Upsampling Data Challenge: Object-aware Approach for 3D Object Detection in Rain *

Richard Capraru[1,2], Jian-Gang Wang[1], and Boon Hee Soong[2]

[1] Institute for Infocomm Research, Agency for Science, Technology and Research, Singapore 138632
{richard_capraru, jgwang}@i2r.a-star.edu.sg
[2] Department of Electrical and Electronic Engineering, Nanyang Technological University, Singapore 639798
ebhsoong@ntu.edu.sg

Abstract. Lidar-based 3D object detection has been widely adopted for autonomous vehicles. However, adverse weather conditions, such as rain, pose significant challenges by reducing both detection distance and accuracy. Intuitively, one could adopt upsampling to improve detection accuracy. Nevertheless, the task of increasing the number of target points, especially the key detection points crucial for object detection, remains an open issue. In this paper, we explore how an additional data upsampling pre-processing stage to increase the density of the point cloud can potentially benefit deep-learning object detection. Unlike the state-of-the-art upsampling approaches which aim to improve point cloud appearance and uniformity, we are interested in optimizing the object detection task. The object of interest, rather than full scenarios or small patches, is used to train the network - we call it object-aware learning. Additionally, data collection and labelling are time-consuming and expensive, especially for rain scenarios. To tackle this challenge, we propose a semi-supervised upsampling network that can be trained using a relatively small number of labelled simulated objects. Lastly, we verify a well-established sensor/rain simulator further, using a publicly available database. The experimental results on a database generated by this simulator are promising and have shown that our object-aware networks can extend the detection range in rainy scenarios by several meters and can achieve improvements in Bird-eye-View Intersection-over-Union (BEV IoU) detection accuracy.

Keywords: object detection · point cloud upsampling · generative adversarial network · object-aware learning · semi-supervised learning.

1 Introduction

LiDAR object detection in adverse weather conditions poses a significant challenge in autonomous vehicle research and remains an open issue. Given an unordered sparse point cloud received by LiDAR in the rain, one could explore point

* Supported by A*STAR

cloud upsampling to achieve a denser point cloud in order to improve the detection of different targets. Traditionally, upsampling methods have primarily been employed to support tasks like object classification [3] and surface smoothness reconstruction [4]. For upsampling experiments, the different patches of the samples present in the training dataset are selected for training and down-sampled using methods such as Poisson disk sampling. Subsequently, during the training phase, the network is tasked with upsampling the point cloud and comparing it with the ground truth to evaluate its performance. Traditional approaches for LiDAR point clouds typically employ the farthest point sampling (FPS) method to select seed points, followed by the application of the K-nearest neighbours (KNN) algorithm to acquire input patches. These patches are then merged after the upsampling process. While this approach presents good results for many applications, it has lower performance in scenarios involving adverse weather driving due to the extreme sparsity of data and the lack of emphasis on the downsampling pattern observed in natural settings.

Reconstructing complex geometry or topology from a sparse point cloud is still an open problem. Recently there has been some work done on optimizing object detection improving the resolution from low-resolution LiDAR (32 Ch) to high-resolution LiDAR (64 Ch), using 2D interpolation methods [6]. Their work demonstrated improved mAP (mean Average Precision) for different objects provided by the publicly available Kitti dataset [7]. However, their method is not adequate for rainy scenarios, where LiDARs can only receive point clouds with high sparsity. We aim to enhance and detect objects rather than obtain high-density point clouds. Building upon this motivation, we propose a novel object-aware upsampling approach, trained using the object of interest instead of small patches, to extend the detection range. The main contributions include:

- We propose a few object-aware upsampling strategies (an angle-invariant approach, a semi-supervised approach and an object-aware-traditional-patch-based combined approach) to increase the LiDAR detection range and to overcome the difficulty of collecting labelled data.
- We verify a well-established simulator and generate a rain database, which we can adopt as a benchmark for object detection in the rain.

The paper is organized as follows. In Section 2, we present the rain model used for our experiments. Section 3 provides an overview of existing upsampling technologies. In Section 4, we introduce our novel rain object-aware approach. The experimental results and concluding remarks are presented in Section 5.

2 Sensor and rain models, verification and simulated database

Prior to conducting experiments on actual rain data, simulations can be employed to efficiently verify our methods. For this purpose, we will utilize a state-of-the-art simulator [8] in our experiments. This simulator offers both the physical sensor model and the rain model. However, it should be noted that the simulator does not account for the noise introduced by rain.

The theoretical model used in this simulator regarding the impact of rain on LiDAR measurements can be found in [9]. The power received by a LIDAR sensor [10] (reflected intensity) is,

$$P_r(z) = E_l * \frac{c * \rho(z) * A_r}{2 * R^2} * \tau_T * \tau_R * exp((-2) * \int_0^z \alpha(z^{'})\, dz^{'}\ ) \tag{1}$$

where E_l is the laser pulse energy, c is the speed of light, $\rho(z)$ is the backscattering coefficient of a target, $\alpha(z^{'})$ is the scattering coefficient of rain along the path to a target, A_r is the effective receiver area, τ_T and τ_R are the transmitter and the receiver efficiencies. Without loss of generality, assuming a homogeneous environment, the constant coefficient $C_s = c * E_l * A_r * \tau_T * \tau_R$ represents the particular characteristics of the sensor, and can be ignored when calculating the relative sensor power, which is,

$$P_n(z) = \frac{\rho}{z^2} * e^{(-2)*\alpha*z} \tag{2}$$

Under clear weather conditions, corresponding to α=0.0, at the maximum detection range of a LiDAR (z_{max}), the maximum detectable power for a high reflectivity object (ρ=0.9/π) will be,

$$P_n^{min} = \frac{0.9}{\pi * z_{max}^2} \tag{3}$$

To calculate the relative sensor power measured under rainy conditions, the scattering coefficient, α, can be defined according to the power law [11],

$$\alpha = a * R^b \tag{4}$$

where R is the rainfall rate (mm/h), and a and b are empirical coefficients. The authors obtain the values a=0.01 and b=0.6 using the measurements of another paper [12]. Therefore the final model for the relative intensity returned by the LiDAR as a function of rainfall rate is:

$$P_n(z) = \frac{\rho}{z^2} * e^{(-0.02)*R^{0.6}*z} \tag{5}$$

The last equation is used to simulate rain in terms of rain rate. The points with a power/intensity less than the value defined by (3) will be eliminated.

3 Point Cloud Upsampling for Object Detection

Several upsampling approaches have been proposed in the literature, such as those mentioned in [3] and [14]. However, these methods have not been specifically evaluated or tested in the context of enhancing point clouds for rainy scenes to improve object detection performance. To the best of our knowledge, we are the first to explore upsampling methods to improve detection in rain.

The first algorithm explored in our experiments is the Multi-step Point cloud Upsampling network (MPU) [14]. The MPU network offers a notable advantage through its feature extraction unit, which employs K-nearest neighbours (KNN) search based on feature similarity to transform inputs into a fixed number of features. These features are further refined through a densely connected multilayer perceptron (MLP) chain. Furthermore, MPU's multi-step progressive upsampling and dynamic graph convolution capabilities make it particularly suitable for our research objectives, leading us to adopt it in our work.

Given the increasing interest in generative adversarial networks (GANs) and their remarkable success in image generation tasks, we further investigate the PU-GAN algorithm [3]. PU-GAN offers a distinct advantage over MPU in terms of its ability to generate uniform point clouds. This improved performance is attributed to the utilization of uniform loss and adversarial loss mechanisms within its framework. Considering these strengths, we include PU-GAN as the second algorithm explored in this paper.

4 Rain object-aware upsampling

Building upon the insights provided in the previous section, it has been established that traditional upsampling approaches primarily emphasize patches and treat individual points equally, overlooking the specific objects of interest. To enhance the accuracy of object detection, this study proposes a novel rain object-aware upsampling technique that extends the capabilities of patch-based approaches. This approach leverages the objects of interest to guide the upsampling process, thereby improving the overall effectiveness of object detection.

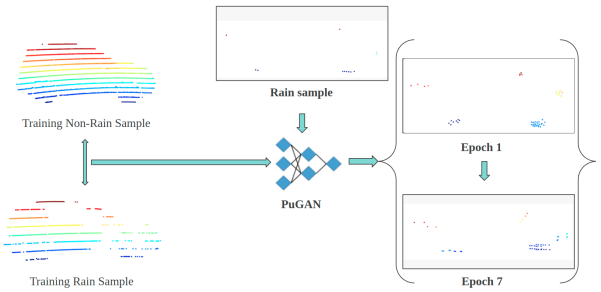

Fig. 1. Upsampling methodology

In this paper, we leverage the up-and-down sampling structure and adversarial training strategy of PuGAN to facilitate our research. Specifically, we implement PuGAN in Pytorch [19], which has been verified in [20]. As shown in Figure: 1, our upsampling network learns weights from some rain and non-rain data. The training data was generated as follows:

- Samples were recorded using the MSU Autonomous Vehicle Simulator (MAVS) [8] simulator on a clear day by placing the target at some distances and angles, e.g. 7 meters and 90 degrees in Fig. 1, which served as ground truth. Each sample in the ground truth training data contains the same number of points, e.g. 4096 points in this work.
- An angle-invariant upsampling approach was used: The network was trained using measurements of the car positioned at different angles, e.g. 400 samples are used in our experiments, which are generated with angle values from 0 to 90 degrees in intervals of 10 degrees.

Without loss of generality, an upsampling ratio of 4 is applied in our experiments. The angle-invariant approach has been trained for 19 epochs (with two different settings: 1. unsupervised trained using only non-rain object samples, 2. semi-supervised using 90% non-rain samples and 10% rain samples). In the final pre-processing experiment, we merged the points obtained from the semi-supervised approach with the pre-trained MPU upsampled points.

4.1 Rain object-aware framework

The first stage of our rain object-aware framework involves enhancing the point cloud by upsampling it to restore and augment 3D object features. Taking inspiration from Wu et al. [16], we employ a density-based spatial clustering with noise (DBSCAN) method for vehicle clustering. As presented in their approach, a recommended value for ϵ is 1.1 meters, as prior research has demonstrated that the average headway of vehicles in saturated flow conditions at signalized intersections is 2.18 meters [17].

After identifying the car cluster, the target point cloud is passed through the pre-trained MPU to assess its initial performance. For performance comparison with traditional methods, we employ the MPU (implemented in Tensorflow as described in [18]) pre-trained using the Sketchfab dataset. Without loss of generality, we can set the minimum number of points to 33, an up ratio of 16, a step ratio of 3, and a patch num ratio of 400.

In the second stage of our framework, we employ object detection on the enhanced point cloud data. To evaluate the effectiveness of our upsampling approach, we measure the object detection performance using the state-of-the-art network, CenterPoint [15] (utilizing the Second backbone [21] and voxel size of 0.1). We choose CenterPoint for its simplicity, high speed, and good accuracy.

5 Experimental Results

5.1 Simulation setting and analysis

To validate the effectiveness of the simulator adopted in this work, we conducted experiments using a publicly available database. Firstly, we focused on verifying the rain model using NuScenes [13]. In their paper [9], Goodin et al. validated the rain model by measuring the maximum detection range based on the rain

rate and object reflectivity. To assess the capabilities of our rain model, we reproduced these experiments and compared our results with the real-world findings presented by the BMW Research group in [22]. The experimental results, shown in Figure 2, illustrate the distance to the farthest point in a random NuScenes scene after applying the model (Equation (5)). We assumed uniform reflectivity for all points, employing the same values as [22]: 0.2 for the red line and 0.07 for the blue line.

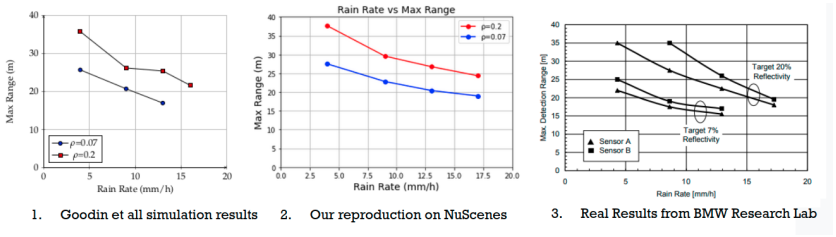

Fig. 2. Max Range vs Rain Rate for 1. Goodin et al simulation, 2. Our reproduction on NuScenes, 3. Real Results from BMW Research Lab

As Figure: 2, the validity of the model matches with the real experiments, motivating us to use the MAVS [8] simulator for this paper.

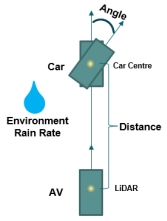

Fig. 3. Simulation setting

To assess the effectiveness of our upsampling method in enhancing LiDAR object detection performance, we generated a new database using the MAVS simulator. The MAVS simulator is developed with an MPI-based framework that enables parallel process coupling, along with a physics-based sensor simulator for LIDAR, GPS, cameras, and radars [8]. In this study, we utilized a Velodyne HDL-32E sensor positioned 1.84 meters above the host vehicle (AV) to generate simulation data. This LiDAR position was selected to match the recording configuration of the NuScenes dataset. The data recording process is depicted in Figure 3.

- A simulated car is placed in front of the AV's LiDAR at 10 meters with an orientation of 0 degrees (parallel to the direction of the host vehicle)
- We incrementally placed the car 1 meter away from the AV until it became undetectable. This iterative process was repeated for various angles of the car, ranging from 0 to 90 degrees (perpendicular to the AV), with an interval of 10 degrees. Additionally, we varied the rain rate from 30 mm/s to 70 mm/s during these experiments.

A total of 5,170 scenarios have been recorded (with nine samples for each). The hierarchy of this dataset is a series of subtrees which represent the data from different rain rates, angles and distances.

We evaluated the impact of rain on the above-mentioned dataset using the CenterPoint detector. The detector was applied to predict bounding box (bbox) scores and coordinates. For each sample scenario, we considered a minimum distance of 5 meters below the maximum detectable range, focusing specifically on challenging cases where detection confidence and accuracy are significantly reduced. The threshold for the bbox confidence score was set to 0.28, which is in close proximity to the commonly chosen value of 0.3 [23].

5.2 Semi/UnSupervised learning for Angle-invariant upsampling analysis

In the subsequent stage of our experiments, we aimed to determine the impact of different training approaches on detection performance. Specifically, we investigated whether employing a semi-supervised approach with a combination of rain non-rain pair training samples and non-rain training samples, or an unsupervised approach using only non-rain training samples, would influence the detection performance.

Through our experimental findings, we have observed that non-rain training samples provide advantages for smaller angles, while rain-non-rain pair training samples provide the greatest benefits for larger angles. As illustrated in Figure 4, the semi-supervised object-aware approach offers substantial benefits for both small angles (e.g., 0 degrees) and large angles (e.g., 80 degrees).

5.3 Best distance improvements

Moreover, we visually depicted the enhancement of the detection range for three different methods: (1) PuGAN semi-supervised object-aware method combined with MPU points, (2) Traditional state-of-the-art patch-based approach (MPU), and (3) PuGAN semi-supervised object-aware method. The visualization, presented in Figure 5, showcases the improvements achieved for two angles, namely 0 degrees and 80 degrees. The results illustrate the promising performance of combining points from different upsampling methods (object-aware and patch-based) in improving detection confidence in certain cases, while the semi-supervised method displays better improvements in other cases. Notably, both methods outperform traditional approaches in terms of performance.

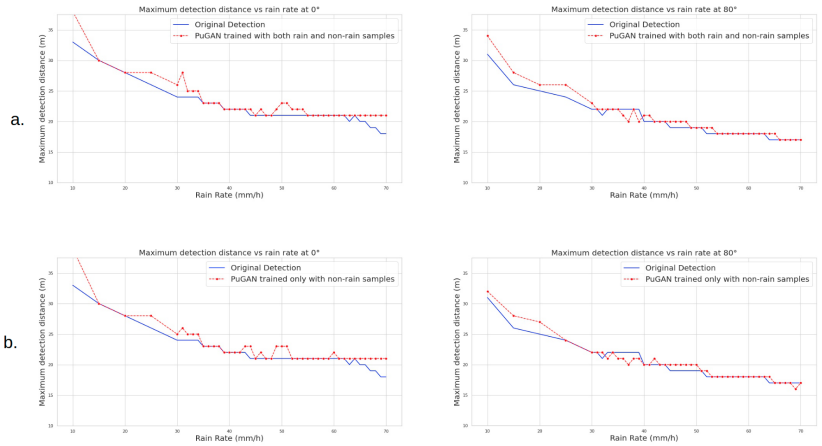

Fig. 4. Maximum detection distance vs rain rate for the target placed at an orientation of 0 and 80 degrees for a. Unsupervised PuGAN object-aware method trained only with non-rain samples, b. Semisupervised PuGAN object-aware method trained with both rain non-rain pairs and non-rain samples

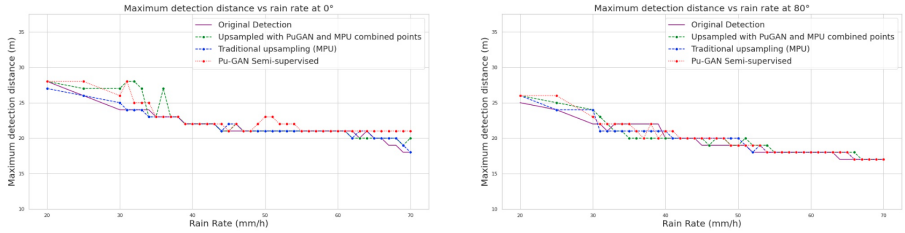

Fig. 5. Maximum detection distance vs rain rate for the target placed at an orientation of 0 and 80 degrees for: PuGAN semi-supervised object-aware method combined with MPU points (green line); Traditional state-of-the-art patch-based approach (MPU) (blue line) and PuGAN semi-supervised object-aware method (red line)

Lastly, it is also relevant to note that the PuGAN Angle-invariant method achieved a speed of 18 FPS using NVIDIA Corporation GM200 [GeForce GTX TITAN X], further bolstering the relevance of these methods. The speed improvement can be attributed to the exclusion of multiple patches, which would otherwise result in a substantial increase in computational power requirements.

5.4 Bird-Eye-View IoU

The evaluation of our method's improvements includes the use of the BEV IoU metric. Our methods can achieve good improvements in various scenarios, particularly in situations with extreme sparsity. Table 1 highlights the most significant

Table 1. Highest Improvement of BEV IoU for different angles, rain rates and distances

Rain	Dist.	Ang.	Max 2D IoU	Up Max 2D IoU	Impr.
10	32	90°	72.1%	88.3%	16.2%
15	31	70°	14.5%	65.3%	50.7%
20	24	90°	26.5%	89.2%	62.7%
25	25	80°	0%	85.6%	85.6%
34	25	70°	3.3%	69.0%	65.7%
35	24	70°	7.6%	60.4%	52.8%
37	25	60°	0%	66.7%	66.7%
39	21	70°	28.5%	81.2%	52.7%
41	20	80°	27.4%	88.1%	60.7%
42	21	80°	26.7%	87.6%	61.0%
45	23	0°	0%	60.8%	60.8%
50	23	20°	0%	89.2%	**89.2%**
53	22	0°	0%	69.8%	69.8%
62	19	60°	29.6%	83.6%	54.0%
70	19	0°	23.9%	78.6%	54.7%
Average Improvement					**60.2%**

improvements in IoU concerning rain rate, achieved through the combination of the PuGAN semi-supervised angle invariant object-aware method and the MPU points method. These improvements can be attributed to the inherent capability of our approach to enhance crucial object features that are essential for existing object detection methods.

Table 1 illustrates the upsampling improvements across various rain rates. In several instances, the CenterPoint model failed to detect the car before upsampling, resulting in an IoU value of 0. However, following the application of upsampling techniques, the network successfully detected the car with an accuracy ranging from 60% to 89%.

5.5 Experiment on Real Data

Finally, we conducted experiments on the Kitti dataset [7] to further evaluate the benefits of upsampling for object detection. In Figure 6, we present an example of the detection results obtained using the 3D-SSD detector [24] on a randomly selected scene. Initially, the original detection (with a bbox threshold score of 0.6) failed to detect certain cars. However, by individually selecting and upsampling the non-detected cars using the semi-supervised method, successful detection was achieved. This result emphasizes the significance of upsampling in enhancing detection confidence, particularly for objects situated at long distances. In our experiments, the cars were positioned at distances of 38, 43, and 45 meters, respectively.

Despite the promising nature of the findings, we acknowledge that there are anticipated limitations in our method. Expected and current limitations encompass the following aspects: potential instability in performance, leading to a

degradation in object detection confidence or accuracy when employed with various object detectors across diverse scenarios, necessitating further fine-tuning (e.g. For experiments conducted on Kitti, there is a minimal-to-none improvement in the average precision metric using the semi-supervised learning method); lack of generalizability, for example, domain knowledge or a segmentation algorithm are required to select the targets of interest. These limitations primarily stem from the ongoing research challenge of point cloud upsampling, as well as the fact that existing approaches have not been specifically designed to address object detection in adverse weather conditions. Specifically, the reconstruction of a target's shape and the utilization of LiDAR measurements for learning purposes prove challenging due to the inherent sparsity of the data. Moreover, introducing an excessive number of points or misplacing points during the upsampling process can adversely affect the detection performance. To overcome these challenges, it is imperative to devise a more robust upsampling algorithm that takes into account the unique difficulties associated with upsampling sparse LiDAR point clouds.

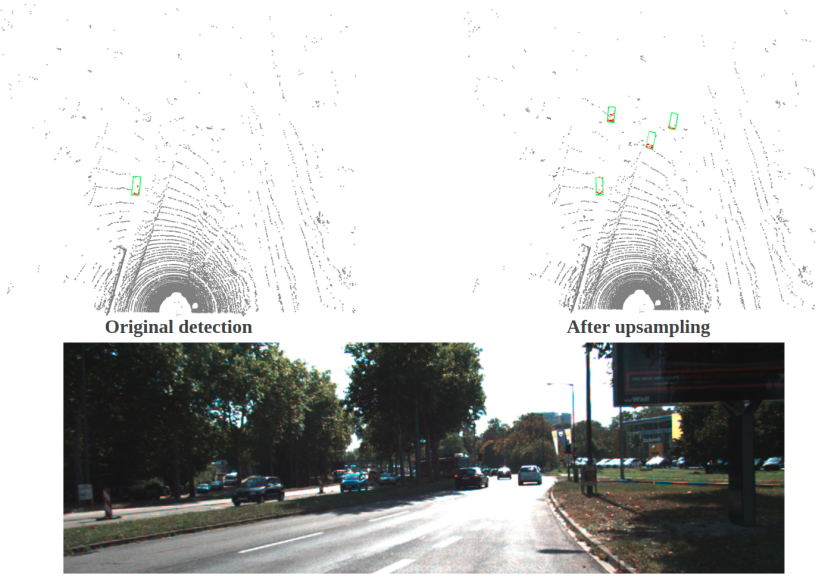

Fig. 6. Experiment on Real Kitti Data

6 Conclusion

Lidar object detection range as well as accuracy will be reduced significantly in rain. In this paper, we explore an object-aware upsampling method to increase

the LiDAR object detection range and accuracy in the rain. Different from the existing upsampling approach, which increases the density of different patches equally, we aim to detect an object. We verified a well-established simulator and the experiments on a database generated by this simulator have shown that our object-aware networks can extend the detection range from traditional patch-based upsampling approaches by several meters in rain conditions. In addition, it can improve object detection accuracy in terms of BEV IoU.

Although some preliminary results obtained by applying our novel approach presented in this study are very encouraging, more experiments and optimization are needed to make it a stable solution to perception in rain. Currently, limited experimentation has been conducted to ascertain the characteristics including the conditions under which the approach is effective or unsuccessful, as well as the accuracy limitations associated with it.

A benchmark database could be built based on more simulation data in the near future. Furthermore, experiments on real rain data, collected using our autonomous vehicle, will allow us to evaluate our methods under challenging adverse weather conditions. Finally, We could extend our work to improve robustness and adaptability across various targets, angles, and sparsity characteristics in future.

Acknowledgements This research is supported in part by grant no. I2001E0063 from the Singapore government's Research, Innovation and Enterprise 2020 plan (Advanced Manufacturing and Engineering domain) and administered by the Agency for Science, Technology and Research. The authors would like to extend their sincere gratitude to Dr. Teoh Eam Khwang for his invaluable suggestions and contributions during the writing of this paper. The authors would also like to thank the anonymous reviewers for their valuable comments.

References

1. J. Andrey and S. Yagar, "A temporal analysis of rain-related crash risk," Accident Analysis Prevention, vol. 25, no. 4, pp. 465–472, 1993. [Online]. Available: https://www.sciencedirect.com/science/article/ pii/0001457593900769
2. A. Kordani, O. Rahmani, A. Nasiri, and S. Boroomandrad, "Effect of adverse weather conditions on vehicle braking distance of highways," Civil Engineering Journal, vol. 4, p. 46, 02 2018.
3. R. Li, X. Li, C.-W. Fu, D. Cohen-Or, and P.-A. Heng, "Pu-gan: a point cloud upsampling adversarial network," 2019. [Online]. Available: https://arxiv.org/abs/1907.10844
4. H. Zhou, K. Chen, W. Zhang, H. Fang, W. Zhou, and N. Yu, "Dup-net: Denoiser and upsampler network for 3d adversarial point clouds defense." [Online]. Available: https://arxiv.org/abs/1812.11017
5. L. Yu, X. Li, C.-W. Fu, D. Cohen-Or, and P.-A. Heng, "Pu- net: Point cloud upsampling network," 2018. [Online]. Available: https://arxiv.org/abs/1801.06761
6. J. You and Y.-K. Kim, "Up-sampling method for low-resolution lidar point cloud to enhance 3d object detection in an autonomous driving environment," Sensors, vol. 23, no. 1, 2023. [Online]. Available: https://www.mdpi.com/1424-8220/23/1/322

7. A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: the kitti dataset," The International Journal of Robotics Research, vol. 32, pp. 1231–1237, 09 2013.
8. "Msu autonomous vehicle simulator," accessed: 2023-01-29. [Online]. Available: https://www.cavs.msstate.edu/capabilities/mavs.php
9. C. Goodin, D. Carruth, M. Doude, and C. Hudson, "Predicting the influence of rain on lidar in adas," Electronics, vol. 8, p. 89, 01 2019.
10. C. Dannheim, C. Icking, M. Mader, and P. Sallis, "Weather detection in vehicles by means of camera and lidar systems," in 2014 Sixth International Conference on Computational Intelligence, Communication Systems and Networks, 2014, pp. 186–191.
11. 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, no. 3, pp. 656 – 664, 2009. [Online]. Available: https://journals.ametsoc.org/view/journals/atot/26/3/2008jtecha1122 1.xml
12. A. Filgueira, H. Gonzalez-Jorge, S. Laguela, L. Diaz-Vilarino, and P. Arias, "Quantifying the influence of rain in lidar performance," Measurement, vol. 95, pp. 143–148, 2017. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0263224116305577
13. H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuscenes: A multimodal dataset for autonomous driving," 2019. [Online]. Available: https://arxiv.org/abs/1903.11027
14. W. Yifan, S. Wu, H. Huang, D. Cohen-Or, and O. Sorkine-Hornung, "Patch-based progressive 3d point set upsampling," 2018. [Online]. Available: https://arxiv.org/abs/1811.11286
15. T. Yin, X. Zhou, and P. Krahenb uhl, "Center-based 3d object detection and tracking," 2020. [Online]. Available: https://arxiv.org/abs/2006.11275
16. J. Wu, H. Xu, J. Zheng, and J. Zhao, "Automatic vehicle detection with roadside lidar data under rainy and snowy conditions," IEEE Intelligent Transportation Systems Magazine, vol. 13, no. 1, pp. 197–209, 2021.
17. "The influence of road familiarity on distracted driving activities and driving operation using naturalistic driving study data," pp. Traffic Psychology and Behaviour 52: 75–85, 2018.
18. Mpu tensorflow implementation," accessed: 2023-01-29. [Online]. Available: https://github.com/yifita/3PU
19. "Pytorch unofficial implementation of pu-net and pugan." accessed: 2023-01-29. [Online]. Available: https://github.com/UncleMEDM/PUGAN-pytorch
20. Q. Yang, Y. Zhang, S. Chen, Y. Xu, J. Sun, and Z. Ma, "Mped: Quantifying point cloud distortion based on multiscale potential energy discrepancy," IEEE Transactions on Pattern Analysis and Machine Intelligence, pp. 1–18, 2022.
21. Y. Yan, Y. Mao, and B. Li, "Second: Sparsely embedded convolutional detection," Sensors, vol. 18, p. 3337, 10 2018
22. R. Rasshofer, M. Spies, and H. Spies, "Influences of weather phenomena on automotive laser radar systems," Advances in Radio Science, vol. 9, 07 2011.
23. "Mmdetection3d: Openmmlab next-generation platform for general 3d object detection," accessed: 2023-01-29. [Online]. Available: https://github.com/open-mmlab/mmdetection3d
24. Z. Yang, Y. Sun, S. Liu, and J. Jia, "3dssd: Point-based 3d single stage object detector," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020.