(R). \tag{5}$$

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

$$H_C(R) = \frac{T^2(R)}{R^2} \xi(R). \tag{6}$$

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 $\xi(R)$ can be geometrically derived from Fig. 2 as

$$\xi(R) = \begin{cases} 0, & R \le R_1 \\ \frac{R - R_1}{R_2 - R_1}, & R_1 < R < R_2 \\ 1, & R_2 \le R. \end{cases} \tag{7}$$

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

$$H_T(R) = \rho_0 \delta(R - R_0), \tag{8}$$

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.

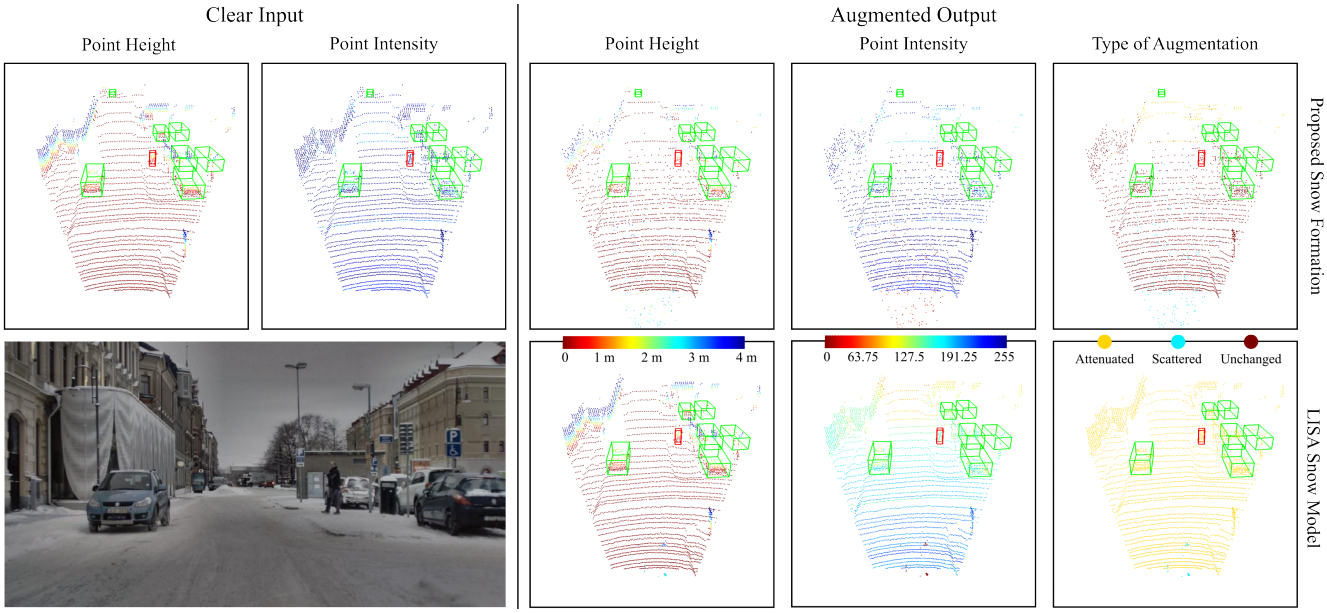


Figure 3. Simulated snowfall corresponding to a snowfall rate of $r_s = 2.5\,\mathrm{mm/h}$. The left block shows the clear undisturbed input. The right block shows our snowfall simulation (top) and the snowfall simulation in LISA [22] (bottom). Note that we simulate the scattering realistically and only attenuate points which are affected by individual snowflakes instead of attenuating all points based on their distance.

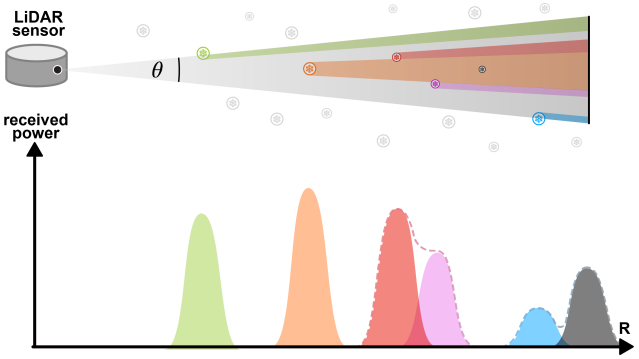


Figure 4. Snow particles interfering a single LiDAR beam (top). Schematic plot of corresponding received power echoes (bottom). Note how the received power of individual targets can overlap with each other ($c\tau_H \approx 3\,\mathrm{m}$ with $\tau_H = 10\,\mathrm{ns}$).

We model snow particle j as a spherical object with reflectivity ρ_s, diameter D_j following the distribution introduced in [15] and distance R_j from the sensor, placed uniformly at random around the sensor so that it does not intersect with any other particle. The number of particles is chosen according to the snowfall rate, typically ranging in 0-2.5 mm/h (see supplementary materials for more details). Particles can occlude each other and the target object, as illustrated in Fig. 4 (top). Thus, each particle j reflects only a fraction θ_j/Θ of the opening angle Θ of the beam, also letting a fraction θ_0/Θ of the beam reach the target. Details on calculating the ratios θ_j/Θ are given in the supplementary materials as well.

Assuming $D_j \ll c\tau_H$ for all j, we can write

$$H_T(R) = \frac{1}{\Theta}\left(\rho_0\,\theta_0\delta(R-R_0) + \rho_s\sum_{j=1}^{n}\theta_j\delta(R-R_j)\right), \quad (9)$$

with $\Theta = \theta_0 + \sum_{j=1}^{n}\theta_j$. Plugging (4), (5), (6) and (9) into (3), the received power in snowfall is

$$P_{R,\mathrm{snow}}(R) = P^0_{R,\mathrm{snow}}(R) + \sum_{j=1}^{n} P^j_{R,\mathrm{snow}}(R), \quad (10)$$

where

$$\begin{aligned} &P^j_{R,\mathrm{snow}}(R) \\ &= \frac{C_A P_0 \rho_s \theta_j \xi(R_j)}{\Theta {R_j}^2}\int_0^{2\tau_H}\sin^2\left(\frac{\pi}{2\tau_H}t\right)\delta(R-\frac{ct}{2}-R_0)dt \\ &= \begin{cases} \frac{C_A P_0 \rho_s \theta_j \xi(R_j)}{\Theta {R_j}^2}\sin^2\left(\frac{\pi(R-R_j)}{c\tau_H}\right), & R_j \le R \le R_j + c\tau_H \\ 0 & \text{otherwise.} \end{cases} \end{aligned} \quad (11)$$

$P^0_{R,\mathrm{snow}}(R)$ can be derived by substituting $(\theta_j,\ R_j,\ \rho_s)$ with $(\theta_0,\ R_0, \rho_0)$ on the right-hand side of (11).

The received power is thus a superposition of multiple echoes, each associated with an object (snow particle or target object), as depicted in Fig. 4 (bottom). Crucially, the magnitude of each echo depends on the angle θ_j and the inverse square of the distance R_j of the respective object from the sensor. In this work, we retrieve the maximum peak of the received power as the LiDAR return. Thus, if a peak owing to snow particles is higher than the peak associated to the target object, the true echo is missed and a cluttered point is added to the simulated point cloud at the range of the former peak.

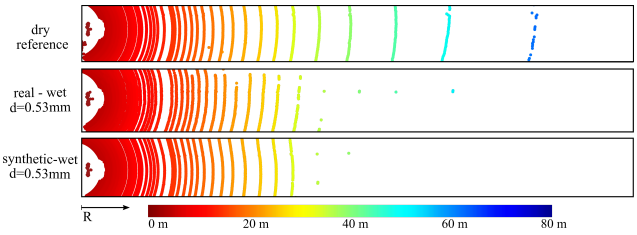


Figure 5. A real-world capture on a dry highway (top), a real-world capture with a water height of $d_w = 0.53\,\text{mm}$ (middle) and the synthesized road wetness from the clear reference (bottom).

Algorithm 1 LiDAR snowfall simulation

procedure SNOWFALL($\mathbf{pc}, n_l, \tau_H, R_{\max}, \Theta, r_s, \rho_s, \rho_0$)
  for l in n_l **do** ▷ for each layer l
    $\mathbf{pc}_l \leftarrow \mathbf{pc}$.SELECT(layer $= l$)
    $f_d, f_s, i_{\max} \leftarrow$ LOAD_CALIB(l)
    $f_o \leftarrow (\frac{1-f_d}{13100})^2$ ▷ focal offset [50]
    $\mathbf{s} \leftarrow$ SAMPLE_SNOWFLAKES($R_{\max}, r_s$) ▷ in 2D [15]
    for $\mathbf{p}$ in $\mathbf{pc}_l$ **do** ▷ for each point in layer l
      $x, y, z, i \leftarrow \mathbf{p}$
      $R_0 \leftarrow \|\mathbf{p}\|_2$
      $\mathbf{t} \leftarrow$ GET_PARTICLES_IN_BEAM($\mathbf{s}, x, y, R_0, \Theta$) ▷ in 2D
      if len($\mathbf{t}$) > 1 **then** ▷ otherwise no interference
        $\mathbf{P}_{R,\text{snow}} \leftarrow \mathbf{0}$ ▷ initialize with zeros
        for i, R, θ in $\mathbf{t}$ **do** ▷ for each target
          if $R = R_0$ **then** ▷ original target
            $P_R \leftarrow i - f_s \left| f_o - \left(1 - \frac{R}{R_{\max}}\right) \right|^2$
            $C_A P_0 \leftarrow \frac{P_R}{\rho_0} {R_0}^2$ ▷ follows from Eq. (11)
          else ▷ snowflake
            $P_R \leftarrow \rho_s\, i_{\max}$
            $C_A P_0 \leftarrow \frac{P_R}{\rho_0}$
          end if
          $\mathbf{P}_{R,\text{snow}}$ += Eq.(11)$(C_A P_0, R, \rho, \tau_H, \theta, \Theta)$
        end for
        $P_R \leftarrow \max(\mathbf{P}_{R,\text{snow}})$
        $R^* \leftarrow \arg\max(\mathbf{P}_{R,\text{snow}}) - c\frac{\tau_H}{2}$
        $i \leftarrow P_R + i_{\max} f_s \left| f_o - \left(1 - \frac{R^*}{R_{\max}}\right) \right|^2$
        $(x, y, z) \leftarrow \frac{R^*}{R_0} \times (x, y, z)$
        $\mathbf{p} \leftarrow x, y, z, i$
      end if
    end for
  end for
  return $\mathbf{pc}$
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 reflected specularly from wet ground leading to significantly 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).

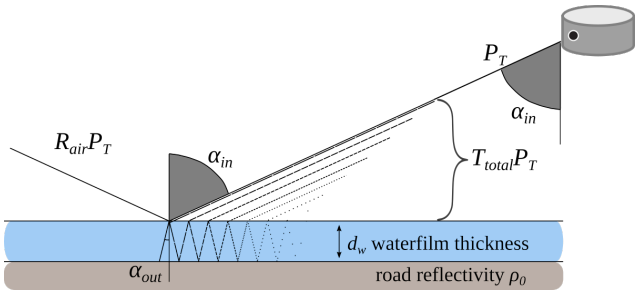


Figure 6. Visualization showing the geometrical optical model which describes the reflection on a wet road surface.

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 $(n_{\text{in}}, n_{\text{out}})$ and angles $(\alpha_{\text{in}}, \alpha_{\text{out}})$ at the transition point. The angle α_{out} can be calculated based on Snell's law:

$$\sin(\alpha_{\text{out}}) = \frac{n_{\text{in}}}{n_{\text{out}}} \sin(\alpha_{\text{in}})\,. \tag{12}$$

The change in pulse amplitude is modeled by the Fresnel equations described by the perpendicular and parallel transmissions $t^{\perp}, t^{=}$ and reflections $r^{\perp}, r^{=}$ with respect to the ground, respectively (details in the supplementary materials). Going from amplitude to transmitted power, we can deduce the power reflection R_k^p and transmission T_k^p:

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

$$T_k^p = \frac{n_{\text{in}} \cos(\alpha_{\text{in}})}{n_{\text{out}} \cos(\alpha_{\text{out}})} (t_k^p)^2\,, \tag{14}$$

where $p \in (\perp, =)$ 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

$$T_{\text{total}}^p = T_{\text{air}}^p \rho_0 T_{\text{water}}^p \sum_{k=0}^{\infty} (\rho_0 R_{\text{water}}^p)^k, \tag{15}$$

which can be simplified for $(\rho_0 R_{water}) < 1$ to

$$T_{\text{total}}^p = T_{\text{air}}^p \rho_0 T_{\text{water}}^p / (1 - \rho_0 R_{\text{water}}^p). \tag{16}$$

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

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

Algorithm 2 LiDAR wet ground simulation

procedure WET_GROUND($\mathbf{p}, i, \mathbf{w}, h, \epsilon_g, d_w, d_p, P_T, i_n$)
  if $|\mathbf{p} \cdot \mathbf{w} - h| < \epsilon_g$ **then**
    $R_0 \leftarrow \|\mathbf{p}\|_2$
    $\alpha_{in} \leftarrow \arccos\left(\frac{\mathbf{p}\cdot\mathbf{w}}{R_0\|\mathbf{w}\|_2}\right)$
    $\rho_0 \leftarrow \frac{i}{\cos(\alpha_{\text{in}})P(R_0)}$ ▷ see Eq. (18)
    $\gamma = \min\left(\max\left(d_w/d_p, 0\right), 1\right)$ ▷ see Eq. (19)
    $\rho_w \leftarrow (1-\gamma)\cdot\rho_0 + \gamma \cdot T_{\text{total}}/\cos\left(\alpha_{\text{in}}\right)$ ▷ see Eq. (20)
    $i \leftarrow \rho_w \cos\left(\alpha_{\text{in}}\right) P_T(R_0)$
    if $i > i_n\left(R_0\right)$ **then**
      return $\mathbf{p}$, i
    else
      return *None*
    end if
  else
    return $\mathbf{p}$, i
  end if
end procedure

Based on this formulation, we design Algorithm 2. A pre-processing step involves estimation of the ground plane normal $\mathbf{w}$ and intercept h with RANSAC [12] as detailed in the supplementary materials. These parameters allow to identify all points belonging to the ground as well as the beam incident angle α_{in}. Then the modified intensity can be reconstructed from the measured "dry" intensity, taking into account that the beam divergence Θ has been corrected by a factory side calibration leading to $i \propto \rho_0 P_T(R)\cos\left(\alpha_{in}\right)$. Normalizing i with the incident angle $i(R)/\cos\left(\alpha_{in}\right)$ returns a linear correspondence between the measured intensities i and the range R, such that we can approximate the power $P_T(R)$ and noise floor $i_n(R)$ linearly from the known given echoes. Then we can obtain the reflectivity ρ_0 for each point $\mathbf{p}$:

$$\rho_0 = i/\left(\cos\left(\alpha_{in}\right)P(R_0)\right). \tag{18}$$

These corrected reflectivities are augmented by weighing dry and wet reflections, assuming a road thread profile of depth d_p filled with water of depth d_w, which yields

$$\gamma\left(d_w, d_p\right) = \min\left(\max\left(d_w/d_p, 0\right), 1\right), \tag{19}$$

$$\rho_w(\gamma) = (1-\gamma)\cdot\rho + \gamma\cdot T_{total}/\cos\left(\alpha_{in}\right). \tag{20}$$

Finally, the measured intensity i is updated based on the modified reflectivity ρ_w and the resulting point $\mathbf{p}$ is only kept if its intensity i is greater than the noise floor i_n.

4. Experiments

Dataset. Our experiments are carried out on the STF dataset [1]. It provides 12997 annotated samples with accurate 3D bounding boxes for object detection of cars, pedestrians, and cyclists in various weather conditions including light fog, dense fog and snow. Without denying the importance of the pedestrian and cyclist classes, in the main paper, we focus on the most dominant class, i.e. cars. In the supplementary materials we provide additional results. In total, 3469 frames in clear conditions can be used for training and 3916 frames in snowy conditions are provided. We split the 3916 samples in the snow test set based on the intensity of snowfall into two different subsets, termed light snowfall and heavy snowfall, with 2512 and 1404 samples respectively. Inspired by [29], we perform this split by leveraging the DROR algorithm [6]. The light snowfall split contains frames where DROR [6] would filter 10-79 points from a $10\times2\times2$m box in front of the ego vehicle, while the heavy snowfall contains at least 80 of such points within this box (further details in the supplementary materials).

Evaluation setting. For evaluation, we use the 3D object detection metrics defined in the KITTI evaluation framework [13] and in [58]. Specifically, [58] introduces an extension to the KITTI metrics by reporting the results with respect to the object distance. Since the weather effects detailed in Sec. 3 are distance-dependent, we opt for following their extension and report results in the intervals as in [20]. Additionally, we follow [45] and report average precision (AP) at 40 recall positions to provide a fair comparison. Other than that we use the typical overlap thresholds defined in [13]. To mitigate potential statistical fluctuations, we report for each experiment the average performance over three independent training runs.

Baseline methods. In total, we investigate the effectiveness of our snowfall simulation scheme for seven well-known 3D object detection methods [9, 25, 39, 41, 42, 57, 59]. We compare our approach to a clear-weather baseline and two competing adverse weather simulation methods, one for fog [16] and one for snowfall [22]. Additionally, we compare to denoising the point clouds using DROR [6]. To train the detection models, we use OpenPCDet [47] and follow the default training configurations for each method. All methods are trained from scratch.

Data augmentation. We choose to apply our simulation(s) to every 10-th training sample, for which the snowfall rate is sampled from [0, 0.5, ..., 2.5]mm/h, and set the sensor constants $\tau_H = 10\,\text{ns}$, $R_{\text{max}} = 120\,\text{m}$, $\Theta = 0.003\,\text{rad}$, $\rho_s = 0.9$ and $\rho_0 = \frac{1\times10^{-6}}{\pi}$. The exact same settings are used for [22]. For the wet ground simulation we use an exponential distribution and sample d_w from the interval 0.1-1.2 mm, while fixing d_p to 1.2 mm and setting ϵ_g to 0.5m.

4.1. Quantitative Results

We present the quantitative results in Table 1. In reading Table 1, the reader should first focus on the columns showing AP across the entire evaluation range of 0-80 m. The main experimental finding from Table 1 is that our full simulation including both the snowfall and the wet ground

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

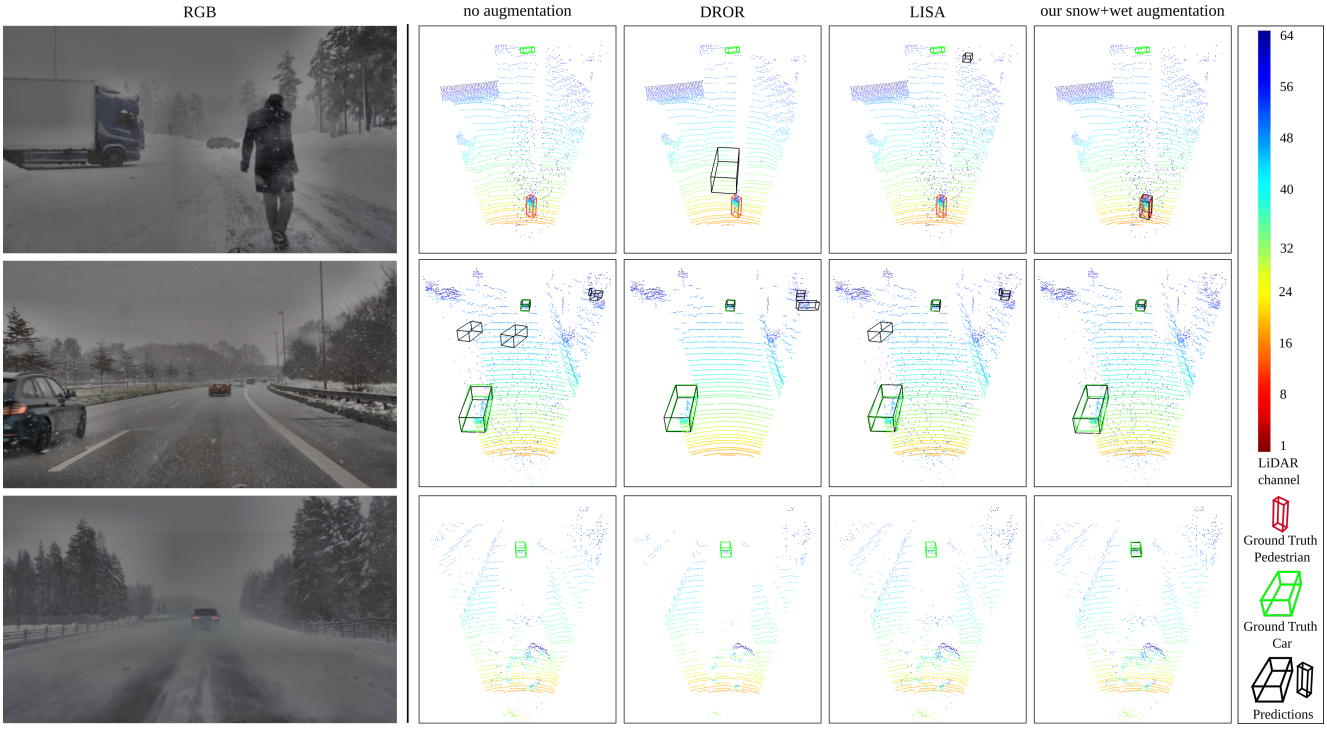


Figure 7. Qualitative comparison of PV-RCNN [39] on samples from STF [1] containing heavy snowfall. The leftmost column shows the corresponding RGB images. The rest of the columns show the LiDAR point clouds with ground-truth boxes and predictions using the clear-weather baseline ("no augmentation"),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

our snowfall simulation for training increases detection performance on severe snowy conditions, while not sacrificing 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 first 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 difficult 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-Zürich. 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-Verification 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.

References

[1] Mario Bijelic, Tobias Gruber, Fahim Mannan, Florian Kraus, Werner Ritter, Klaus Dietmayer, and Felix Heide. Seeing through fog without seeing fog: Deep multimodal sensor fusion in unseen adverse weather. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 2, 5, 6, 7, 8

[2] Mario Bijelic, Tobias Gruber, and Werner Ritter. A benchmark for LiDAR sensors in fog: Is detection breaking down? In *IEEE Intelligent Vehicles Symposium (IV)*, 2018. 1, 2

[3] Holger Caesar, Varun Bankiti, Alex Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuScenes: A multimodal dataset for autonomous driving. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2

[4] Alexander Carballo, Jacob Lambert, Abraham Monrroy Cano, David Robert Wong, Patiphon Narksri, Yuki Kitsukawa, Eijiro Takeuchi, Shinpei Kato, and Kazuya Takeda. LIBRE: The multiple 3D LiDAR dataset. In *IEEE Intelligent Vehicles Symposium (IV)*, 2020. 1, 2

[5] Johan Casselgren, Mikael Sjödahl, and James LeBlanc. Angular spectral response from covered asphalt. *Applied Optics*, 46(20):4277–4288, 2007. 2, 5

[6] Nicholas Charron, Stephen Phillips, and Steven Waslander. De-noising of LiDAR point clouds corrupted by snowfall. In *Conference on Computer and Robot Vision (CRV)*, 2018. 2, 6, 7, 8

[7] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3D object detection network for autonomous driving. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2

[8] European Commission. Road safety in the European Union. https://doi.org/10.2832/169706, 2018. Accessed: 15 November 2021. 1

[9] Jiajun Deng, Shaoshuai Shi, Peiwei Li, Wengang Zhou, Yanyong Zhang, and Houqiang Li. Voxel R-CNN: Towards high performance voxel-based 3D object detection. In *AAAI Conference on Artificial Intelligence*, 2021. 2, 6, 7

[10] Ayush Dewan, Tim Caselitz, Gian Diego Tipaldi, and Wolfram Burgard. Motion-based detection and tracking in 3D LiDAR scans. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2016. 1

[11] Qiong Ding, Wu Chen, Bruce King, Yanxiong Liu, and Guoxiang Liu. Combination of overlap-driven adjustment and phong model for LiDAR intensity correction. *Photogrammetry and Remote Sensing*, 75:40–47, 2013. 3

[12] Martin Fischler and Robert Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24:381–395, 1981. 6

[13] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the KITTI vision benchmark suite. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012. 2, 6

[14] Christopher Goodin, Daniel Carruth, Matthew Doude, and Christopher Hudson. Predicting the influence of rain on LiDAR in ADAS. *Electronics*, 8, 2019. 1, 2

[15] Kevin Gunn and John Marshall. The distribution with size of aggregate snowflakes. *Atmospheric Sciences*, 15(5):452 – 461, 1958. 2, 4, 5

[16] Martin Hahner, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Fog simulation on real LiDAR point clouds for 3D object detection in adverse weather. In *IEEE International Conference on Computer Vision (ICCV)*, 2021. 1, 2, 6, 7, 8

[17] Robin Heinzler, Florian Piewak, Philipp Schindler, and Wilhelm Stork. CNN-based LiDAR point cloud de-noising in adverse weather. *IEEE Robotics and Automation Letters*, 2020. 2

[18] Robin Heinzler, Philipp Schindler, Jürgen Seekircher, Werner Ritter, and Wilhelm Stork. Weather influence and classification with automotive LiDAR sensors. In *IEEE Intelligent Vehicles Symposium (IV)*, 2019. 1, 2

[19] Maria Jokela, Matti Kutila, and Pasi Pyykönen. Testing and validation of automotive point-cloud sensors in adverse weather conditions. *Applied Sciences*, 9, 2019. 1, 2

[20] Frank Julca-Aguilar, Jason Taylor, Mario Bijelic, Fahim Mannan, Ethan Tseng, and Felix Heide. Gated3D: Monocular 3D object detection from temporal illumination cues. In *IEEE International Conference on Computer Vision (ICCV)*, 2021. 2, 6

[21] Alireza Kashani, Michael Olsen, Christopher Parrish, and Nicholas Wilson. A review of LiDAR radiometric processing: From ad hoc intensity correction to rigorous radiometric calibration. *Sensors*, 15(11), 2015. 3

[22] Velat Kilic, Deepti Hegde, Vishwanath Sindagi, A. Brinton Cooper, Mark Foster, and Vishal Patel. LiDAR light scattering augmentation (LISA): physics-based simulation of adverse weather conditions for 3D object detection. *arXiv preprint 2107.07004*, 2021. 2, 4, 5, 6, 7, 8

[23] Jason Ku, Melissa Mozifian, Jungwook Lee, Ali Harakeh, and Steven Waslander. Joint 3D proposal generation and object detection from view aggregation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018. 2

[24] Matti Kutila, Pasi Pyykönen, Maria Jokela, Tobias Gruber, Mario Bijelic, and Werner Ritter. Benchmarking automotive LiDAR performance in arctic conditions. In *IEEE International Conference on Intelligent Transportation Systems (ITSC)*, 2020. 1, 2

[25] Alex Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. PointPillars: Fast encoders for object detection from point clouds. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 3, 6, 7

[26] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C. Berg. SSD: Single shot multibox detector. In *European Conference on Computer Vision (ECCV)*, 2016. 3

[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 Giguère. 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. In *IEEE International Conference on Computer Vision (ICCV)*, 2021. 2

[38] R. S. Sekhon and R. C. Srivastava. Snow size spectra and radar reflectivity. *Atmospheric Sciences*, 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

[53] Dingkang Wang, Connor Watkins, and Huikai Xie. MEMS mirrors for LiDAR: A review. *Micromachines*, 11(5), 2020. 1

[54] Yan Wang, Wei-Lun Chao, Divyansh Garg, Bharath Hariharan, Mark Campbell, and Kilian Weinberger. Pseudo-LiDAR from visual depth estimation: Bridging the gap in 3D object detection for autonomous driving. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2

[55] George Williams. Optimization of eyesafe avalanche photodiode LiDAR for automobile safety and autonomous navigation systems. *Optical Engineering*, 56(3):1 – 9, 2017. 3

[56] Jacek Wojtanowski, Marek Zygmunt, Mirosława Kaszczuk, Zygmunt Mierczyk, and Michał Muzal. Comparison of 905 nm and 1550 nm semiconductor laser rangefinders' performance deterioration due to adverse environmental conditions. *Opto-Electronics Review*, 22, 2014. 2

[57] Yan Yan, Yuxing Mao, and Bo Li. SECOND: Sparsely embedded convolutional detection. *Sensors*, 18, 2018. 2, 3, 6, 7

[58] Bin Yang, Wenjie Luo, and Raquel Urtasun. Pixor: Real-time 3D object detection from point clouds. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 6

[59] Tianwei Yin, Xingyi Zhou, and Philipp Krahenbuhl. Center-based 3D object detection and tracking. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2, 3, 6, 7

[60] Yurong You, Yan Wang, Wei-Lun Chao, Divyansh Garg, Geoff Pleiss, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-LiDAR++: Accurate depth for 3D object detection in autonomous driving. In *International Conference on Learning Representations (ICLR)*, 2020. 2

[61] Ji Zhang and Sanjiv Singh. LOAM : LiDAR odometry and mapping in real-time. *Robotics: Science and Systems Conference (RSS)*, 2014. 1

[62] Yin Zhou and Oncel Tuzel. VoxelNet: End-to-end learning for point cloud based 3D object detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2, 3