

Enhancing Robustness of LiDAR-Based Perception in Adverse Weather using Point Cloud Augmentations

Sven Teufel, Jörg Gamberdinger, Georg Volk, Christoph Gerum and Oliver Bringmann
{sven.teufel, joerg.gamberdinger, georg.volk, christoph.gerum, oliver.bringmann}@uni-tuebingen.de
Department of Computer Science
University of Tübingen
Tübingen, Germany

Abstract—LiDAR-based perception systems have become widely adopted in autonomous vehicles. However, their performance can be severely degraded in adverse weather conditions, such as rain, snow or fog. To address this challenge, we propose a method for improving the robustness of LiDAR-based perception in adverse weather, using data augmentation techniques on point clouds. We use novel as well as established data augmentation techniques, such as realistic weather simulations, to provide a wide variety of training data for LiDAR-based object detectors. The performance of the state-of-the-art detector Voxel R-CNN using the proposed augmentation techniques is evaluated on a data set of real-world point clouds collected in adverse weather conditions. The achieved improvements in average precision (AP) are 4.00 *p.p.* in fog, 3.35 *p.p.* in snow, and 4.87 *p.p.* in rain at moderate difficulty. Our results suggest that data augmentations on point clouds are an effective way to improve the robustness of LiDAR-based object detection in adverse weather.

I. INTRODUCTION

Since autonomous vehicles (AVs) promise to make future mobility more comfortable and efficient, autonomous driving is of tremendous interest to both society and current research. Additionally, a major benefit of autonomous driving is an increase in traffic safety. As presented by the European Union, 90% of accidents are due to human error including factors such as alcohol, drugs or distraction [1]. Therefore, utilizing AVs to replace human drivers could reduce the number of accidents.

An AV must be able to perceive its surroundings completely and comprehensively at all times in order to genuinely achieve the safety enhancement. Therefore, the sensors and algorithms must be able to handle a wide range of external influences or disturbances. The ISO26262 recommends a testing scope for AVs of 10^9 km to verify the sensors and algorithms. However, even with such extensive testing, not all scenarios are necessarily covered and cannot be taken into account for neural network-based object detection training and verification. Since the robustness of large networks cannot be formally verified [2], the robustness verification must be carried out using a testing procedure with as much input as possible regarding scenarios, environmental influences, and potential corruptions due to sensor degradation and failures. In order to improve

the robustness of object detection a retraining of object detectors on an extended data set containing the aforementioned influences can be performed [3]. Therefore, we propose a set of data augmentation techniques integrated in an optimization flow operating on LiDAR point clouds. This method creates a comprehensive and manifold data set to train state-of-the-art object detectors and increase their robustness. Numerous augmentation techniques, including geometric modifications such as scaling or transformations, intensity shifts, and augmentations through a physically realistic weather simulation, are considered and evaluated.

An overview of our approach is given in Figure 1. The object detector is embedded in an optimization loop and trained using a set of augmentation techniques on clear weather recordings. Then the detector is evaluated on recordings of real fog, snow, and rain. To optimize the augmentation parameters the Optuna hyperparameter optimization framework [4] is used. In Section II, we investigate current state-of-the-art research related to our proposed methods. The proposed LiDAR point cloud augmentation strategies are presented in Section III. Afterwards, we explain our conducted experiments in Section IV and the corresponding results in Section V. Finally, in Section VI we conclude our work and give an outlook to further research topics.

II. RELATED WORK

Hahner et al. [5] evaluated the impact of different data augmentations on the detection performance of an object detector on the KITTI data set [6]. For this, they trained the object detector PointPillars [7] with different augmentation techniques separately as well as combinations of the augmentations which achieved the best results. They showed, that point cloud augmentations can significantly improve the detection performance of a neural network object detector. However, they only evaluated a small set of fixed parameters for each augmentation, but we showed that the achieved results for each augmentation heavily depends on the selected parameters. Recently they extended their study by training multiple state-of-the-art detectors with the augmentation policies, that

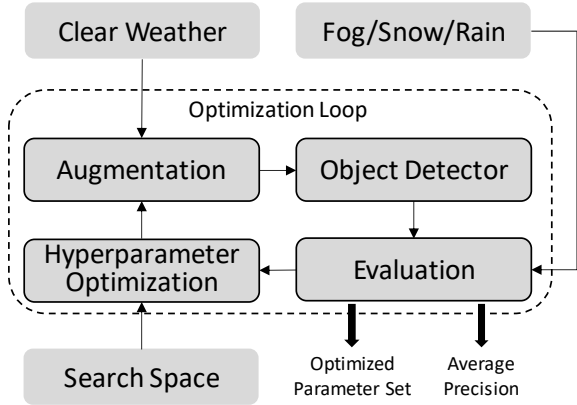


Fig. 1: Overview of the proposed optimization flow. The object detector is trained on clear weather recordings using different augmentation techniques (Section III). Then the detector is evaluated on fog, snow, and rain recordings (Section IV). To optimize the augmentation parameters the Optuna hyperparameter optimization framework [4] was used (Section IV).

achieved the best results for PointPillars. The other detectors showed similar performance increase like PointPillars. They also evaluated these augmentation policies on the STF data set [8], but focused only on the clear weather recordings. In different works Hahner et al. used a fog simulation [9] and a snow simulation [10] to augment point clouds to improve the detection performance in adverse weather. They trained multiple detectors on the STF clear weather data set and used the STF fog and snow data sets for evaluation. For both snow and fog they could show an improved performance when augmenting the point clouds with fog and snow respectively. However, they did not evaluate the influence of other common augmentation techniques separately and instead trained the detectors using the OpenPCDet [11] default augmentations for KITTI together with their simulations, so the influence of each individual augmentation on the detection performance is unknown. Rivero et al. [12] augmented point clouds in adverse weather by placing objects in the spray region of vehicles. The points caused by spray droplets behind driving vehicles caused false positive detections which triggered the emergency braking of the vehicle. To train a detector to not detect false positives in the spray region, they actually placed objects in this region, so that the detector learns to discriminate between spray points and actual object points.

III. POINT CLOUD DATA AUGMENTATIONS

This section describes the augmentation strategies which are evaluated in this work. Besides well established geometric data augmentation techniques, which are widely used on data sets recorded in good weather conditions, we also evaluate novel data augmentations that specifically add effects that can occur in adverse weather. Furthermore, we use weather simulations to directly augment the training data with expected input variations caused by the corresponding weather condition.

Weather Effects

Noise Augmentation: The noise augmentation mimics the occurrence of false points in adverse weather caused by backscattering. Therefore, noise is added to the point cloud in the form of additional points. The point positions x , y and z are sampled independently from a uniform distribution within the specified detection ranges. Moreover, each added point requires an intensity value i , for which we evaluate four different selection strategies. The intensity is either set to the minimum or the maximum value, or the intensity is sampled from a uniform distribution within the minimum and maximum values. The fourth strategy is the salt and pepper variant, where one half of the points gets the minimum intensity and the other half gets the maximum intensity. The number of points n added to the point cloud is sampled from a normal distribution $\mathcal{N}(0, \sigma^2)$ with $n \geq 0$.

Drop Out: The drop out augmentation mimics the point loss which occurs in adverse weather due to absorption and scattering, therefore points from the point cloud are randomly deleted. The fraction of points f which are deleted is sampled from a normal distribution $\mathcal{N}(0, \sigma^2)$ with $f \in [0, 1]$, where $f = 0$ means that no point is deleted and $f = 1$ means that all points are deleted.

Intensity Shift: The intensity shift augmentation shifts every points intensity by the same amount such that a point $P(x, y, z, i)$ becomes $P'(x, y, z, i + \Delta i)$, where Δi is sampled from a normal distribution $\mathcal{N}(0, \sigma^2)$. If the intensity exceeds the minimum or maximum intensity, it is set to the minimum or maximum respectively. This augmentation imitates the intensity alterations caused by changes in the environment or sensor settings.

Geometric Transformations

Random Translation: The random translation augmentation translates every point by the same amount in a random direction, such that a point $P(x, y, z, i)$ becomes $P'(x + \Delta x, y + \Delta y, z + \Delta z, i)$, where Δx , Δy and Δz are sampled independently from a normal distribution $\mathcal{N}(0, \sigma^2)$. The bounding box labels need to be translated by the same amount, such that a bounding box $BB(x, y, z, w, l, h, \Theta)$ becomes $BB'(x + \Delta x, y + \Delta y, z + \Delta z, w, l, h, \Theta)$.

Random Scaling: The random scaling augmentation scales every point by the same factor s , such that a point $P(x, y, z, i)$ becomes $P'(x \cdot s, y \cdot s, z \cdot s, i)$, where s is sampled from a normal distribution $\mathcal{N}(1, \sigma^2)$. The bounding box labels need to be scaled by the same factor, such that a bounding box $BB(x, y, z, w, l, h, \Theta)$ becomes $BB'(x \cdot s, y \cdot s, z \cdot s, w \cdot s, l \cdot s, h \cdot s, \Theta)$.

Local Scaling: Unlike the other augmentations which are applied to the whole point cloud, this augmentation applies only to the points inside bounding boxes. Each point in a bounding box together with the bounding itself is scaled as in the random scaling augmentation.

Random Flip: The random flip augmentation flips every point at the x -axis with a certain probability p , such that a point $P(x, y, z, i)$ becomes $P'(x, -y, z, i)$. The bounding

TABLE I: Combined augmentation policies

Policy	Augmentations
1	Noise Salt-Pepper
	Drop Out
	Random Translation
	Random Scaling
2	Fog Chamfer
	Random Translation
	Random Scaling
3	Rain
	Random Translation
	Random Scaling
4	Snow
	Random Translation
	Random Scaling

box labels need to be flipped as well and the orientation of the box needs to be adjusted, such that a bounding box $BB(x, y, z, w, l, h, \Theta)$ becomes $BB'(x, -y, z, w, l, h, (\Theta + \pi) \bmod 2\pi)$.

Filter Labels: This augmentation deletes all bounding box labels which contain less points than a certain threshold t . The point cloud remains unchanged.

Weather Simulation

Fog Augmentation: This augmentation applies the probabilistic fog simulation, which was developed in a previous work [13]. This fog simulation uses probabilistic calculations to determine individual point modification to match the characteristics of fog influence on the point clouds. The intensity of the fog simulation depends on the selected visibility distance v and the selected parametrization as described in [13]. The visibility distance is sampled from a normal distribution $\mathcal{N}(\mu, \sigma^2)$ with $v \in [0, \mu]$. Then the fog simulation is applied to the point cloud with probability p .

Rain Augmentation: Similar to the fog augmentation, this augmentation strategy applies the rain simulation from [13]. This rain simulation is based on raytracing on a generated rain volume. For each point in the point cloud multiple diverging rays are traced in a circular pattern to approximate the beam divergence of the laser beam. The point cloud is then modified based on different intersection ratios. A point is selected for modification, if its total intersection ratio is larger than a certain threshold T_{all} . Then this point is moved towards the sensor if the intersection ratio of a single rain drop is larger than the threshold T_{most} , otherwise this point is deleted. The intensity of this augmentation is adjusted by the precipitation rate r and the number of drops per unit volume n . The precipitation rate is sampled from a normal distribution $\mathcal{N}(0, \sigma_r^2)$ with $r \in [0, 20]$. The number of drops per unit volume is also sampled from a normal distribution $\mathcal{N}(0, \sigma_n^2)$

TABLE II: Parameter search results for each augmentation.

Augmentation	Parameters	Best Values
Noise Uniform	σ^2	9977.06
Noise Minimum	σ^2	9940.64
Noise Maximum	σ^2	9440.40
Noise Salt-Pepper	σ^2	9325.73
Drop Out	σ^2	0.29
Intensity Shift	σ^2	33.74
Random Translation	σ^2	2.04
Random Scaling	σ^2	0.04
Local Scaling	σ^2	1.92
Random Flip	p	0.19
Filter Labels	t	14
Fog Chamfer	σ^2, p	11.80, 0.80
Fog Distance	σ^2, p	92.07, 0.99
Fog Mean	σ^2, μ	64.37, 153.18
Rain	$\sigma_r^2, \sigma_n^2, p$	6.28, 97.40, 0.86
Snow	$\sigma_r^2, \sigma_n^2, p, s$	9.29, 150.50, 0.85, 9.85

with $n \in [0, 1200]$. The rain simulation is then applied with probability p to the point cloud.

Snow Augmentation: The snow augmentation applies the snow simulation from [13]. This snow simulation uses the same raytracing approach as the rain simulation to alter the point cloud but with a different volume generation. Since the size of the snowflakes depends on the simulated snow type this simulation has an additional scaling parameter s which scales the snowflakes from their molten diameter to their real diameter as described in [13]. The other parameters are the same as the parameters from the rain augmentation but with $r \in [0, 10]$. Like for the other weather augmentations, the snow simulation is then applied to the point cloud with probability p .

IV. EXPERIMENTS

Data Set

For evaluation the STF data set [8] was used, as it features recordings in clear weather conditions as well as in rain, snow and fog. Each frame in the STF data set is labeled with the weather condition it was recorded in. These labels are: clear, rain, light fog, dense fog and snow. In total there are 6306 frames in clear conditions, 4588 in snow, 1045 and 875 in light and dense fog respectively and 118 frames in rain. For training and validation a random 50/50 split of the recordings in clear weather was used. For testing the recordings in rain, light fog, dense fog and snow are used separately.




Fig. 2: Parameter search results for the translation and the scaling augmentation on the light fog data set. The points show the highest achieved average precision (AP) for independent training runs.

3D Object Detection

As object detector the state-of-the-art detector Voxel R-CNN [14] was used. For the evaluation we used the 3D object detection metrics from the KITTI data set [6]. For the reported average precision values the 40 recall points interpolation was used. For a detection to be counted as true positive the Intersection over Union (IoU) needs to be at least 0.7. Like in the KITTI 3D object detection evaluation, we report our results for the three difficulty levels *easy*, *moderate* and *hard*, which are based on the bounding box height, the occlusion and the truncation of objects.

Training

For training the HANNAH framework [15] was used with the implementation of the object detector from OpenPCDet [11]. The object detector was trained solely on the clear weather training set with one augmentation at a time. For each augmentation a hyper parameter search using the Optuna framework [4] was conducted. The model and training hyper parameter were kept the same for all training runs and only the augmentation parameter mentioned in Section III were optimized. For each training run the detector was trained for 80 epochs. Additionally, to the single augmentation evaluation we trained the detector with combinations of different augmentations. The combined augmentation policies are shown in Table I. As parameters we used the best parameter from the single augmentation evaluation. Moreover, we trained the detector using the Voxel R-CNN standard augmentation policy from [11].

Fixed Parameters

For the fog simulation there are two sets of fixed parameters which were obtained from two optimizations using different metrics as explained in [13]. The first parameter set is from the optimization of the fog simulation using the chamfer

distance to compare the simulated foggy point clouds to real point clouds in fog. For the other parameter set the point density by distance was used to compare the simulated to the real point clouds. In the following we call these fog simulation variants fog chamfer and fog distance. The mean μ of the normal distribution from which the visibility distance is sampled was set to 200 m. For the fog distance variant also a hyper parameter search was conducted including μ with a fixed probability of $p = 1$. This variant is called fog mean.

For the rain and snow simulation, other than in [13] the parameters, that determine the generated ray pattern for the beam divergence approximation are chosen to be $N_r = 2$ and $N_c = 5$, for performance reasons, which results in a total of 11 rays to be traced per point in the point cloud. The intersection thresholds, which determine if a point is deleted, moved or remains unchanged, depending on the intersection ratios on the rain/snow volume, are chosen to be $T_{all} = 0.15$ and $T_{most} = 0.8$ for both the rain and snow simulation. A detailed description of all the weather simulation parameters is given in [13].

Since in the STF data set the intensities are not normalized and the sensor output for the intensity is one byte, the minimum intensity in the STF data set is 0 and the maximum intensity is 255. These minimum and maximum intensity values were also used as minimum and maximum for the intensity shift. For the other evaluated augmentations there were no fixed parameter choices.

V. RESULTS

Parameter Search

The parameter search results are shown in Table II for each augmentation strategy. The given best values are the parameter values which were used in the training run that achieved the best results on the light fog data set or in case of the weather augmentations on the corresponding weather data set. As an

TABLE III: Results of the data augmentation evaluation on the STF data set. Average precision for the Car class for each augmentation and weather condition, the best results for each augmentation category are underlined and the over all best results are highlighted in bold.

		AP Car 3D [%]														
		Clear			Light Fog			Dense Fog			Snow			Rain		
		Easy	Moderate	Hard	Easy	Moderate	Hard	Easy	Moderate	Hard	Easy	Moderate	Hard	Easy	Moderate	Hard
Baseline		29.78	22.74	19.14	25.53	20.50	18.42	13.04	11.37	9.92	26.66	19.97	16.95	25.61	17.67	14.37
Voxel R-CNN Policy		41.23	32.32	28.01	33.59	27.49	25.64	19.69	16.89	14.92	36.24	28.26	24.36	33.72	23.93	20.31
Noise Uniform		28.10	21.25	18.09	23.75	19.21	16.74	8.41	7.37	6.36	24.80	18.79	15.83	27.04	18.34	15.05
Noise Minimum		29.20	22.72	19.50	25.29	19.82	18.34	10.28	9.17	8.06	26.44	20.24	17.14	29.02	18.93	16.25
Noise Maximum		29.92	22.47	19.06	23.57	19.07	17.55	10.49	8.76	7.39	25.49	19.18	16.02	28.99	19.73	16.80
Noise Salt-Pepper		29.90	22.18	18.76	24.57	19.67	17.94	10.32	8.60	7.80	25.73	18.86	15.75	22.51	14.52	13.09
Drop Out		<u>36.92</u>	<u>28.65</u>	<u>24.72</u>	<u>32.26</u>	<u>26.12</u>	<u>24.22</u>	<u>16.30</u>	<u>13.71</u>	<u>11.85</u>	<u>34.80</u>	<u>26.33</u>	<u>22.77</u>	<u>30.64</u>	<u>19.97</u>	<u>17.75</u>
Intensity Shift		30.03	23.11	19.62	25.14	20.47	18.83	11.69	10.12	8.90	27.65	20.89	17.69	25.96	16.53	13.79
Random Translation		44.21	34.34	29.29	38.82	30.85	28.91	<u>20.31</u>	<u>17.57</u>	<u>15.70</u>	41.00	31.61	27.80	41.69	28.80	<u>25.29</u>
Random Scaling		40.38	31.23	26.52	36.62	28.93	26.46	18.38	15.46	13.24	36.53	27.60	23.83	36.19	25.07	21.83
Local Scaling		30.44	22.74	18.92	24.63	20.22	17.36	10.65	9.54	8.16	24.14	18.18	15.76	22.33	14.53	12.07
Random Flip		30.19	23.74	20.62	24.28	20.49	18.35	14.14	11.61	10.79	26.25	20.30	17.43	23.29	16.37	14.59
Filter Labels		28.63	21.93	17.46	23.13	18.98	17.11	13.35	11.77	10.43	25.56	19.37	16.10	25.88	16.92	14.34
Fog Chamfer		<u>41.15</u>	<u>32.13</u>	<u>28.19</u>	<u>36.41</u>	<u>29.51</u>	<u>27.56</u>	24.33	20.89	17.96	<u>38.41</u>	<u>29.43</u>	<u>25.82</u>	38.73	26.84	23.67
Fog Distance		39.90	31.24	27.44	35.31	28.38	26.64	21.74	18.71	16.64	37.95	29.18	25.41	36.10	24.54	23.35
Fog Mean		39.56	31.11	27.24	35.78	29.03	27.25	23.60	20.30	17.41	38.06	29.36	25.65	<u>39.99</u>	<u>27.12</u>	25.51
Rain		32.74	24.28	20.52	25.63	20.55	18.93	10.63	9.44	8.10	27.92	21.18	17.76	22.72	14.56	12.58
Snow		36.74	28.58	23.93	30.65	25.13	22.75	18.86	16.24	14.60	34.88	26.59	22.95	37.23	25.81	22.38
Policy 1		42.89	32.60	28.54	36.99	29.97	28.06	21.77	18.44	15.96	39.64	30.33	26.33	35.33	24.87	22.22
Policy 2		41.24	32.06	26.94	36.52	29.70	27.91	21.85	18.91	16.32	36.81	28.98	25.26	36.78	25.41	22.16
Policy 3		44.21	<u>33.37</u>	<u>29.01</u>	<u>37.22</u>	<u>30.18</u>	<u>28.28</u>	20.49	17.52	15.60	<u>40.56</u>	<u>30.53</u>	<u>26.45</u>	37.01	25.11	22.56
Policy 4		42.14	32.62	28.40	36.23	29.48	27.52	<u>22.90</u>	<u>19.64</u>	<u>17.50</u>	39.71	30.34	26.19	<u>38.03</u>	<u>25.87</u>	<u>22.91</u>

example detailed results from the parameter search are given in Figure 2 for the translation and scaling augmentation. Both parameter search results show how important it is to choose the right parameters for augmentation. While it is sufficient to choose a sufficiently large value for σ^2 , to achieve good results in the random translation augmentation, σ^2 must lie within a relatively small interval to achieve good results for the random scaling augmentation. This shows that even a small change in parameter values can have a large impact on the obtained results and for an expressive comparison of augmentation strategies an extensive parameter search is necessary.

Detection Performance

Table III shows the results of the data augmentation evaluation for each augmentation under all weather conditions. From the weather effect augmentations only the drop out augmentation could gain a significant improvement compared to the baseline. The achieved increase in AP was 5.91 percentage points *p.p.*, 5.62 *p.p.*, 2.34 *p.p.*, 6.36 *p.p.*, and 2.30 *p.p.* for clear, light fog, dense fog, snow, and rain respectively at the moderate difficulty level. The noise augmentation variants and the intensity shift augmentation instead behaved poorly and could not improve the detection performance compared to the baseline, the noise augmentations even caused a significant decrease in AP in some conditions.

For the geometric augmentations the random translation achieved the best results on all data sets with an AP increase of

11.60 *p.p.*, 10.35 *p.p.*, 6.20 *p.p.*, 11.64 *p.p.*, and 11.13 *p.p.* for clear, light fog, dense fog, snow, and rain respectively at moderate difficulty. Moreover the random scaling augmentation also significantly improved the detection performance under all conditions but with somewhat lower results than the random translation augmentation. The other geometric augmentations could only marginally improve the baseline or even reduced the detection performance for some conditions.

Of the weather simulations the fog chamfer augmentation outperformed the other weather simulations on the clear, light fog, dense fog, and snow data set with an increased AP compared to the baseline by over 9 *p.p.* on all data sets at moderate difficulty. Only on the fog mean augmentation could achieved better results on the rain data set than the fog chamfer augmentation. The fog mean augmentation improved the baseline in rain by 14.38 *p.p.*, 9.45 *p.p.*, and 11.14 *p.p.* for the *easy*, *moderate* and *hard* difficulty respectively. The fog distance augmentation achieved similar results like the other fog augmentation variants with a marginally lower detection performance. The snow augmentation could also significantly improve the baseline in all conditions. However, the snow augmentation is outperformed by the fog augmentation variants on all data sets except the rain data set where it outperformed the fog distance variant at easy and moderate difficulty, but still lacks behind the fog mean and fog chamfer augmentation. The rain simulation performed poorly and could only achieve slight improvements in some conditions or reduced the detection

performance.

The combined augmentation policies all performed well on all data sets and reached high performance improvements, however none of the combined policies could improve the single augmentations. The best combined policy was policy 3, which achieved the same result as the random translation in clear weather at easy difficulty. On the other data sets and difficulty levels the results of policy 3 are only marginally lower than the random translation augmentation. Policy 4 showed similar results as policy 3 but outperformed policy 3 on the dense fog and rain data set.

As shown beforehand there are augmentations in each category that can significantly improve the baseline without augmentations. However, when training a neural network object detector, there is always a form of data augmentation applied. To show the real improvement of our proposed methods we now compare the best augmentations from each category to the standard augmentation policy from Voxel R-CNN.

In clear weather and light fog the over all best performing augmentation was the random translation with an increase in AP of 2.02 *p.p.* in clear weather and 3.36 *p.p.* in fog compared to the Voxel R-CNN policy at moderate difficulty. In dense fog the fog simulation augmentation was clearly the best performing augmentation with an improvement of 4.00 *p.p.* compared to the Voxel R-CNN policy at moderate difficulty. On the snow data set the best performing augmentation was again the random translation with an increased AP by 3.35 *p.p.* in comparison with the Voxel R-CNN policy at the moderate difficulty level. Also for the rain data set the largest improvement could be achieved using the random translation augmentation with an increase by 4.87 *p.p.* compared to the Voxel R-CNN policy at moderate difficulty. However, for the hard difficulty level the fog mean augmentation outperformed the random translation with an AP increase of 5.20 *p.p.* compared to the Voxel R-CNN policy.

VI. CONCLUSION & FUTURE WORK

In this work we improved the detection performance of a state-of-the-art object detector in clear weather as well as in adverse weather conditions using data augmentation techniques for LiDAR point clouds. For this we trained the detector Voxel R-CNN on the STF data set with recordings in clear weather and extended it using augmentations. The trained detector was tested on recordings in clear weather, fog, snow and rain. We evaluated both common and novel augmentation techniques, for each augmentation we conducted an extensive parameter search in order to get an expressive comparison of the different augmentations. Our parameter search results showed, that even small changes of the augmentation parameter can have a huge impact on the detection performance. Furthermore, we showed, that data augmentations can drastically improve the detection performance under all weather conditions. Moreover, we could improve the Voxel R-CNN augmentation policy significantly with an improvement in AP of 4.00 *p.p.* in fog, 3.35 *p.p.* in snow, and 4.87 *p.p.* in rain at moderate difficulty

For future research this evaluation will be extended to more augmentations and object detectors as well as other data sets. Furthermore, the fixed parametrization of the weather simulations will be included into the parameter search and the influence of different sampling strategies of the augmentation parameters will be investigated.

ACKNOWLEDGMENT

This work has been partially funded by the German Research Foundation (DFG) in the priority program 1835 under grant BR2321/5-2.

REFERENCES

- [1] European Commission, "Road safety: Commission elcomes agreement on new EU rules to help save lives," 2019, [Online]. Available: https://ec.europa.eu/commission/presscorner/detail/en/IP_19_1793, [Accessed: December 5, 2022].
- [2] G. Katz, C. Barrett, D. L. Dill, K. Julian, and M. J. Kochenderfer, "Towards proving the adversarial robustness of deep neural networks," *arXiv preprint arXiv:1709.02802*, 2017.
- [3] G. Volk, S. Müller, A. v. Bernuth, D. Hospach, and O. Bringmann, "Towards Robust CNN-based Object Detection through Augmentation with Synthetic Rain Variations," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, Oct 2019, pp. 285–292.
- [4] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, "Optuna: A next-generation hyperparameter optimization framework," in *Proceedings of the 25rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019.
- [5] M. Hahner, D. Dai, A. Liniger, and L. Van Gool, "Quantifying data augmentation for lidar based 3d object detection," *arXiv preprint arXiv:2004.01643*, 2020.
- [6] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *2012 IEEE conference on computer vision and pattern recognition*. IEEE, 2012, pp. 3354–3361.
- [7] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, "Pointpillars: Fast encoders for object detection from point clouds," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 12 697–12 705.
- [8] M. Bijelic, T. Gruber, F. Mannan, F. Kraus, W. Ritter, K. Dietmayer, and F. Heide, "Seeing through fog without seeing fog: Deep multimodal sensor fusion in unseen adverse weather," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11 682–11 692.
- [9] M. Hahner, C. Sakaridis, D. Dai, and L. Van Gool, "Fog simulation on real lidar point clouds for 3d object detection in adverse weather," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 15 283–15 292.
- [10] M. Hahner, C. Sakaridis, M. Bijelic, F. Heide, F. Yu, D. Dai, and L. Van Gool, "Lidar snowfall simulation for robust 3d object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 16 364–16 374.
- [11] O. D. Team, "Openpcdet: An open-source toolbox for 3d object detection from point clouds," <https://github.com/open-mmlab/OpenPCDet>, 2020.
- [12] J. R. Vargas Rivero, T. Gerbich, B. Buschardt, and J. Chen, "Data augmentation of automotive lidar point clouds under adverse weather situations," *Sensors*, vol. 21, no. 13, p. 4503, 2021.
- [13] 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 *2022 IEEE 95th Vehicular Technology Conference (VTC2022-Spring)*. IEEE, 2022, pp. 1–7.
- [14] J. Deng, S. Shi, P. Li, W. Zhou, Y. Zhang, and H. Li, "Voxel r-cnn: Towards high performance voxel-based 3d object detection," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 2, 2021, pp. 1201–1209.
- [15] C. Gerum, A. Frischknecht, T. Hald, P. P. Bernardo, K. Lübeck, and O. Bringmann, "Hardware accelerator and neural network co-optimization for ultra-low-power audio processing devices," *arXiv preprint arXiv:2209.03807*, 2022.