

Energy-based Detection of Adverse Weather Effects in LiDAR Data

Aldi Piroli¹, Vinzenz Dallabetta², Johannes Kopp¹, Marc Walessa²,
Daniel Meissner², Klaus Dietmayer¹

Abstract—Autonomous vehicles rely on LiDAR sensors to perceive the environment. Adverse weather conditions like rain, snow, and fog negatively affect these sensors, reducing their reliability by introducing unwanted noise in the measurements. In this work, we tackle this problem by proposing a novel approach for detecting adverse weather effects in LiDAR data. We reformulate this problem as an outlier detection task and use an energy-based framework to detect outliers in point clouds. More specifically, our method learns to associate low energy scores with inlier points and high energy scores with outliers allowing for robust detection of adverse weather effects. In extensive experiments, we show that our method performs better in adverse weather detection and has higher robustness to unseen weather effects than previous state-of-the-art methods. Furthermore, we show how our method can be used to perform simultaneous outlier detection and semantic segmentation. Finally, to help expand the research field of LiDAR perception in adverse weather, we release the SemanticSpray dataset, which contains labeled vehicle spray data in highway-like scenarios. The dataset is available at <https://semantic-spray-dataset.github.io>

I. INTRODUCTION

LiDAR sensors are commonly used in autonomous driving applications alongside cameras and radars. LiDARs offer precise and rich depth information independently of the lighting conditions, helping an autonomous vehicle to understand and navigate its surroundings. The biggest drawback of LiDAR sensors is their sensitivity to adverse weather conditions such as rain, snow, and fog. Most perception algorithms developed for LiDAR sensors are trained and tested in favorable weather conditions. However, their performance is seen to degrade in adverse weather [1], [2], severely reducing their reliability. With the increasing number of autonomous vehicles on the roads, it is fundamental that they can navigate the environment and safely interact with other road users regardless of weather conditions.

In the literature, only a few approaches have been proposed to address this problem [4], [5], [6], [7], [8]. Sebastian et al. [7] classify weather conditions from a LiDAR scan. Zhang et al. [5] measure the degradation of LiDAR measurements in rainy conditions. Both methods only provide information on a scan-wise level, which does not allow the identification of adverse weather effects in the point cloud. However, detecting which points belong to solid obstacles and which are derived from adverse weather effects is essential for the reliable operation of an autonomous vehicle.

¹ Institute of Measurement, Control, and Microtechnology, Ulm University, Germany {firstname.lastname}@uni-ulm.de

² BMW AG, Petuelring 130, 80809 Munich, Germany {vinzenz.dallabetta, marc.walessa}@bmw.de and daniel.da.meissner@bmwgroup.com

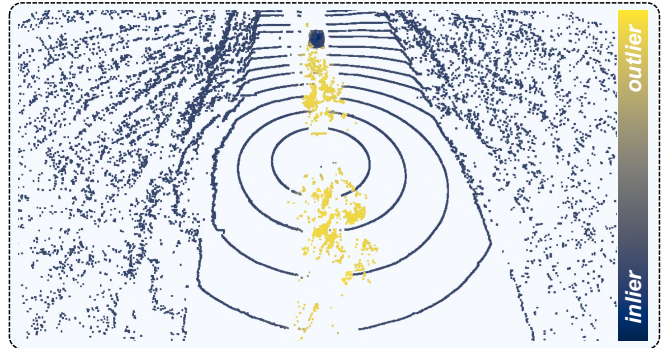


Fig. 1. We propose a method for detecting adverse weather effects in LiDAR data based on energy outlier detection. In the figure, we show a scene from the RoadSpray [3] dataset, where both a leading vehicle and the ego vehicle are driving on a wet surface, generating trailing spray. Our model is trained to associate low energy scores (darker colors) with inliers and high energy scores (lighter colors) with outliers, allowing for robust classification of adverse weather effects.

Heinzler et al. [6] propose a semantic segmentation network to classify adverse weather on a point-wise level. Due to the challenge of labeling adverse weather data, they test their model on data collected in a weather chamber, where a static scenario is created and artificial fog and rain are overlaid. However, real-world scenarios are much more challenging since complex dynamic scenarios can occur under different weather conditions.

In this work, we address the limitations mentioned above by proposing a method for the point-wise detection of adverse weather conditions. Instead of directly classifying adverse weather points, we reframe the task as an outlier detection problem. For this purpose, we adapt the energy-based outlier detection framework proposed in [9] to the 3D point cloud domain. More specifically, we rewrite the energy score formulation proposed in [9] for the point cloud domain and extend the proposed energy loss function to account for large class imbalances. We use this approach to differentiate between inlier points (buildings, vehicles, pedestrians, etc.) and outlier points (spray, rain, snow, fog, etc.). Fig. 1 shows a qualitative result of our approach in detecting spray. In extensive experiments, we show that our method performs better than previous state-of-the-art methods on real and simulated data. When training on datasets containing a single adverse weather effect, our approach shows higher robustness to unseen adverse weather effects, making it more applicable to real-world applications where multiple weather effects can occur simultaneously. Furthermore, our approach allows for the combined semantic segmentation of inlier

points and the detection of outlier points. When we compare our method with a state-of-the-art network for semantic segmentation [10], we see that our solution performs better in detecting points generated by adverse weather effects.

The lack of publicly available labeled data is the primary reason for the limited research on LiDAR in adverse weather. With this work, we help increase the research opportunities in this field by releasing the SemanticSpray dataset, which provides semantic labels for the RoadSpray [3] dataset. The dataset contains scenes of vehicles driving on wet surfaces, which can generate a trailing spray corridor. This effect is seen to be highly problematic in high-speed scenarios where a large number of spray points are introduced in the LiDAR measurements [1], impeding the sensors' field of view and, in extreme cases, causing perception systems to fail.

In summary, our main contributions are:

- We adapt the energy-based outlier detection framework [9] to detect adverse weather in LiDAR point clouds by formulating the point energy score and extending the energy loss function to account for large class imbalances.
- We show that our method outperforms the previous state-of-the-art approaches in detecting adverse weather effects in LiDAR data and has greater robustness to unseen weather effects.
- We show how our method can be adapted to perform both outlier detection and semantic segmentation of the inlier classes using different network architectures.
- We help expand the critical research field of LiDAR perception in adverse weather conditions by releasing the SemanticSpray dataset.

II. RELATED WORK

A. Adverse Weather Detection in LiDAR Data

Walz et al. [1] perform a study on the effect of vehicle spray on LiDAR and camera-based object detectors using a vehicle-mounted setup to simulate spray. The results show that state-of-the-art detectors are negatively affected by this weather effect, introducing misdetections and ghost objects in the scene. Zhang et al. [5] propose a method for estimating the degree of degradation for LiDAR scans in rainy conditions. They use an auto-encoder trained with the DeepSAD framework [11] to output an anomaly score for an input LiDAR scan. Sebastian et al. [7] propose a CNN-based method for classifying weather and road surface conditions in LiDAR data. Similar to [5], the method is not meant for classification on a point-wise level but rather on the entire LiDAR scan. Heinzler et al. [6] use a lightweight CNN architecture to detect rainfall and fog on a point-wise level. They train their network on data recorded inside a weather chamber where a static scenario is created using vehicles and mannequins, and then artificial fog and rain are overlaid. Although the proposed method performs well on the test data recorded inside the weather chamber, a generalization to dynamic real-world scenarios is not trivial [12]. Stanislas et al. [8] propose a CNN and voxel-based approach for detecting airborne

particles in robotic applications. Their CNN-based method uses a state-of-the-art CNN architecture developed for image segmentation. Their voxel-based approach instead uses a fully 3D convolutional backbone, resulting in high inference times. Although the primary goal of [8] is to detect smoke and dust, we test their approaches on adverse weather since airborne particles also include snow, rain, fog, and spray.

A variety of filtering-based methods has also been developed for the detection of outliers in point clouds. This includes voxel grid down-sampling, statistical outlier removal (SOR), and radius outlier removal (ROR). However, these general-purpose filters perform poorly on adverse weather condition detection [13]. For this reason, Charron et al. [13] propose the DROR filter, which aims to detect snow points by dynamically adjusting the radius of neighbor search to compensate for the non-uniform density of LiDAR point clouds. Kurup et al. [14] propose DSOR, which further improves DROR by considering the mean distance between neighbors. Overall, filtering-based approaches have a significant advantage over learning-based methods of not requiring training data. However, they have high computational complexity and generalize poorly to unmodeled weather effects, limiting their use in real-world applications.

B. Anomaly Detection Methods

Many approaches have been developed for anomaly detection, mainly for image data. The outlier exposure method introduced in [15] proposes to train a network with an auxiliary dataset of outliers, allowing the network to generalize to unseen anomalies. Ruff et al. [11] use a similar approach, redefining the loss function to map inlier data inside a hypersphere and outliers outside it. The anomaly score is then derived by computing the distance of the input from the center of the hypersphere. Recently, energy-based methods (EBMs) have emerged as the state-of-the-art for anomaly detection both on an image [9] and pixel [16] level. These methods rely on the energy function [9] to map the classification logits to a single real number. The network is trained to create an energy gap between inlier and outlier data which can be used for classification. In our work, we use a similar approach for the point-wise detection of adverse weather effects in LiDAR data. Liu et al. [17] propose the abstention learning paradigm, where a model is trained to abstain from making a classification when it is not certain that the input is an inlier. For this purpose, an additional class is added to the last classification layer of a network, and its output value is used for uncertainty estimation. Similar to [16], we use the idea of an additional class to extend our proposed network. However, different from [16], we do not optimize the abstention learning loss function and instead use an adapted version of the energy loss function proposed in [9] to detect outliers.

III. METHOD

This section describes our proposed method for detecting adverse weather points in LiDAR data. More specifically, our goal is to classify in a binary way if a point in a point

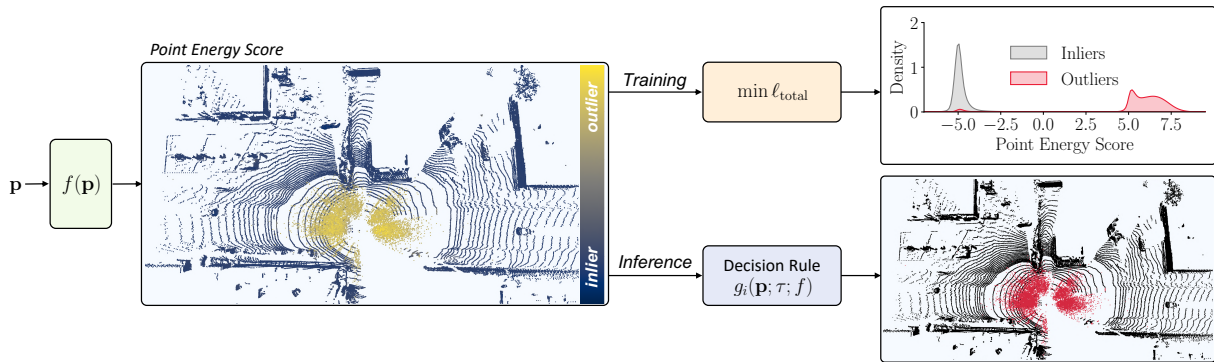


Fig. 2. Given an input point cloud \mathbf{p} , we aim to detect if a point $\mathbf{x}_i \in \mathbf{p}$ is caused by adverse weather. We reframe the problem as an outlier detection task and use the proposed point energy score to detect outlier points. During training, we minimize the loss function ℓ_{total} (3), which results in the model $f(\mathbf{p})$ associating a low energy score with inlier points and a high energy score with outliers. This creates an energy gap between the two categories, which can be used to select a classification threshold τ . The top-right plot shows an example of the energy gap between inlier and outlier points on the test set of the WADS [14] dataset when training $f(\mathbf{p})$ on snowy conditions (WADS training set). During inference, points are classified as outliers (red points in the bottom-right plot) if their energy score is greater than the threshold τ .

cloud is an inlier (not adverse weather) or an outlier (adverse weather). In Fig. 2, we show an overview of the approach.

A. Preliminaries

Brief Introduction to EBMs. Given a classification model $f(\mathbf{x}) : \mathbb{R}^D \rightarrow \mathbb{R}^K$, which maps an input $\mathbf{x} \in \mathbb{R}^D$ to K -logits, the *energy function* is defined as:

$$E(\mathbf{x}; f) = -\log \sum_{k \in \{1, \dots, K\}} e^{f(k|\mathbf{x})}, \quad (1)$$

with $f(k|\mathbf{x})$ being the k -th output logit. The output of $E(\mathbf{x}; f) : \mathbb{R}^D \rightarrow \mathbb{R}$ represents the *energy score* of the input \mathbf{x} . EBMs are characterized by their objective function, which aims to associate low energy scores with inlier inputs and high energy scores with outliers. The learned energy gap can then be used to classify inliers and outliers using an appropriate energy threshold [9].

Point Energy Score. We rewrite the energy function to process 3D inputs on a point-wise level. Let $\mathbf{p} = [\mathbf{x}_1, \dots, \mathbf{x}_N] \in \mathbb{R}^{N \times C}$ be a point cloud with N total points and $\mathbf{x}_i \in \mathbb{R}^C$ a single point with C features. Given a classification model $f(\mathbf{p}) : \mathbb{R}^{N \times C} \rightarrow \mathbb{R}^{N \times K}$ which outputs K -logits for each point $\mathbf{x}_i \in \mathbf{p}$, we define the *point energy function* as:

$$E_i(\mathbf{p}; f) = -\log \sum_{k \in \{1, \dots, K\}} e^{f_i(k|\mathbf{p})}. \quad (2)$$

The output of the function $E_i(\mathbf{p}; f) : \mathbb{R}^{N \times C} \rightarrow \mathbb{R}$ represents the *point energy score* of \mathbf{x}_i , with $i = \{1, \dots, N\}$. Here, $f_i(k|\mathbf{p})$ is the k -th output logit for point \mathbf{x}_i .

B. Point-wise Adverse Weather Detection

Dataset. Given a point cloud \mathbf{p} , we define the associated point-wise label set as $\mathbf{y} = \{y_1, \dots, y_N\} \in \mathbb{N}^N$. We reserve the labels $y_i \in \mathcal{Y}^{\text{in}} = \{1, \dots, Y\}$ for inlier classes (buildings, vehicles, pedestrians, etc.) and the label $y_i \in \mathcal{Y}^{\text{out}} = \{Y + 1\}$ for the outlier class (rain, snow, fog, etc.).

Network Architecture. Using the abstention learning (AL) framework, we extend the last classification layer of

$f(\mathbf{p})$ with an additional class. This class is not used to classify outliers directly but instead allows the model to abstain from classification when it is not confident that a point belongs to one of the \mathcal{Y}^{in} classes.

Although our method is not restricted to a particular network architecture, we propose a model for the specific case of adverse weather detection, which we name AdverseWeatherNet (AWNet). Fig. 3 shows an overview of the architecture. AWNet first voxelizes the input point cloud and then passes it through a stack of squeeze-excitation attention layers [18]. Each attention layer independently weights the channels of each voxel, and then a combined weight is derived by summing the individual weights. This value modulates the voxelized input channels via element-wise multiplication. The result is then processed using the sparse convolutions backbone presented in [19], which allows for rich feature extraction without the high computational cost of full 3D convolutions. Finally, the extracted features are classified using a set of fully connected layers.

Loss Function. The network $f(\mathbf{p})$ is trained using the following loss function:

$$\ell_{\text{total}} = \ell_{\text{cls}} + \lambda \ell_{\text{energy}}. \quad (3)$$

The first term in (3) is a semantic segmentation term, which in our case is the standard negative log-likelihood (NLL) loss function defined for an input point cloud \mathbf{p} , and associated point-wise labels \mathbf{y} as:

$$\ell_{\text{cls}} = \mathbb{E}_{\mathbf{x}_i \in \mathbf{p}} \left[-\log \frac{f_i(y_i|\mathbf{p})}{\sum_{k \in \{1, \dots, K\}} f_i(k|\mathbf{p})} \mid y_i \in \mathcal{Y}^{\text{in}} \right], \quad (4)$$

with $f_i(y_i|\mathbf{p})$ being the y_i -th output logit. As shown in [9], minimizing the NLL results in the model associating a low energy score with the inlier inputs. Moreover, the loss function allows to learn the semantic segmentation of the inlier classes. The second term in (3) is defined for an input

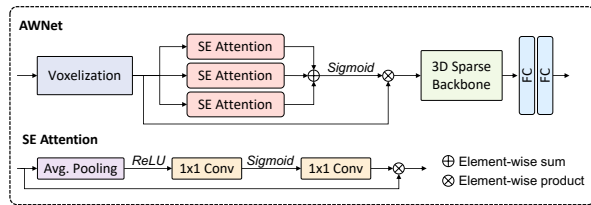


Fig. 3. AWNet architecture. We use a stack of three squeeze excitation [18] (SE) attention layers to weight the channels of each voxel. Then, we extract features using the 3D sparse convolution proposed in [19] and classify each voxel using two fully connected (FC) layers.

point cloud \mathbf{p} and associated point-wise labels \mathbf{y} as:

$$\begin{aligned} \ell_{\text{energy}} = & \\ & \mathbb{E}_{\mathbf{x}_i \in \mathbf{P}} \left[\omega_{in}^{-1} \left(\max(0, E_i(\mathbf{p}; f) - m_{in}) \right)^2 \mid y_i \in \mathcal{Y}^{in} \right] + \\ & \mathbb{E}_{\mathbf{x}_i \in \mathbf{P}} \left[\omega_{out}^{-1} \left(\max(0, m_{out} - E_i(\mathbf{p}; f)) \right)^2 \mid y_i \in \mathcal{Y}^{out} \right] \end{aligned} \quad (5)$$

with:

$$\omega_{in} = 1 + \sum_{y_i \in \mathcal{Y}^{in}} \mathbb{1}_{y_i \in \mathcal{Y}^{in}}, \quad \omega_{out} = 1 + \sum_{y_i \in \mathcal{Y}^{out}} \mathbb{1}_{y_i \in \mathcal{Y}^{out}}. \quad (6)$$

The ℓ_{energy} (5) term is a hinge loss function that penalizes inlier energies greater than the margin parameter $m_{in} \in \mathbb{R}$ and outlier energies lower than $m_{out} \in \mathbb{R}$ [9]. This creates an energy gap that can be used to define an adequate threshold to differentiate between inlier and outlier points. The task of point-wise anomaly detection in the case of adverse weather conditions can be an unbalanced problem. Even in severe precipitations, the number of non-adverse weather effect points can largely outnumber adverse weather ones. This large unbalance causes the ℓ_{energy} (5) term to saturate, resulting in insufficient supervision during training for the outlier energy association. We address this problem using the terms described in (6), which weight the energy loss terms proportionally to the number of inlier and outlier points present in the point cloud. The parameter $\lambda \in \mathbb{R}$ is used to weight the loss term ℓ_{energy} (5).

C. Training and Inference

Training. Like other energy-based methods, we use a fine-tuning approach for training [9], [16]. Given a model trained on a dataset containing only inlier points, we first extend its final classification layer using the AL framework described in Section III-B. Then, we fine-tune the model weights on a dataset that contains both inlier and outlier points using the loss function ℓ_{total} (3). This approach has the advantage of being efficient since it does not require the entire retraining of the network.

Inference. At inference time, our model returns the energy score for each point in the point cloud. An appropriate energy threshold $\tau \in \mathbb{R}$ can be chosen depending on the application at hand. For example, in autonomous driving applications, τ can be chosen so that a high fraction (e.g., 95%) of inlier points is correctly classified. The threshold τ is common for

TABLE I

EVALUATION RESULTS OF ADVERSE WEATHER DETECTION WHEN TRAINING AND TESTING ON THE SAME WEATHER EFFECT. VALUES ARE IN PERCENTAGE. \uparrow MEANS HIGHER VALUES ARE BETTER, AND \downarrow THAT LOWER VALUES ARE BETTER.

Method	D_{test}	AUROC \uparrow	AUPR \uparrow	FPR95 \downarrow
Particle-UNet [8]	WADS	97.18	97.17	12.62
	SemanticSpray	99.69	99.24	0.31
	NuScenes-Fog	99.98	99.98	0.02
	DENSE	99.61	98.96	0.68
	average	99.11	98.84	3.41
Particle-VoxelNet [8]	WADS	98.71	98.48	2.02
	SemanticSpray	99.16	98.96	0.70
	NuScenes-Fog	99.88	99.87	0.15
	DENSE	92.02	82.01	37.57
	average	97.44	94.83	10.11
WeatherNet [6]	WADS	97.40	96.68	10.59
	SemanticSpray	99.46	95.47	2.88
	NuScenes-Fog	99.89	99.60	0.04
	DENSE	99.49	98.90	1.94
	average	99.06	97.66	3.86
AWNet (ours)	WADS	98.26	96.89	1.24
	SemanticSpray	99.84	99.25	0.02
	NuScenes-Fog	99.99	99.98	0.01
	DENSE	99.27	97.55	0.77
	average	99.34	98.42	0.51

all points during inference. The decision rule g_i for a point $\mathbf{x}_i \in \mathbf{p}$ can then be formulated as:

$$g_i(\mathbf{p}; \tau; f) = \begin{cases} \text{inlier} & \text{if } E_i(\mathbf{p}; f) \leq \tau, \\ \text{outlier} & \text{else.} \end{cases} \quad (7)$$

IV. THE SEMANTICSPRAY DATASET

Within this work, we also release the SemanticSpray dataset. The data is based on the recently released RoadSpray [3] dataset, which performs large-scale testing of the effect of vehicle spray on LiDAR, camera, and radar sensors. The RoadSpray dataset provides unlabeled scenes of vehicles driving on wet a surface at different speeds in highway-like scenarios. For the sensor setup and additional information on the experiments, we refer the reader to the original publication [3]. To create the SemanticSpray dataset, we manually label each point in a LiDAR scan as one of three categories: *background*, *vehicle*, and *spray*. All static objects (road, vegetation, buildings, etc.) are labeled as *background*. The moving vehicles in the scene are labeled as *vehicle*. The *spray* class contains the spray generated by the ego vehicle and other moving vehicles in the scene. In total, we provide semantic labels for 16565 dynamic scenes, with approximately $6.23 \cdot 10^6$ *background*, $4.84 \cdot 10^4$ *vehicle* and $2.15 \cdot 10^4$ *spray* points. The dataset is available for download at <https://semantic-spray-dataset.github.io>. The dataset toolkit is instead available at https://github.com/aldipirolis/semantic_spray_dataset.

V. EXPERIMENTS

A. Experiment Setup

Baselines. We evaluate our approach against the other methods for adverse weather [6], [13], [14] and airborne

particle detection in LiDAR data [8]. Each learning-based baseline is trained to classify inlier (non-adverse weather), and outlier (adverse weather) points using the default parameters. Previous methods evaluate only on a single dataset which usually includes a single adverse weather condition. However, this testing approach is unrealistic since, in many real-world scenarios, multiple weather effects can occur simultaneously, i.e., snowfall and wet surface spray. Therefore, we test each method on the adverse weather effects used during training and on unseen ones.

Datasets. We evaluate the performance in a variety of different adverse weather conditions. SemanticSpray is the dataset presented in Section IV, which contains spray effects in highway-like scenarios. We use 7898 scans for training and 8667 for testing. The WADS [14] dataset was recorded in snowy conditions while driving in urban environments. We use 1011 scans for training and 918 for testing¹. The DENSE [6] dataset contains static scenarios inside a weather chamber where artificial rainfall and fog are generated. We use the official train (61900 scans), and test (19787 scans) splits. To test our proposed method on more complex scenarios under foggy conditions, we use the simulation method proposed in [20] to augment with fog the NuScenes dataset [21] (NuScenes-Fog). The dataset has 27449 scans for training and 6019 for testing. Since our main goal is to differentiate between inliers and outliers, unless otherwise stated, we consider all the possible non-adverse weather classes (e.g., vehicle, pedestrian, building, etc.) as the inlier class. Similarly, we consider all of the adverse weather classes (e.g., snow, fog, and spray) as the outlier class.

Evaluation Metrics. Following prior work [22], [15], [9], [16], we evaluate the outlier detection task using the area under receiver operating characteristics (AUROC), area under the precision-recall curve (AUPR) and the false positive rate at 95% true positive rate (FPR95). To compare against statistical filters, we use precision and recall metrics. Additionally, we use the intersection over union (IoU) and the mean IoU (mIoU) to evaluate semantic segmentation performances.

Implementation Details. We pretrain Awnet using the NuScenes [21] dataset on the foreground (vehicles, pedestrians, cyclists, etc.) and background class (everything else). We train for 30 epochs on 20% of the training set using the NLL loss, Adam optimizer [23], and constant learning rate of 10^{-4} . For the energy-based fine-tuning, we train for a maximum of 20 epochs using Adam optimizer and training parameters $m_{in} = -5$, $m_{out} = 5$ and $\lambda = 0.1$. When training on the SemanticSpray and DENSE datasets, we set ω_{in} and ω_{out} equal to 1 since there is a lower imbalance between inlier and outlier points. We implement Awnet with voxel size of $[0.1, 0.1, 0.2]$ m in the x , y and z direction respectively. A full description of the sparse backbone parameters is given in [19]. The final classification layer comprises two fully connected layers of size 256.

¹Used training/test splits: $\{13, 14, 17, 20, 23, 26, 30, 34, 35, 36\} / \{11, 12, 15, 16, 18, 22, 24, 28, 37, 76\}$.

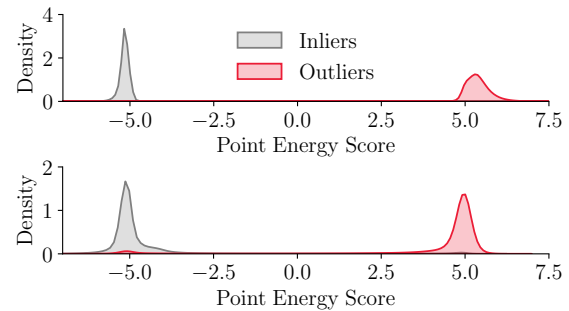


Fig. 4. Point energy score distributions of Awnet trained on the SemanticSpray dataset and tested on the SemanticSpray (top) and WADS (bottom) datasets. Although the network is trained only on spray data, the model associates similar energy scores to snow and spray points. This highlights the robustness of our method to unseen weather effects since a common classification threshold τ can be chosen to classify both effects.

B. Results

Training and Testing on the Same Weather Effect.

We start the evaluation by comparing the performance of Awnet against the other methods by training and testing in the same weather conditions. We report the results in Tab. I. Our method achieves the best average performance across all tested datasets in terms of AUROC and FPR95, improving the latter by 2.9% points compared to the second-best method (Particle-UNet). This last result shows that our method is well suited for safety-critical applications like autonomous driving, where a high percentage of adverse weather effects needs to be detected without having a large number of false positives. Fig. 1 and Fig. 5-top show qualitative results of the point energy score output of Awnet trained and tested on vehicle spray and snow respectively.

Training and Testing on Different Weather Conditions.

In Tab. II, we report the results of the methods trained on a single weather condition \mathcal{D}_{train} and tested on unseen (during training) weather conditions \mathcal{D}_{test} . We see that Awnet outperforms the other methods by a large margin across all metrics. When training on snowy conditions (WADS) and testing on spray (SemanticSpray), Awnet achieves an improvement of 3.17% AUROC, 24.12% AUPR and 17.29% FPR95 points against the second best network (Particle-VoxelNet). We see similar performances when training on spray and testing on snowfall. Although snowfall and spray are very different effects, our method can detect both as adverse weather, even when training with only one of the weather effects. This result shows that compared to the other approaches, the proposed point-wise energy-based outlier detection has a higher robustness to unseen adverse weather effects. This can also be observed in Fig. 4, where Awnet associates similar energy scores for both seen and unseen weather effects. This property is desirable for real-world applications where it is not possible to cover all possible weather variations during training. A qualitative example of Awnet trained on SemanticSpray and tested on WADS is reported in Fig. 5-bottom. We also see that both voxel-based approaches (Awnet and Particle-VoxelNet) perform better than the CNN-based methods when training and testing on different sensors.

TABLE II

EVALUATION RESULTS OF ADVERSE WEATHER DETECTION WHEN TRAINING ON AN ADVERSE WEATHER EFFECT AND TESTING ON AN UNSEEN (DURING TRAINING) WEATHER EFFECT. FOR EACH RESULT, $\mathcal{D}_{\text{TRAIN}}$ REPRESENTS THE TRAINING SET, AND $\mathcal{D}_{\text{TEST}}$ THE TEST SET. VALUES ARE IN PERCENTAGE. \uparrow MEANS HIGHER VALUES ARE BETTER, AND \downarrow THAT LOWER VALUES ARE BETTER.

Method	$\mathcal{D}_{\text{train}}$	$\mathcal{D}_{\text{test}}$	AUROC \uparrow	AUPR \uparrow	FPR95 \downarrow
Particle-UNet [8]	WADS	SemanticSpray	84.33	52.14	59.59
	SemanticSpray	WADS	74.57	63.75	70.17
	NuScenes-Fog	SemanticSpray	61.94	55.40	83.22
	NuScenes-Fog	WADS	93.30	92.74	35.58
	average		78.53	66.01	62.14
Particle-VoxelNet [8]	WADS	SemanticSpray	96.13	69.49	18.15
	SemanticSpray	WADS	93.02	90.14	23.10
	NuScenes-Fog	SemanticSpray	87.54	82.09	28.57
	NuScenes-Fog	WADS	95.32	95.06	10.07
	average		93.00	84.19	19.97
WeatherNet [6]	WADS	SemanticSpray	94.40	71.72	33.94
	SemanticSpray	WADS	81.58	68.53	58.64
	NuScenes-Fog	SemanticSpray	85.17	57.34	65.12
	NuScenes-Fog	WADS	88.83	87.00	58.99
	average		87.49	71.15	54.17
AWNNet (ours)	WADS	SemanticSpray	99.30	93.61	0.86
	SemanticSpray	WADS	96.74	91.93	7.19
	NuScenes-Fog	SemanticSpray	96.29	70.52	21.24
	NuScenes-Fog	WADS	97.35	95.26	4.52
	average		97.42	87.83	8.45

TABLE III

EVALUATION OF OUR PROPOSED METHOD AGAINST STATISTICAL FILTERS. OUR METHOD IS TRAINED AND TESTED ON THE SAME DATASET. VALUES ARE IN PERCENTAGE. \uparrow MEANS HIGHER VALUES ARE BETTER, AND \downarrow THAT LOWER VALUES ARE BETTER.

Method	$\mathcal{D}_{\text{test}}$	Precision \uparrow	Recall \uparrow
DROR [13]	WADS	85.48	88.11
	SemanticSpray	50.68	64.24
	NuScenes-Fog	70.00	84.09
	DENSE	57.92	67.65
	average	66.02	76.02
DSOR [14]	WADS	76.92	88.85
	SemanticSpray	51.57	56.58
	NuScenes-Fog	75.47	89.72
	DENSE	64.56	69.76
	average	67.13	76.23
AWNNet (ours)	WADS	91.11	95.45
	SemanticSpray	87.49	99.32
	NuScenes-Fog	99.22	99.84
	DENSE	91.43	95.73
	average	92.31	97.58

For example, the WADS dataset was recorded using a 64 layers LiDAR whereas the SemanticSpray dataset with a 32 layers sensor. CNN-based methods rely on a range image projection which is dependent on both the sensor resolution and sampling rate. In contrast, voxel-based methods are less affected by the sensor properties since they use a fixed voxel partitioning. In our experiments, we also tested the generalizability of models trained on the DENSE dataset. However, similar to what [12] reported, the restricted field of view and the internal structure of the fog chamber yields poor performances on outdoor datasets with 360° field of view.

Training With Simulated Data. All four methods per-

form well when training and testing with simulated fog data (NuScenes-Fog). When testing on snowfall (WADS), AWNet improves by 5.55% FPR95 points against the second-best method (Particle-VoxelNet). However, when testing on the SemanticSpray dataset, AWNet has lower performances than Particle-VoxelNet in terms of AUPR, suggesting that there is still a margin for improvement.

Comparison Against Statistical Filters. In Table III, we show the comparison between our proposed method and the statistical filters. Our method largely outperforms both filters in all datasets. DROR and DSOR perform well on snow (design task) and fog detection. When testing on spray, we see lower performances instead. Unlike snowfall and fog, spray is detected in dense clusters, making the use of statistical properties like the number of neighbors less effective for point-wise filtering.

Simultaneous Semantic Segmentation and Outlier Detection. In many applications, in addition to outlier detection, the segmentation of the inlier class might also be needed. Our method is well suited for this task since, in addition to differentiating between inliers and outliers, we also learn the semantic segmentation of the inlier classes using ℓ_{cls} (4). We test the performances of AWNet against Cylinder3D [10], a state-of-the-art semantic segmentation network, on the WADS dataset. Both networks are trained from scratch using only the WADS dataset. We train AWNet for 30 epochs using only ℓ_{cls} (4) and then fine tune for 20 epochs using ℓ_{total} (3) with parameters $m_{\text{in}} = -4.5$, $m_{\text{out}} = 0$ and $\lambda = 0.1$. During inference, we determine the semantic class of a point by choosing the maximum softmax logit among the $K + 1$ outputs. In addition, we classify as *snowfall* all points \mathbf{x}_i which have $E(\mathbf{x}_i; f) > \tau$, with $\tau = -0.25$. We train Cylinder3D using the default parameters described in [10]. Additionally, we train the same network using our proposed

TABLE IV

SEMANTIC SEGMENTATION EVALUATION WHEN TRAINING AND TESTING ON WADS. VALUES ARE IN PERCENTAGE.

Model	unlabeled	car	road	other-ground	IoU				mIoU
					building	vegetation	other-obstacle	snowfall	
Cylinder3D [10]	29.76	63.33	51.31	62.34	68.41	53.32	15.41	78.75	52.83
Cylinder3D-E [10] + ours	32.58	65.29	47.76	61.71	66.84	52.09	17.12	79.48	52.86
AWNet (ours)	12.44	56.43	58.06	60.77	67.64	60.06	17.32	87.03	52.47

approach (Cylinder3D-E). First, we train for 30 epochs with the default parameters, then we fine-tune for 20 epochs using both the original loss function and ℓ_{energy} (3) with parameters $m_{\text{in}} = -5$, $m_{\text{out}} = 5$ and $\lambda = 1$. For inference, we use the same approach described for Awnet with $\tau = 4.5$. During training and evaluation, we also use the *unlabeled* class since we observe that a small portion of the *snowfall* points is incorrectly labeled as such. We show the results in Tab. IV. Overall, Cylinder3D achieves a higher mIoU across the inlier classes than Awnet. However, in the adverse weather class (*snowfall*), Awnet outperforms Cylinder3D by 8.28% IoU points. Using our fine-tuning approach (Cylinder3D-E) improves the performances of Cylinder3D on the outlier class (*snowfall*) by 0.73% IoU points, showing that our method can be used with different network architectures.

Inference Time. We test the inference time of Awnet and the other methods on the SemanticSpray dataset using an NVIDIA RTX 2080 Ti. To process a single LiDAR scan Awnet takes on average 15.37 ms, Particle-VoxelNet 409.14 ms, Particle-UNet 4.32 ms, WeatherNet 4.18 ms, DROR 199.27 ms and DSOR 71.50 ms. Although both CNN methods have faster inference, Awnet is still real-time capable given the 10-20 Hz acquisition frequency of common LiDAR sensors. Furthermore, compared to the previous voxel-based state-of-the-art approach (Particle-VoxelNet), Awnet has more than 26-times faster inference time.

Discussion. Our approach demonstrates promising results in distinguishing between inliers and outliers while training and testing under similar weather conditions. Additionally, it offers good generalization performance for detecting unseen weather effects. However, our method has some limitations. For instance, when the point distribution of inliers and outliers is similar (e.g., snow and vegetation in Fig. 5-top), some inlier points incorrectly receive a high point energy score. This issue is further accentuated when testing under unseen weather conditions (Fig. 5-bottom).

C. Ablation Studies

Components Contribution. Tab. V shows the contribution of each component used when training Awnet on the WADS dataset. When we pretrain Awnet on inlier data, we see a small improvement in snowfall detection. When testing the same model on spray (unseen during training), we see an improvement of 0.73% AUROC, 7.54% AUPR, and 4.30% FPR95 points instead. This can be associated with the pretrained model already having an internal representation of the inlier class, allowing for a better generalization of the outlier class during training. This result is important since most LiDAR perception datasets are recorded in good weather

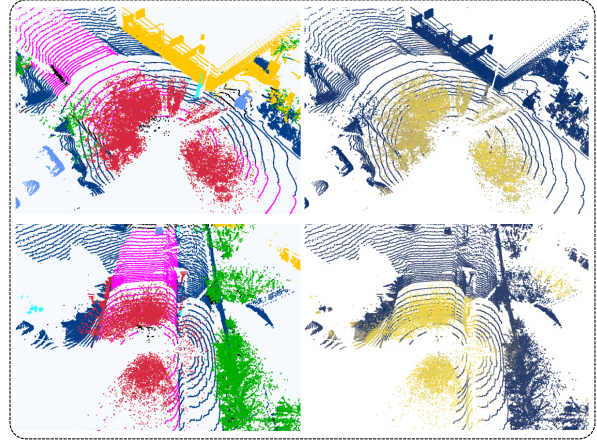


Fig. 5. The top row shows an example of Awnet trained and tested on the WADS dataset. The bottom row shows instead an example of Awnet trained on SemanticSpray and tested on WADS. On the left column, we show the semantic ground truth labels, whereas on the right, the output point energy score where lighter colors represent inlier scores and brighter colors outliers. The ground truth classes are represented with the following colors: ● unlabeled, ● car, ● road, ● other-ground, ● building, ● vegetation, ● other-obstacle, ● snowfall.

TABLE V

ABLATION STUDIES OF OUR PROPOSED METHOD WHEN TRAINING ON WADS. VALUES ARE IN PERCENTAGE. \uparrow MEANS HIGHER VALUES ARE BETTER, AND \downarrow THAT LOWER VALUES ARE BETTER.

$\mathcal{D}_{\text{test}}$	Pretrained Model	Weighted Energy Loss	AL	SE Attention	AUROC \uparrow	AUPR \uparrow	FPR95 \downarrow
WADS	-	-	✓	✓	97.97	96.75	1.33
	✓	-	✓	✓	98.23	97.05	1.22
	✓	✓	-	✓	98.08	96.54	1.40
	✓	✓	✓	-	98.13	96.67	1.28
	✓	✓	✓	✓	98.26	96.89	1.24
SemanticSpray	-	-	✓	✓	98.11	85.88	5.69
	✓	-	✓	✓	98.84	93.42	1.39
	✓	✓	-	-	98.96	87.82	3.36
	✓	✓	✓	-	98.60	90.89	2.53
	✓	✓	✓	✓	99.30	93.61	0.86

conditions. Our method shows that this large amount of data can be leveraged to improve the detection of adverse weather without the cost of additional labeled data. When we compare the results of Awnet trained from scratch on WADS and tested on SemanticSpray to Particle-VoxelNet, we still see an improvement of 1.98% AUROC, 16.39% AUPR and 12.46% FPR95 points, showing that even without pretraining our method has better robustness to unseen weather effects. The point-wise energy loss weighting term described in (6) further improves the generalizability performances of Awnet when testing on spray. In our experiments, we observed that in datasets where a large imbalance between inliers and outliers exists (WADS and NuScenes-Fog), the energy

TABLE VI

APPLICATION OF OUR PROPOSED LOSS FUNCTION TO DIFFERENT NETWORK ARCHITECTURES. ALL MODELS ARE TRAINED FROM SCRATCH ON THE WADS DATASET. VALUES ARE IN PERCENTAGE. \uparrow MEANS HIGHER VALUES ARE BETTER, AND \downarrow THAT LOWER VALUES ARE BETTER.

Method	$\mathcal{D}_{\text{test}}$	AUROC \uparrow	AUPR \uparrow	FPR95 \downarrow
Particle-UNet [8]	WADS	97.62	97.09	9.82
Particle-VoxelNet [8]		98.72	98.35	2.19
WeatherNet [6]		97.51	96.43	8.66
Particle-UNet [8]	SemanticSpray	87.43	52.56	53.37
Particle-VoxelNet [8]		98.48	88.48	5.03
WeatherNet [6]		93.77	72.15	31.60

weighting term helps the model consistently converge during training. Finally, we see that both the AL framework and the SE Attention modules increase performance when testing on the WADS and SemanticSpray datasets.

Different Network Architectures. In Tab. VI we test the performance of our proposed method using different network architectures. All models are trained from scratch, using the proposed loss function (3) and class-wise weighting described in (6). We set $\lambda = 0.01$ and margin parameters $m_{in} = -5$, $m_{out} = 5$ in all experiments. The results indicate that our approach can improve performance when evaluating on known and unknown adverse weather conditions. For instance, our method improves Particle-UNet’s FPR95 score on WADS by 2.8% points. Similarly, Particle-VoxelNet’s FPR95 score on SemanticSpray improves by 13.12% points. The performance gains can be attributed to the energy loss function, which provides additional supervision during training. By penalizing energy outputs in the range $[m_{in}, m_{out}]$, the network leans a more distinguishable representation between inlier and outliers points.

VI. CONCLUSION

In this work, we propose a new method for detecting adverse weather effects in LiDAR point clouds. We reframe the task as an outlier detection problem and use the recently proposed energy-based outlier detection framework to robustly detect adverse weather points. Extensive experiments on datasets containing spray, rain, snow, and fog show that our method performs better than previous state-of-the-art methods. Furthermore, our proposed method has higher robustness to unseen weather effects, increasing its applicability to real-world applications. Finally, we contribute to the expansion of the critical research field of LiDAR perception in adverse weather conditions by releasing the SemanticSpray dataset, which contains labeled scenes of vehicle spray data in highway-like scenarios.

REFERENCES

[1] S. Walz, M. Bijelic, F. Kraus, W. Ritter, M. Simon, and I. Doric, “A benchmark for spray from nearby cutting vehicles,” in *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2021, pp. 188–195.

[2] M. J. Mirza, C. Buerkle, J. Jarquin, M. Opitz, F. Oboril, K.-U. Scholl, and H. Bischof, “Robustness of object detectors in degrading weather conditions,” *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, pp. 2719–2724, 2021.

[3] C. Linnhoff, L. Elster, P. Rosenberger, and H. Winner, “Road spray in lidar and radar data for individual moving objects,” 2022-04. [Online]. Available: <https://tudatalib.ulb.tu-darmstadt.de/handle/tudatalib/3537>

[4] A. Piroli, V. Dallabetta, M. Walessa, D. A. Meissner, J. Kopp, and K. C. J. Dietmayer, “Robust 3d object detection in cold weather conditions,” *2022 IEEE Intelligent Vehicles Symposium (IV)*, pp. 287–294, 2022.

[5] C. Zhang, Z. Huang, M. H. Ang, and D. Rus, “Lidar degradation quantification for autonomous driving in rain,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 3458–3464.

[6] R. Heinzler, F. Piewak, P. Schindler, and W. Stork, “Cnn-based lidar point cloud de-noising in adverse weather,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2514–2521, 2020.

[7] G. Sebastian, T. Vattem, L. Lukic, C. Bürgy, and T. Schumann, “Rangeweathernet for lidar-only weather and road condition classification,” in *2021 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2021, pp. 777–784.

[8] L. Stanislas, J. Nubert, D. Dugas, J. Nitsch, N. Sünderhauf, R. Siegwart, C. Cadena, and T. Peynot, “Airborne particle classification in lidar point clouds using deep learning,” in *Field and Service Robotics*. Springer, 2021, pp. 395–410.

[9] W. Liu, X. Wang, J. Owens, and Y. Li, “Energy-based out-of-distribution detection,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 21 464–21 475, 2020.

[10] X. Zhu, H. Zhou, T. Wang, F. Hong, Y. Ma, W. Li, H. Li, and D. Lin, “Cylindrical and asymmetrical 3d convolution networks for lidar segmentation,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 9939–9948.

[11] L. Ruff, R. A. Vandermeulen, N. Görmitz, A. Binder, E. Müller, K.-R. Müller, and M. Kloft, “Deep semi-supervised anomaly detection,” in *International Conference on Learning Representations*, 2020.

[12] J. Egelhof, P. Wolf, and K. Berns, “Disturbance and particle detection in lidar data,” in *IECON 2022 – 48th Annual Conference of the IEEE Industrial Electronics Society*, 2022, pp. 1–6.

[13] N. Charron, S. Phillips, and S. L. Waslander, “De-noising of lidar point clouds corrupted by snowfall,” in *2018 15th Conference on Computer and Robot Vision (CRV)*. IEEE, 2018, pp. 254–261.

[14] A. Kurup and J. Bos, “Dsor: A scalable statistical filter for removing falling snow from lidar point clouds in severe winter weather,” *arXiv preprint arXiv:2109.07078*, 2021.

[15] D. Hendrycks, M. Mazeika, and T. Dietterich, “Deep anomaly detection with outlier exposure,” in *International Conference on Learning Representations*, 2019.

[16] Y. Tian, Y. Liu, G. Pang, F. Liu, Y. Chen, and G. Carneiro, “Pixel-wise energy-biased abstention learning for anomaly segmentation on complex urban driving scenes,” in *European Conference on Computer Vision*, 2021.

[17] Z. Liu, Z. Wang, P. P. Liang, R. R. Salakhutdinov, L.-P. Morency, and M. Ueda, “Deep gamblers: Learning to abstain with portfolio theory,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[18] J. Hu, L. Shen, and G. Sun, “Squeeze-and-excitation networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7132–7141.

[19] 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*, vol. 43, no. 8, pp. 2647–2664, 2020.

[20] 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 *IEEE International Conference on Computer Vision (ICCV)*, 2021.

[21] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, “nusenes: A multimodal dataset for autonomous driving,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 621–11 631.

[22] D. Hendrycks and K. Gimpel, “A baseline for detecting misclassified and out-of-distribution examples in neural networks,” in *International Conference on Learning Representations*, 2017.

[23] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *CoRR*, vol. abs/1412.6980, 2014.