Degradation of Lidar-Based 3D Object Detection under the Influence of Artificial Rain

Christoph Rohmann *, Bashar Hoary, Silas Scholz, and Harald Konrad Bachem

Teaching and Research Area Vehicle Safety, Faculty of Automotive Engineering,
Ostfalia University of Applied Sciences, Wolfsburg, Germany
Email: ch.rohmann@ostfalia.de (C.R.); ba.hoary@ostfalia.de (B.H.); sila.scholz@ostfalia.de (S.S.);
h.bachem@ostfalia.de (H.K.B.)
*Corresponding author

***Abstract*—Ensuring robust and safe 3D object detection in outdoor environments requires addressing the challenges posed by adverse conditions such as rain, fog, snow, and varying lighting. However, the scarcity of diverse, labeled training, and test data reflecting these conditions hinders progress in this area. Synthetic data augmentation offers a promising solution to bridge this gap. In this paper, we evaluate the effectiveness of an existing physical rain model for augmenting lidar-based datasets. First, we validate the model by comparing its results to those generated by state-of-the-art simulation software, AURELION. Next, we apply the rain model to augment the KITTI 3D object detection dataset with varying rain intensities and assess the impact on a lidar-based object detection framework. Our results demonstrate that the physical rain model produces outputs nearly identical to AURELION. Furthermore, the augmented data reveal a significant degradation in detection performance across all evaluated object classes under increasing rain intensities. Retraining the detection model with the augmented data set substantially improves its robustness, even under heavy rainfall. These findings highlight the potential of synthetic data augmentation for enhancing the resilience of lidar-based 3D object detection systems in adverse weather conditions.**

***Keywords*—3D object detection, sensor robustness, lidar simulation, data augmentation, adverse weather, rain simulation, autonomous systems**

I. Introduction

Object detection and classification are among the fundamental tasks in the perception pipeline of mobile autonomous machines. Common sensor technologies used for this purpose are optical sensors such as camera and lidar [1]. While these methods are well-studied, their application in outdoor environments still poses a great challenge, as the sensors are exposed to adverse environmental effects such as rain, snow, fog and different lighting conditions. For lidar, these effects can have a substantial negative effect on the resulting data quality [2–4]. The emitted rays in the near-infrared spectrum (usually 905 nm or 1550 nm) are subject to atmospheric scattering and attenuation when interfering with particles in the atmosphere. Rain, which is the focus of this paper, consists of droplets with sizes ranging from 0.1 mm to 6 mm in diameter [5], making them substantially larger than the wavelength of rays emitted by lidar sensors. This is characterized through the size parameter x as shown in Eq. (1), where r is the radius of the particle (here the droplet) and λ the wavelength of the lidar. In this case of $x \gg 1$, the scattering behaviour falls into the domain of geometrical optics (e.g., through ray tracing). However, there have also been adaptations of the original Mie-Theory [6], which is traditionally applied for cases in which $x \approx 1$, to predict the intensity and distribution of scattered lidar rays in rain [7].

$$x = \frac{2\pi r}{\lambda} \tag{1}$$

II. Literature Review

As extensive data sets are required to test lidar performance and train the underlaying model under the influence of rain, various efforts have been made to model said rain behavior to artificially augment pre-existing data sets for validation purposes. A variety of approaches in this field are based on physical models that primarily rely on simplifications of the general lidar equation and experimental studies [8–11]. Other methods use derivations of the Mie-Theory [12–14] to model the scattering behavior of rays traversing through rain. Additionally, there are approaches that utilize complementary simulation methods e.g. for the dynamic splash simulation [15] and photon motion simulation [16]. Teufel *et al*. [17] implemented their own rain model to augment pre-existing data sets and investigated the degradation in object detection at different intensities using the PointPillars detection model [18]. While these investigations give insight into the general applicability of such models, there is a research gap in how they fare against contemporary simulation tools for autonomous mobile machines and how they affect the object detection task when being used for the training process of detection models.

In this paper, we apply the rain model proposed by Goodin *et al*. [9] to augment the KITTI 3D Object

Manuscript received February 12, 2025; revised March 20, 2025; accepted April 27, 2025; published July 25, 2025.

Detection data set [19] with artificial rain of varying intensities. The goal is to investigate its applicability to create synthetic training and test data for 3D object detection models and how the performance of these models holds up when being exposed to artificial rain effects. We specifically chose this rain model, as it is not computationally demanding, thus making it suitable for both augmentation of pre-existing data as well as real-time usage. Furthermore, unlike other of the models mentioned, it only requires the desired rain intensity as an input, which makes it also applicable to recreate rain conditions based on actual historic data, which is usually represented in intensities of mm/h.

For this purpose, we first compare the rain noise with that generated from a simulation software called AURELION to analyze the rain model's general behavior. We then investigate the degradation behavior of an object detection model when subjected to inputs with the augmented KITTI data sets of varying rain intensities. Following the results from the analysis, we retrain the underlaying model with a mixture of original and noise-applied data to compare the resulting performance with that of the original model.

III. Rain Model Implementation and Analysis

A. Physics-Based Rain Model Integration

Goodin *et al*. [9] utilize a simplified variant of the general lidar equation to calculate the return signal power of each ray (represented as a point in the point cloud) as a function of the distance R, the reflectivity ρ and the rain intensity I as in Eq. (2).

$$P(R) = \frac{\rho}{R^2} \cdot e^{-0.02I^{0.6}R} \tag{2}$$

The factors used in this equation are based on real outdoor experiments conducted with a Velodyne VLP-16 sensor under rain intensities of up to 8 mm/h by Filgueira *et al*. [20]. As shown in Fig. 1, the return signal power of each point is compared to the minimum power threshold P_{min} of the lidar sensor, which is characterized by its maximum range, usually stated in combination with the minimum required Lambertian reflection (e.g., 120 m@0.80% for Velodyne HDL-64E). The points are then subjected to normally distributed noise as shown in Eq. (3).

$$R' = R + N(0,0.02R(1 - e^{-I})^2) \tag{3}$$

If the return signal of a given point is below the threshold P_{min} given by the sensor, it is removed from the point cloud. All the remaining points are added to the new rain-applied output cloud. Our implementation allows both for real-time operation using the Robot Operating System (ROS) with input clouds in the PointCloud2 format and augmentation of common lidar formats such as pcd-files. Additionally, there is a plugin to request historic or forecasted rain intensities through the Open-Meteo API [21] by inputting a desired location and time. This feature is especially useful for obtaining realistic rain intensity values for a given location.

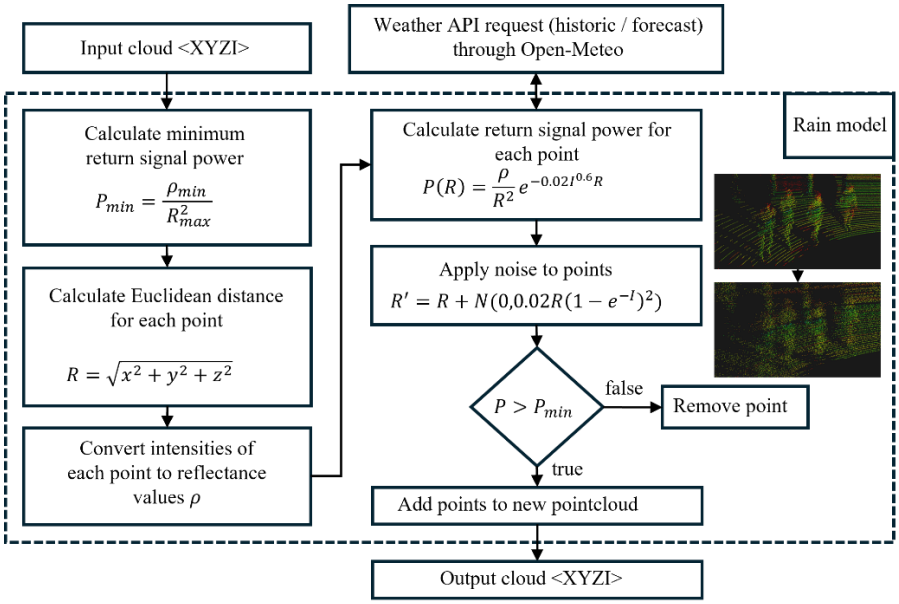

Fig. 1. The principle of the data augmentation with artificial rain.

There is no mutual international classification of rain and their corresponding intensities. In this work, we use the classifications as given by the World Meteorological Organization (WMO) [22] and the UK Meteorological Office [23]. They differentiate between rain and showers. Showers are characterized by a short duration and rapid fluctuations of intensity. Rain on the other hand has a more constant characteristic and usually lasts for hours. The corresponding intensities of slight, moderate, heavy and violent rain and showers respectively are listed in Table I.

To define the maximum hourly rainfall intensity for this study, we analyzed historic meteorological precipitation data from 2016 to 2024 [21] from the city of Toyohashi in Japan, where most of this research was conducted. The highest recorded hourly intensity of rain in this period was 38.8 mm/h in 2020, while the average annual maximum

was 31.38 mm/h. In 2023, the maximum daily rainfall reached 200 mm/day with a peak hourly intensity of 37.1 mm/h. As this event resulted in major flooding [24] and taking into account previous maximum intensities, we set the upper limit for our investigations at a realistic maximum rainfall intensity of 50 mm/h.

TABLE I. RAIN CLASSIFICATIONS ACCORDING TO WMO

Class	Rain intensity (mm/h)	Shower intensity (mm/h)
Slight	<0.5	<2
Moderate	0.5, …, 4	2, …, 10
Heavy	>4	10, …, 50
Violent	-/-	>50

B. *Simulated Rain Noise*

As the suggested rain model is based on experimental data collected with a Velodyne VLP-16 sensor under intensities of up to 8 mm/h, we validate its noise behavior at intensities beyond that level through simulations done in the simulation software AURELION. AURELION is software based on Unreal Engine developed by the company dSPACE that allows sensor-realistic simulation of lidar sensors. It is mainly used for the development of automated driving systems and has a built-in rain simulation that affects the lidar data output.

AURELION comes with a library of Velodyne sensors, including the HDL-64 lidar used in the KITTI data set. We chose a common urban 3D scene (Fig. 2) as testing ground with diverse objects of varying intensity (e.g., traffic signs and number plates) and captured a single frame of the simulated point cloud, once without any rain and additionally with the integrated rain simulation toggled on. We compare intensities of 1 mm/h to 20 mm/h, since 20 mm/h is the maximum supported rain intensity in AURELION. The frame captured without any rain simulation is subjected to the rain model.

We compare the resulting point clouds with respect to effective range, point count, and spatial distribution. While AURELION simulates rainfall up to 20 mm/h, the model extends to 50 mm/h to reflect historical precipitation extremes. To validate the model, we focused on the overlapping intensity range (0–20 mm/h), where our method demonstrated strong agreement with AURELION in terms of point cloud degradation patterns. This confirms that the model's core mechanisms—signal-dependent point removal and Gaussian noise—align with an established simulation framework under equivalent conditions. Although higher intensities cannot be directly validated due to software limitations, the model's behavior remains plausible due to its physical grounding and consistent trends observed at lower intensities. As shown in Fig. 3, both the number of points and the maximum range degrade equally with both methods.

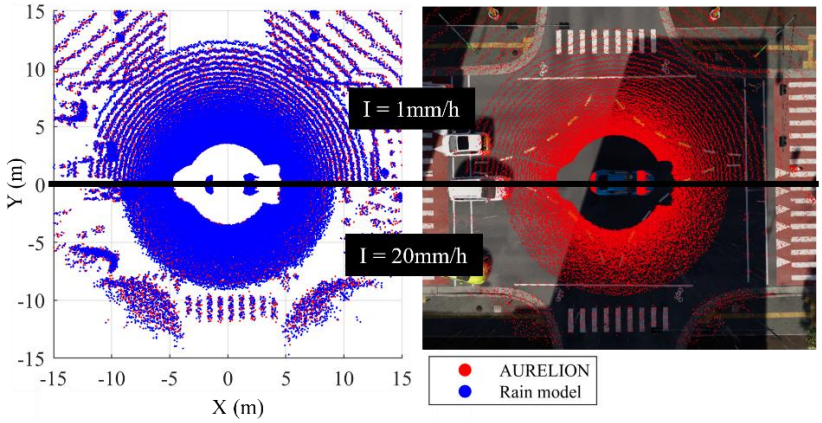

Fig. 2. (Left) Comparison of the point distribution between the rain model and the simulated rain in AURELION. (Right) Simulated scene and lidar data in AURELION.

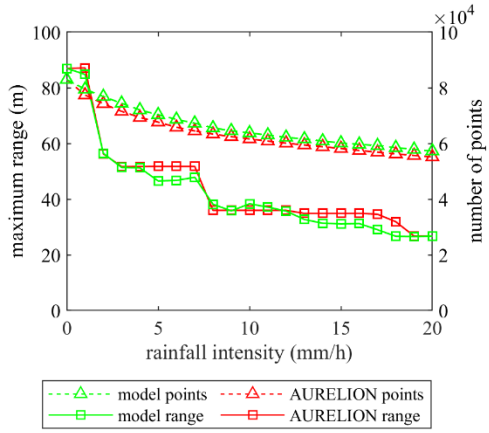

Fig. 3. Comparison of the number of points and the maximum range between therain model and AURELION.

In addition to assessing the maximum effective range and point count, we evaluated the global similarity of both point clouds using the Chamfer Distance (CD) and Earth Mover's Distance (EMD) metrics. For both metrics, we used the implementations given in [25].

The Chamfer Distance (CD) quantifies similarity by computing the Nearest Neighbor (NN) in point cloud B for each point in cloud A, and vice versa. The sum of these distances provides a measure of similarity as shown in Eq. (4).

$$CD(A,B) = \frac{1}{2n}\sum_{i=1}^{n}|x_i - NN(x_i,B)| + \frac{1}{2m}\sum_{j=1}^{n}|x_j - NN(x_j,A)| \quad (4)$$

As rain intensity increases, the point cloud density rises, leading to shorter distances between corresponding points and thus a lower CD value. For an intensity I = 1 mm/h, the CD is 0.111 m, decreasing to 0.105 m at I = 20 mm/h, with an average CD of 0.114 m across all intensities. Since CD does not explicitly account for spatial density variations, we also employed EMD to capture these effects.

The Earth Mover's Distance (EMD) considers the spatial distribution of points by measuring the minimum cost required to transform one point cloud into the other as described in Eq. (5).

$$EMD(A,B) = \min_{\pi \in \Pi(A,B)} \sum_{i=1}^{n} \sum_{j=1}^{m} \pi_{i,j} |a_i - b_j| \quad (5)$$

Similar to CD, EMD decreases with increasing rain intensity, ranging from 0.195 m at I = 1 mm/h to 0.125 m at I = 20 mm/h, with an average of 0.160 m. Considering both absolute distances and point cloud size, these metrics indicate a high degree of similarity between the two point clouds.

IV. Object Detection Setup

A. 3D Object Detection Framework

We use the OpenPCDet [26] framework for the object detection task, as it offers a variety of built-in testing and training tools as well pre-trained detection models. We opted for the Part-A^2 net model [27], as it provides the best classification results on the KITTI test split for Vulnerable Road Users (VRU), consisting of pedestrians and cyclists. We put special emphasis on these targets, as their detection and classification are generally more challenging for the underlying models and interactions with VRU pose a particularly high level of safety criticality. We also tested the more recent voxel-based model VoxelNeXt [28], which in comparison provided worse results as shown in Table II. The Average Precision (AP) is calculated over 40 recall positions at the medium difficulty KITTI benchmark for the test split.

TABLE II. Object Detection Performance on KITTI 3D Part-A^2 vs. VoxelNeXt

Model	AP_{Ped} (%)	AP_{Cycle} (%)	AP_{Car} (%)
Part A^2-net	66.40	75.33	80.31
VoxelNeXt	54.36	64.88	80.25

B. Model and Data Preparation

We input the KITTI 3D Object Detection test split, consisting of 3769 point cloud samples, into the rain model with rain intensities of I = {1, 5, 10, 20, 50} mm/h resulting in a total of five noise-applied test datasets.

As shown in Fig. 4, we run these noise-applied data sets on a pre-trained version of the Part-A^2 model and two retrained versions. The pre-trained model is trained on the unmodified (raw) training split of the KITTI 3D Object Detection set. For the first retrained model, we input the training split into our rain model with the rain intensity set to I = 5 mm/h. The resulting noise-applied data set is used to train the model ($retrained_{noise}$). For the second version ($retrained_{mixed}$), we use a mixture of both raw KITTI data as well as the noise-applied modification with I = 5 mm/h. We opted for a ratio of 25% raw and 75% noise-applied data, as it provided the best results. Higher percentages of raw data result in reduced performance under rain.

Each noise-applied test data set is used as an input for the OpenPCDet built-in evaluation tool, giving both a visual output of the interference (qualitative) and the average precision for the detection performance (quantitative).

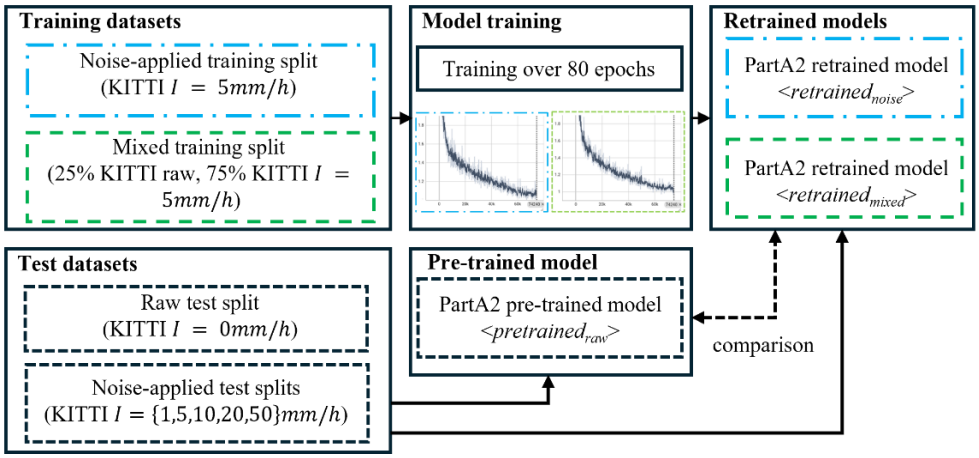

Fig. 4. Data and detection model preparation based on raw and noise-applied training and test data.

V. Detection Results

A. Qualitative Evaluation

For the qualitative comparison, we plot the interference results of the detection models in a Birds-Eye View (BEV) perspective. We chose a point cloud frame from the test split (000017.bin) in which a lot of pedestrians are present, as this object class is the focus of this evaluation.

Fig. 5 shows the results from the pre-trained model (top) and the model that was trained on the noise-only data set (bottom). Looking at the point clouds resulting from applying artificial rain of increasing intensity, it becomes clear that heavy rain starting from rain intensities of I =

5 mm/h onwards, the noise behavior is visually identical. However, there is still a noticeable degradation in the effective range from I = 5 mm/h up to I = 50 mm/h. With the pre-trained model, object detection provides good results when no noise is introduced to the input cloud, but quickly falls apart when exposed to rain intensities of I = 5 mm/h and higher. This is due to the heavy amount of noise introduced to the input cloud, which makes the usually clearly defined layers of the cloud indistinguishable from one another. As the pre-trained model was trained solely on noiseless data, it has no reference on how to work with inputs that have such a fundamental difference in structure.

The model that was trained exclusively on noise-applied data ($\text{retrained}_{\text{noise}}$) shows a significant increase in performance. Looking at the inputs ranging from I = 1 mm/h (moderate rain) to I = 50 mm/h (heavy shower), the model still detects pedestrians consistently, especially when compared to the pretrained model. The main influencing aspect of the number of detected pedestrians in the frame is the effective range of the point cloud that is significantly shorter at I = 50 mm/h than with the raw data as input. However, since the model was trained on noise-applied data exclusively, the performance at I = 0 mm/h (noise-free data) is worse compared to that of the pretrained model. There are several false positive detections of objects that are not present in the original frame. Fig. 6 shows a close-up of this scene as a comparison between the retrained models $\text{retrained}_{\text{noise}}$ (left) and $\text{retrained}_{\text{mixed}}$ (right). While the former shows said behavior of false positive detections, the model trained with a mixed set of raw and noise-applied data behaves visually identically to the pre-trained model. For inputs with rain intensities of I = 1 mm/h and higher, the mixed model also performs better, which becomes evident in the quantitative evaluation.

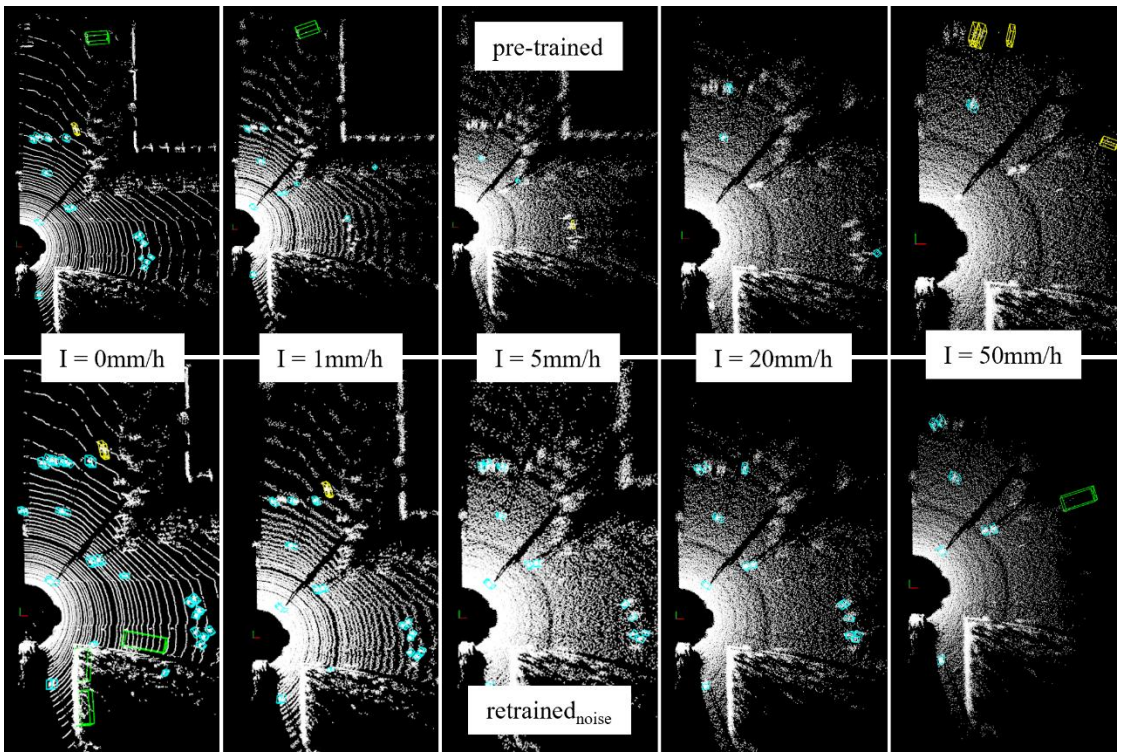

Fig. 5. Qualitative comparison of the PartA2 detection model on test data with increasing rain intensities. Teal: pedestrians; yellow: cyclists; green: cars.

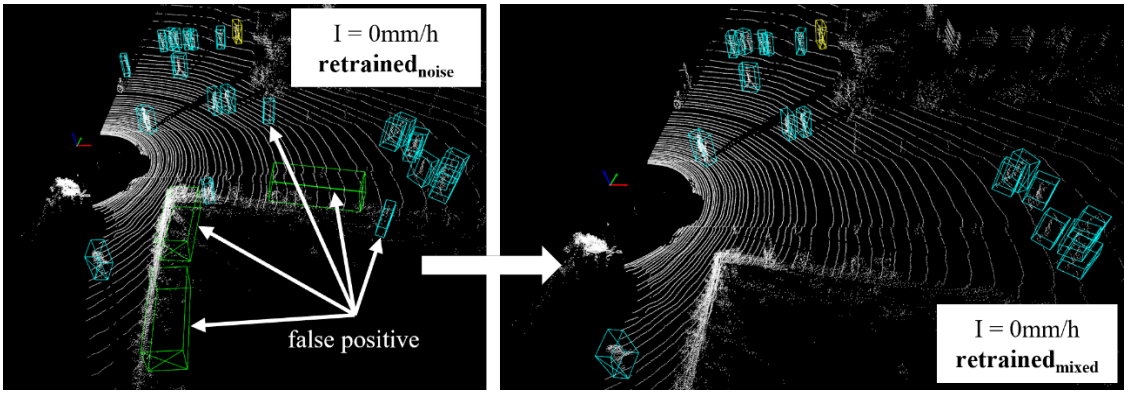

Fig. 6. (Left) False positive detections of the model exclusively trained on noise-applied data. (Right) Improved detection performance by using the model with mixed training data.

In summary, these qualitative results indicate that incorporating both augmented and non-augmented data during training is an effective strategy to prevent false-positive detections in noise-free input data while preserving high accuracy when applied on noise-augmented data.

B. Quantitative Evaluation

For the quantitative evaluation, we separately input the raw and the noise-applied test splits into each of the three detection models and compare their Average Precision (AP). The AP is based on 40 recall positions at medium difficulty in compliance with the KITTI detection benchmark. Cars require a bounding box overlap of 70% and pedestrians and cyclists of 50% to be recognized as detected objects.

Fig. 7 and Table III show the AP of the pre-trained model $\text{pretrained}_{\text{raw}}$ and the retrained models $\text{retrained}_{\text{noise}}$ and $\text{retrained}_{\text{mixed}}$ under varying rain intensities.

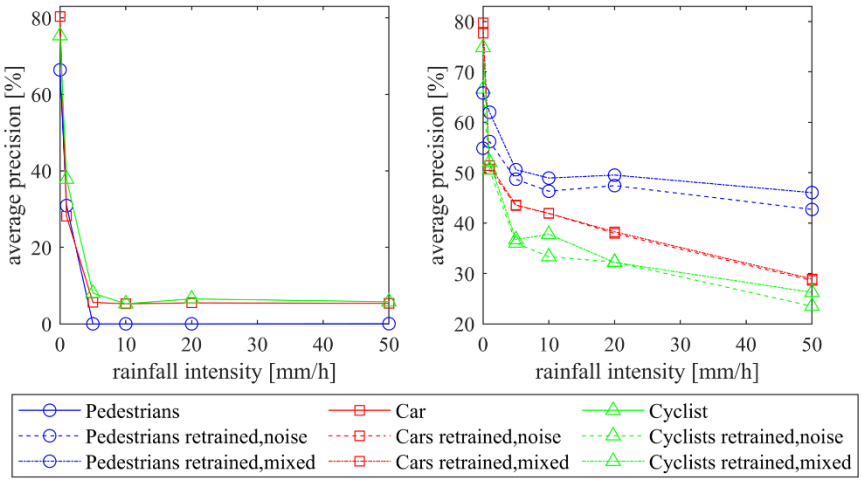

Fig. 7. (Left) Performance of the pre-trained detection model. (Right) Performance of the two retrained models.

TABLE III. COMPARISON OF THE AP BETWEEN THE PRE-TRAINED AND THE TWO RETRAINED DETECTION MODELS ON INPUT DATA WITH INCREASING RAIN INTENSITIES. VALUES REFLECT AP@R40 MEDIUM DIFFICULTY FROM THE KITTI BENCHMARK METRIC

I (mm/h)	$\text{pretrained}_{\text{raw}}$ (%)			$\text{retrained}_{\text{noise}}$ (%)			$\text{retrained}_{\text{mixed}}$ (%)		
	AP_{Ped}	AP_{Car}	AP_{Cycle}	AP_{Ped}	AP_{Car}	AP_{Cycle}	AP_{Ped}	AP_{Car}	AP_{Cycle}
0	66.40	80.31	75.33	54.86	77.65	66.60	65.80	79.73	74.84
1	30.92	28.13	37.84	56.18	50.66	50.66	61.99	51.42	52.09
5	0.0097	5.62	8.02	48.70	43.45	36.03	50.60	43.57	36.73
10	0.0092	5.25	5.27	46.33	41.95	33.34	48.94	41.91	37.79
20	0.0086	5.48	6.53	47.44	37.94	32.23	49.52	38.28	32.19
50	0.43	5.37	5.79	42.71	28.61	23.51	46.04	28.93	26.26

As seen in the qualitative evaluation, the performance of the pre-trained model degrades significantly when exposed to rain, resulting in an AP of 30.9% at $I = 1$ mm/h and 0.01% at $I = 5$ mm/h for the pedestrian detection task. Looking at the AP of the $\text{retrained}_{\text{noise}}$ model, it reflects the behavior seen in the qualitative evaluation. At no rain present in the input clouds, the AP is 11.5% lower in the retrained model. When rain is introduced to the input clouds however, the retrained model shows a significant performance increase, reaching a maximum of $\Delta AP = 48.7\%$ at $I = 5$ mm/h. In comparison, the $\text{retrained}_{\text{mixed}}$ model performs even better with $\Delta AP = 50.6\%$ at $I = 5$ mm/h. Additionally, its performance at $I = 0$ mm/h almost reaches that of the pre-trained model, resulting in a deviation of $\Delta AP = -0.6\%$ for the pedestrian detection task.

When subjected to rain noise, the pedestrian detection with the retrained models holds up better at every investigated intensity compared to that of the cars and cyclists, maintaining an AP of around 50%. Pedestrians have a more uniform and predictable shape compared to cars and cyclists, which can have varying shapes, sizes and orientations. As the applied noise algorithm distorts these usually distinct characteristics, the detection of cars and cyclists becomes more challenging.

VI. RESULT AND DISCUSSION

Initial comparisons between state-of-the-art simulation tools and the applied principle for the creation of artificial rain effects on lidar showed that the underlaying physical model is generally feasible for creating synthetic test and training data. When such artificial rain effects are introduced to the input data of 3D object detection models, rain intensities of 5mm/h and higher (heavy rain) cause a significant degradation in detection performance, making the model virtually unusable under these conditions. By retraining the model with a data set augmented with noise representing rain intensities of 5 mm/h exclusively, the detection results became significantly better, maintaining an AP of 48.6% and 42.7% for pedestrian detection under 5 mm/h and 50 mm/h rain intensities, respectively. False-positive detections occurred when using the retrained

model on the raw data set with no noise applied, which was corrected by performing a second training iteration, this time using a combination of both raw and noise-introduced data as input for the training process.

VII. Conclusion

This study investigated the impact of artificial rain on lidar-based 3D object detection and evaluated the effectiveness of a physical rain model for data augmentation. By comparing the model's output with state-of-the-art simulation software, we confirmed its validity in generating realistic rain-induced noise. Our findings demonstrate that increasing rain intensity leads to a significant degradation in object detection performance. However, retraining the detection model with a combination of raw and augmented data substantially improved its robustness, mitigating the negative effects of rain. These results highlight the importance of incorporating diverse environmental conditions in training datasets to enhance the reliability of lidar-based perception systems. Future research should focus on further validating synthetic augmentation methods through real-world experiments and extending the approach to other adverse weather conditions.

Conflict of Interest

The authors declare no conflict of interest.

Author Contributions

C.R. conducted the research, developed the methodology, and wrote the manuscript. B.H. and S.S. contributed equally by performing the simulation experiments in AURELION and generating the corresponding diagrams. H.K.B. provided thematic guidance, reviewed the manuscript, and contributed to refining the research focus and presentation of results. All authors reviewed and approved the final version of the paper.

Funding

This work was supported by the Japan Society for the Promotion of Science (JSPS) through the Pre/Postdoctoral Fellowships for Research in Japan—JSPS Summer Program 2024. Additionally, part of this research was conducted within the scope of the GESAL project, funded by the European Regional Development Fund (ERDF) under grant number 87012037.

Acknowledgment

The authors wish to express their sincere gratitude to Professor Jun Miura from the Toyohashi University of Technology for his thematic advice that greatly contributed to this work.

References

[1] S. Y. Alaba, A. C. Gurbuz, and J. E. Ball, "Emerging trends in autonomous vehicle perception: Multimodal fusion for 3D object detection," *World Electric Vehicle Journal (WEVJ)*, vol. 15, no. 20, 2024.

[2] M. Dreissig, D. Scheuble, F. Piewak, and J. Boedecker, "Survey on LiDAR perception in adverse weather conditions," in *Proc. 2023 IEEE Intelligent Vehicles Symposium (IV)*, 2023.

[3] Y. Zhang, A. Carballo, H. Yang, and K. Takeda, "Perception and sensing for autonomous vehicles under adverse weather conditions: A survey," *ISPRS Journal of Photogrammetry and Remote Sensing* vol. 196, no. 146, 2023.

[4] T. Yang, J. Hu, Y. Li, C. Zhao *et al*., "3D ToF LiDAR for mobile robotics in harsh environments: A review" *Unmanned Systems*, vol. 1, 2024.

[5] A.C. Best, "The size distribution of raindrops," *Quart J. Royal Meteoro Soc.*, vol. 76, no. 16, 1950.

[6] G. Mie, "Contributions to the optics of turbid media, especially colloidal metal solutions," *Annalen der Physik*, pp. 330–377, 1908. https://doi.org/10.1002/andp.19083300302 (in German)

[7] M. Byeon and S. W. Yoon, "Analysis of automotive lidar sensor model considering scattering effects in regional rain environments," *IEEE Access*, vol. 8, 102669, 2020.

[8] P. A. Lewandowski, W. E. Eichinger, A. Kruger, and W. F. Krajewski, "LiDAR-based estimation of small-scale rainfall: Empirical evidence," *Journal of Atmospheric and Oceanic Technology*, vol. 26, 656, 2009.

[9] C. Goodin, D. Carruth, M. Doude, and C. Hudson, "Predicting the influence of rain on LIDAR in ADAS," *Electronics*, vol. 8, 89, 2019.

[10] J. Zhao, Y. Li, B. Zhu, W. Deng, and B. Sun, "Method and applications of Lidar modeling for virtual testing of intelligent vehicles," *IEEE Trans. Intell. Transport. Syst*., vol. 22, 2990, 2021.

[11] J. P. Espineira, J. Robinson, J. Groenewald, P. H. Chan, and V. Donzella, "Realistic LiDAR with noise model for real-time testing of automated vehicles in a virtual environment," *IEEE Sensors J.*, vol. 21, 9919, 2021.

[12] M. Ballesta-Garcia, G. DeMas-Giménez, and S. Royo, "Modeling the use of LiDAR through adverse weather," *GIS and Spatial Analysis*, 2023.

[13] A. Haider *et al*, "A Methodology to model the rain and fog effect on the performance of automotive LiDAR sensors", *Sensors*, vol. 23, 2023.

[14] V. Kilic, D. Hegde, V. Sindagi, A. B. Cooper, M. A. Foster, and V. M. Patel, "Lidar Light Scattering Augmentation (LISA): Physics-based simulation of adverse weather conditions for 3D object detection," in *Proc. 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021.

[15] X. Huang, H. Wu, X. Li, X. Fan, C. Wen, and C. Wang, "Sunshine to rainstorm: Cross-weather knowledge distillation for robust 3D object detection," in *Proc. 2024 AAAI Conference on Artificial Intelligence*, 2024.

[16] R. Yu, X. Li, and T. Bi, "Modelling and validation of LiDAR noise distribution in fog and rain," *Measurement*, vol. 229, 114472, 2024.

[17] S. Teufel, G. Volk, A. von Bernuth, and O. Bringmann, "Simulating realistic rain, snow, and fog variations for comprehensive performance characterization of LiDAR perception," in *Proc. 2022 IEEE 95th Vehicular Technology Conference: (VTC2022-Spring) (IEEE)*, 2022.

[18] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, "PointPillars: Fast encoders for object detection from point clouds," in *Proc. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

[19] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? The KITTI vision benchmark suite," in *Proc. 2012 IEEE Conference on Computer Vision and Pattern Recognition (IEEE)*, 2012.

[20] A. Filgueira, H. González-Jorge, S. Lagüela, L. Díaz-Vilariño, and P. Arias, "Quantifying the influence of rain in LiDAR performance," *Measurement*, vol. 95, 143, 2017.

[21] P. Zippenfenig. (2024). Open-Meteo.com Weather API. Zenodo. [Online]. Available: https://zenodo.org/records/14582479

[22] World Meteorological Organization. (2025). Aviation | Hazards | Precipitation. [Online]. Available: https://community.wmo.int/en/activity-areas/aviation/hazards/precipitation

[23] Met Office. National Meteorological Library and Archive Water., Fact Sheet 3—Water in the Atmosphere. [Online]. Available: https://www.metoffice.gov.uk/binaries/content/assets/metofficegovuk/pdf/research/library-and-archive/library/publications/factsheets/factsheet_3-water-in-the-atmosphere_2023.pdf

[24] The Asahi Shimbun. Water at Thigh Level; Mother, Son Saved from Submerged Car. [Online]. Available: https://www.asahi.com/ajw/articles/14924168
[25] F. Williams. (2022). Point Cloud Utils. [Online]. Available: https://github.com/fwilliams/point-cloud-utils
[26] OpenPCDet Development Team. (2020). OpenPCDet: An Open-source Toolbox for 3D Object Detection from Point Clouds. [Online]. Available: https://github.com/open-mmlab/OpenPCDet
[27] S. Shi, Z. Wang, J. Shi, X. Wang, and H. 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*, 2019, pp. 2647–2664.
[28] Y. Chen Y, J. Liu, X. Zhang, X. Qi, and J. Jia, "VoxelNeXt: Fully sparse VoxelNet for 3D object detection and tracking," in *Proc. 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.