

De-Snowing Algorithm for Long-Wavelength LiDAR

Bharat Jayaprakash[†], Matthew Eagon[†], Lu Zhan[†], William F. Northrop^{†,*}

Abstract— Long wavelength light detection and ranging (LiDAR) sensors have emerged as an essential component for increasing the accuracy and range of perception of autonomous vehicles because they employ directed lasers with wavelengths longer than $1\mu\text{m}$. However, adverse weather conditions like fog, rain, and snow pose a major challenge. Long-wavelength lasers generally exhibit increased absorption and scattering by water-based ambient particles compared to those with short wavelengths, which reduces sensor accuracy. Filtering out ambient particles is crucial for accurately representing the surrounding environment to ensure safe navigation. Despite extensive research on filtering snow particles from LiDAR point clouds, there is little documented research on long-wavelength LiDAR. Furthermore, existing filters that can be used with long-wavelength LiDAR sensors are limited in speed and accuracy, impeding their implementation in autonomous vehicles. In this paper, we propose a Network-Adjusted Reflectance Filter (NARF), a novel two-phase, physics-informed filtering method for long-wavelength LiDAR that outperforms the state-of-the-art geometric filters in terms of both speed and accuracy. The NARF first uses a physics-based range-corrected directional reflectance (RCDR) filter for initial snow particle classification, followed by a CNN-based RestoreNet to refine the RCDR predictions. Due to the lack of open-source datasets collected from long-wavelength LiDAR systems, we use a custom experimental dataset obtained during a snow event to train and validate the proposed filter.

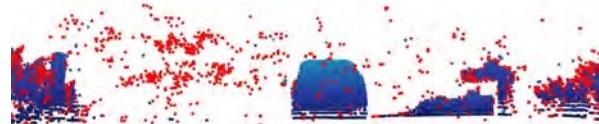
I. INTRODUCTION

The autonomous vehicle industry is rapidly evolving, driving fast advancements in LiDAR sensor technology and capacity, along with steadily decreasing costs. This makes large-scale LiDAR deployment in the near future inevitable [1]. However, as a laser-based sensor, LiDAR is naturally subjected to particle light scattering physics. For vehicles, this means that ambient particles present during adverse weather, such as fog, snow, and rain droplets can potentially reduce sensor accuracy. Contamination in LiDAR measurements due to snow particles remains a challenge.

It is crucial to tackle the degradation of LiDAR perception due to snow particles to achieve safer operation of autonomous vehicles. The understanding of automated LiDAR performance in snow is twofold. On one hand, historical studies investigating the physical and optical properties of snow particles [2] provide a guide for modeling an accurate perception environment from LiDAR. On the other hand, modeling LiDAR performance is indispensable to evaluating real-world measurement. Hence, the crucial aspects to address are: first, how to couple snow statistics with the LiDAR



(a) Experimental setup for custom dataset



(b) Point cloud view

Fig. 1: Spatial filters are typically less capable of detecting clustered snow

performance model; and second, how to integrate the coupled models into the data processing.

Light scattering theory holds that scattering is affected by parameters including particle shape, size distribution, liquid water content, laser wavelength, etc. Mie theory accounts for scattering where particle size is either of the same order as the laser wavelength or much larger. Larger particles result in a scattering diagram that is more peaked in the forward direction [3]. Automated LiDAR utilizes the particle back-scattered energy to differentiate observed objects from background noise. In [4], [5], rain droplet shape and size distribution are modeled, which are then used to characterize laser attenuation in a simulated environment. The correlation between LiDAR reflected intensity and rain intensity is evaluated through experimental studies [6]. Another experimental study quantified the actual laser attenuation at different wavelengths (785 nm and 1550 nm) in fog [7]. Dust particles have a comparable size to snow particles. When dust is present, four levels of LiDAR return signal degradation are derived, observed, and characterized in [8].

Commercially available LiDAR systems typically employ directed lasers with wavelengths of 905 nm (short-wavelength) or 1550 nm (long-wavelength). Despite longer-wavelength lasers having increased absorption by water-based particles such as snow and fog compared to shorter

[†]Department of Mechanical Engineering, University of Minnesota, MN

*Corresponding author.

wavelengths, there is a notable gap in the literature comparing the performance of long-wavelength and short-wavelength LiDAR systems for autonomous vehicle applications in adverse weather conditions. Additionally, existing snow-filtering methods have been developed and validated using datasets captured by short-wavelength LiDAR systems. This necessitates an investigation into the performance of existing filters on data captured by long-wavelength LiDAR systems.

The paper is organized as follows. Prevailing particle filters for LiDAR data are reviewed in Section II. The problem statement is formally defined in Section III. The experimental setup and data collection process are described in Section IV-A. The rest of Section IV explains the proposed Network-Adjusted Reflectance filter (NARF). Section V presents the experimental results using the NARF and compares the results to conventional filters.

II. RELATED WORK

Conventional filters capable of removing snow particles from LiDAR point clouds include spatial or distance-based filters, intensity-based filters, and reflectance-based filters, which we refer to collectively as physics-based filters, and finally, data-driven or learning-based filters.

A. Physics-based Snow Particle Detection Methods

The Radius Outlier Removal (ROR) filter counts the number of neighbors for each point within a specified search radius and filters out points with fewer neighbors than a set minimum [9]. The ROR filter uses a k-d tree data structure for efficient nearest-neighbor search. The Statistical Outlier Removal (SOR) filter calculates the mean distance of each point to its nearby points (the number of neighbors is specified). Points that surpass the weighted sum of the mean distance and standard deviation are filtered [9]. The SOR also uses a k-d tree for efficient search operations. One of the primary drawbacks of the ROR and SOR filters is that their parameters are fixed and not distance-dependent. This results in farther points having a higher probability of being classified as outliers. To address this limitation, the authors of [10] developed the Dynamic Radius Outlier Removal (DROR) filter which implements a distance-dependent search radius for the nearest neighbors search. The farther the point is from the LiDAR, the greater the search radius. The authors reported DROR precision and recall rates of over 90%. The authors of [11] proposed an adaptive ROR using PCA-based dimensionality reduction to remove noise from point clouds, outperforming other spatial and density-based filters. Subsequently, the authors of [12] developed the Low-Intensity Outlier Removal (LIOR) filter that classifies points as snow based on the intensity value of the point captured by the LiDAR and then uses a spatial filter to restore the few incorrectly classified snow points. The authors of [13] presented a de-snowing approach that first filtered out snow using an intensity filter and then leveraged spatiotemporal

information from the point clouds to restore the incorrectly filtered points.

One persistent drawback of spatial filters is their reliance on k-d trees for nearest-neighbor searches, which can be computationally intensive. In the context of autonomous cars where real-time processing is crucial, de-snowing filters should be designed for high-speed execution. Spatial filters are also more likely to miss moving snow clusters (Figure 1b). Moreover, these filters require empirical parameter tuning. A feature of the LIOR filter, discussed in Section V, is that it employs an intensity model tailored for 865 nm wavelength LiDAR systems, making it unsuitable for long-wavelength LiDAR systems. Long-wavelength LiDAR systems typically use higher-powered lasers compared to short-wavelength LiDAR systems. This fundamental difference may impact the performance of intensity-based filters. Moreover, laser power also influences measurement range, which may impact the performance of spatial filters.

B. Learning-based Snow Particle Detection Methods

Data-driven methods used for de-snowing LiDAR point clouds include WeatherNet [14] and LiSnowNet [15], both of which are constructed using convolutional neural networks (CNN). The authors of [16] present 4DenoiseNet, a neural network-based snow-denoising model that also takes in the time dimension as input, unlike any of the previous works. In more recent work, the authors of [17] introduced SLiDE, a self-supervised learning method for snow point removal in LiDAR point clouds. It uses two deep neural networks that exploit the low spatial correlations that snow points have with their neighbors. However, these methodologies require projecting the point clouds onto a spherical coordinate system as range images before filtering. Our current work bypasses this conversion and directly processes the point cloud. Furthermore, unlike the predominantly black-box models used in these approaches, our method incorporates a central physics-based component. This hybrid approach not only ensures a more explainable model but could also potentially provide more robust performance in varied environmental conditions.

C. Contributions

Our contributions are as follows: (1) We propose a novel range-corrected directional reflectance (RCDR) filter for filtering out snow particles from long-wavelength LiDAR point clouds. (2) We propose a Network-Adjusted Reflectance Filter (NARF), a novel, two-phase framework that uses a neural network, RestoreNet, to improve the baseline performance of the proposed RCDR filter.

III. PROBLEM DESCRIPTION

Within the context of de-snowing LiDAR outputs, a point cloud is defined as a set of 4-dimensional points $\{P_i|i = 1, 2, \dots, n\}$, where each point $P_i \in \mathbb{R}^4$ contains the 3-dimensional (3D) Cartesian coordinates and the reflectance value of the point. The number of points in a point cloud, n ,

may vary depending on the environment. Given an unordered set of points comprising a point cloud P , a snow filter should return a processed point cloud $Q = \{Q_i|i = 1, 2, \dots, m\}$, where $Q_i \in \mathbb{R}^4$ are the non-snow points from P . The objective is to find Q with high accuracy and precision. To function in a practical setting, filters should ideally be input order invariant (or permutation invariant) and have low computational time complexity to enable fast real-time processing.

IV. METHODOLOGY

A. Data Collection and Annotation

A stationary experiment was conducted to study LiDAR performance during a natural snow event. The LiDAR sensor used in this work is a prototype product manufactured by Luminar Technologies Inc. with a $1.55 \mu\text{m}$ wavelength. The sensor is capable of scanning 120° and 30° in the horizontal and vertical directions, respectively. The system can acquire approximately 880×64 data points at 10 Hz sampling frequency, achieving an angular resolution of roughly 0.136° and 0.47° in the horizontal and vertical directions, respectively. An outdoor test with the LiDAR mounted on top of a test vehicle resulted in snow accumulation on the laser lens leading to heavily contaminated snapshots. Since this can be circumvented only by systems like a heating unit or a shield-wiper and not by algorithmic post-processing, the LiDAR was placed indoors on a portable table facing the target vehicle approximately 10 meters away (Figure 1a). In addition to snow accumulated on objects, the LiDAR was also able to capture snow suspended in the air, similar to what a vehicle-mounted LiDAR would capture in the absence of intervening objects. This makes the setting a suitable representation of real-world driving.

The LiDAR sensor reported the 3D Cartesian coordinates and reflectance of the captured points. The reflectance value, or RCDR, is the ratio of the total received energy to the total incident energy of each laser beam. The RCDR value is scaled to be one at any range if the reflector is a Lambertian (isotropic) target. About 2837 LiDAR snapshots were collected during the snow event, out of which 200 consecutive snapshots were selected for training and testing the neural network. A two-step approach was adopted to annotate the dataset. Firstly, if the RCDR is below or equal to 0.02, the LiDAR sample point is labeled as snow, otherwise, it is labeled as non-snow. The threshold of 0.02 was derived theoretically as described in the next section. The second step is to manually annotate the snow particles to re-label the false positives on hard targets and the obvious true negatives.

B. RCDR-based filtering

In [18], a formula is established to model the RCDR of $1.55 \mu\text{m}$ wavelength laser from ambient snow particle:

$$I \sim \frac{P_r}{P_i} \frac{L^2}{O(L)} = \left[\frac{c_\tau}{2} A\eta \right] \times \beta_b \times e^{-2\beta_e} \quad (1)$$

where P_r, P_0 are the received and incident power respectively. L is the distance from receiving aperture to a certain snow particle. The term $O(L)$ is the LiDAR geometric factor and the first term on the RHS is LiDAR system factor [19]. These two factors can be treated as constant once the LiDAR is deployed. λ is the laser wavelength of $1.55 \mu\text{m}$. Usually for a cluster of snow particles, back-scatter coefficient β_b and extinction coefficient β_e can be defined as:

$$\beta_b = \int_0^\infty Q_{bck}(x, m) \pi r^2 n(r) \quad (2)$$

$$\beta_e = \int_0^\infty Q_{ext}(x, m) \pi r^2 n(r) \quad (3)$$

where r is the radius of the snow particle, assuming a spherical shape, and $n(r)$ is the number of particles with an equivalent radius of r . x is defined as a unit-less ratio $x = \frac{2\pi r}{\lambda}$ and m is the refractive index of snow, which can be approximated as 1.33 [2]. Q_{bck} and Q_{ext} are a function of snow particle size distribution. The analytical solution of Q_{bck} and Q_{ext} as a function of size parameter x can be computed from the software Mieplot [20].

For an individual snow particle, the RCDR can be reduced to a Beer-Lambertian approach as:

$$I \sim \text{Const.} \times Q_{bck}(x, m) e^{-2Q_{ext}(x, m)} \quad (4)$$

To model the snow size distribution, we adopt the Gunn-Marshall distribution as [21]:

$$n(D_m) = n_m \exp(-\gamma_m D_m) \quad (5)$$

Noted that n_m and γ are defined by following results in [22]:

$$n_m = 2500 R_s^{-0.94} \quad (6)$$

$$\gamma = 22.9 R_s^{-0.54} \quad (7)$$

R_s is the precipitation rate and D_m is the diameter of a water drop to which a snow particle melts. According to [21], the former can be interpolated by using the extinction coefficient as input. The visibility sensor measures the snow extinction coefficient at a fixed sampling spot. And assuming the homogeneity of snow distribution the precipitation rate of snow is estimated to be around 0.3 mm/hr .

In order to simulate the distribution of $n(D_m)$, we first sample snow particle diameter D_x from 0.01 mm to 1 mm uniformly. Based on the Equation 5, an inverse transform sampling technique is performed.

$$r_{inv} = \frac{1}{-2\gamma} \log\left(\frac{D_x}{n_m}\right) \quad (8)$$

where r_{inv} is the simulated snow particle radius randomly sampled by following the Gunn-Marshall distribution. We obtain and substitute Q_{ext} and Q_{bck} value back to Equation 4 to finally compute the simulated snow particle RCDR. The Constant in Equation 4 can be set to one without loss of generality. Comparing the histogram of RCDR from the simulation and the experimental data in Figure 2, a strong agreement is observed. More importantly, both simulated RCDR and experimental data present the same cut-off value at 0.02, which is then set to be the threshold for RCDR-based filtering. The distribution plot for the experimental data for RCDR values larger than 0.02 can be noisy with

multiple local peaks, which can make the determination of an appropriate threshold difficult. This supports the need for a physics-based simulation to determine the cut-off value.

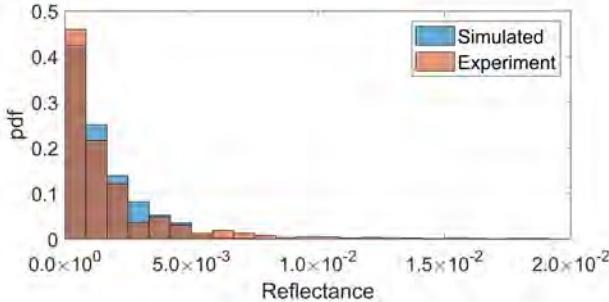


Fig. 2: Histogram of simulated and experimental RCDR values

It was observed that the RCDR filter exhibited varying classification capabilities based on the RCDR values of the points in the cloud. The filter has high accuracy above the 0.02 threshold, which means high confidence in predicting non-snow points. The relatively low accuracy of the RCDR filter below the threshold of 0.02 calls for a second filtering method specifically to re-classify the detected snow particles to correct false positives.

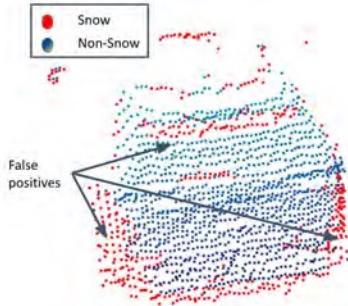


Fig. 3: Visualization of the RCDR filter's output. False positives were detected on the rear side of the vehicle due to the high angle of incidence of laser beams

A notable observation from the output in Figure 3 is the nature of the false positives detected by the RCDR filter. Surfaces like the rear wheel arc of the vehicle are oriented such that the angle between their normal and the line of sight from the LiDAR is high. Such spatial orientations result in the laser beams hitting the surface with high angles of incidence because of which the intensities of the beams returning to the sensor are low. This causes the RCDR filter to classify such points as snow. Hence, it is necessary to use a secondary filter that restores the false positive points back into the filtered point cloud.

C. Restoring False Negatives

Inspired by the accuracy of WeatherNet [14], we implement a CNN-based framework called RestoreNet (Figure 4) that restores false positives detected by the RCDR filter back into the filtered point cloud. Carrying forward the definitions

from the authors of [23], a set of points in \mathbb{R}^n that make up a point cloud has the following important properties:

- 1) The points are unordered. Hence, any class of filter that is applied on such point clouds needs to be input order invariant, i.e., a change in the order of the input must not result in a different point-wise output.
- 2) Neighboring points form meaningful local representations. Points that are close to each other are more likely to share geometric and optical properties than points that are far from each other.

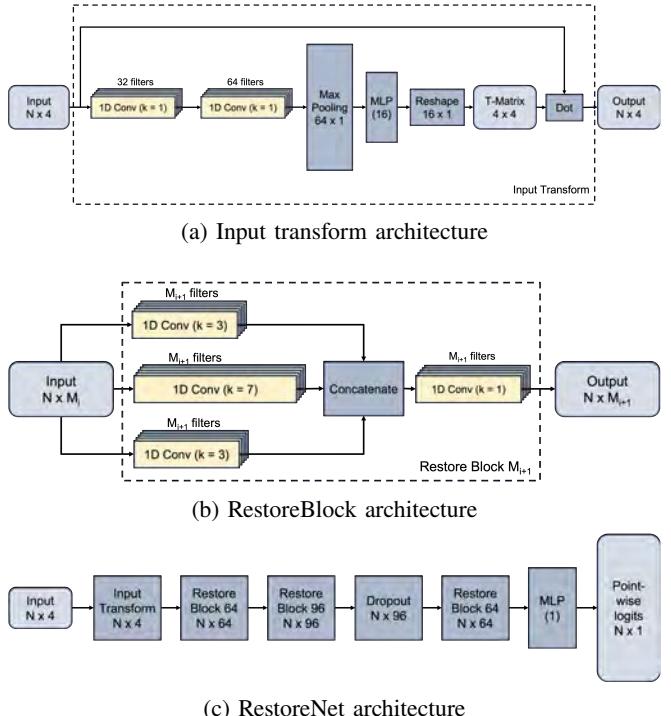


Fig. 4: Illustration of the individual components of RestoreNet inspired by WeatherNet architecture [14]

Considering the two properties, it was necessary to transform the input point cloud to a canonical order before allowing the network to learn local representations. The input transform block in Figure 4 predicts a Cartesian transformation matrix (labeled as T-Matrix in Figure 4a) that transforms the raw input cloud into a canonical order. The transformed point cloud is then fed into a series of RestoreBlocks that form the building blocks of RestoreNet. The RestoreBlocks consist of convolutional filters and transform input point sets from size $N \times M_i$ to $N \times M_{i+1}$. The architectures of the RestoreBlock and RestoreNet are case-specific variants of LiLaBlock and LilaNet [24]. The complexities of the respective architectures are reduced, given that the purpose of the RestoreNet is to correctly reclassify false positive snow particles. A dropout regularization layer is added after the second RestoreBlock to prevent overfitting. The output of the RestoreNet is a set of logits $\{S_i | i = 1, 2, \dots, N\}$, where $S_i \in \mathbb{R}^{N \times 1}$ represents the probability that the i_{th} point being a snow particle.

Algorithm 1: Network-Adjusted Reflectance Filter

```

Input:  $P = \{[P_{ix}, P_{iy}, P_{iz}, P_{ir}] | i = 1, 2, \dots, N\}$ ,  $r_{thres}$ 
Output:  $S, NS$ 
1  $NS \leftarrow \emptyset$  // Points classified as non-snow
2  $S \leftarrow \emptyset$  // Points classified as snow
3 for  $p$  in  $P$  do
4   if  $P_{ir} \leq r_{thres}$  then
5      $S \leftarrow S \cup \{p\}$ 
6   else
7      $NS \leftarrow NS \cup \{p\}$ 
8  $logits \leftarrow \text{Run RestoreNet}(P)$ 
9 for  $p$  in  $S$  do
10  if  $logits[p] < 0.5$  then
11     $S \leftarrow S \setminus \{p\}$  // Remove from S
12     $NS \leftarrow NS \cup \{p\}$  // Restore to NS

```

D. RestoreNet training

Network hyperparameters were optimized using grid search with 3-fold cross-validation. The dropout layer probability was set to 0.3. The batch size b for training was set to 10. The learning rate was scheduled with an exponential decay rate of 0.9 for every 150 epochs. The Adam solver was used for the training process, with the suggested default values $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\varepsilon = 10^{-8}$ [25]. The dataset is split into a training set (75%) and a test set (25%). The loss function used for forward pass evaluation was binary cross entropy as the network is a point-wise binary classifier. Since the dataset consists of point clouds with stationary objects and non-stationary snow particles, we implemented a data augmentation technique in which each point cloud was subjected to 3D axis translation to prevent overfitting.

$$\begin{aligned} P' &= P + d \\ d &\sim U(-D, D) \end{aligned} \quad (9)$$

The translation can be mathematically represented by Equation 9, where P is the original point cloud, P' is the augmented point cloud, and d is a random number generated from a uniform distribution with upper and lower bounds $-D$ and D respectively. The value of D for the training process was set to 5 meters. Additionally, the raw point clouds were padded with dummy points positioned at the origin to ensure all the point clouds to had the same number of points ($N=22000$), ultimately aiding in the training process.

V. RESULTS AND DISCUSSION

A threshold-based RCDR filter was first implemented on the custom dataset. Since the RCDR filter is not a spatial filter and operates on individual points, it remains effective regardless of the LiDAR's angular resolution. RestoreNet was then trained on the dataset to restore false positives back into the filtered point cloud with hyperparameters optimized as discussed in Section IV-D. Figure 5 shows the bin-wise accuracy of RestoreNet, where each bin is 0.025 units long in the output probability space and the bins span the whole output range: $[0.0, 1.0]$. It was observed that RestoreNet was least confident when it classified a point with a probability in

Table I: Performance Metrics Comparison

Metric	DROR [10]	LIOR [12]	RCDR	NARF (<i>proposed</i>)
Accuracy	0.6910	0.9444	0.9517	0.9879
Precision	0.9105	0.8740	0.8767	0.9814
Recall	0.0421	0.9660	0.9887	0.9808
IOU	0.0419	0.8480	0.8680	0.9629
TNR	0.9980	0.9341	0.9342	0.9912
FPR	0.0019	0.1391	0.0658	0.0088
FNR	0.9578	0.0160	0.0112	0.0192
MFT	0.5946	0.3769	0.0005	0.1957

Table II: Effect of shuffled input

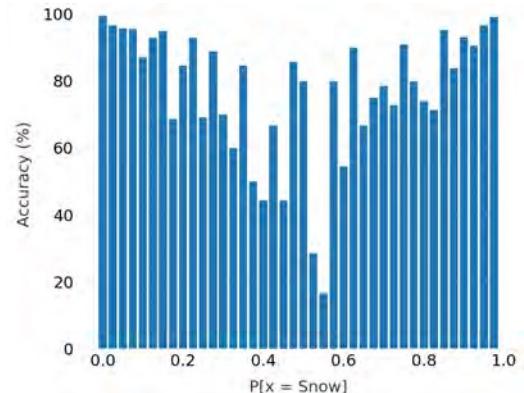
Shuffle %	0%	25%	50%	75%	100%
Accuracy	0.9886	0.9809	0.9740	0.9675	0.9579

the ranges $[0.4, 0.475]$ and $[0.525, 0.625]$. However, it was observed that predictions in these ranges were predominantly for points that are not snow particles, for which the RCDR filter exhibited high accuracy. Figure 7 provides a visual comparison between the outputs of the NARF and the RCDR filter. It can be observed that the false positives that were hard targets on the rear of the vehicle are restored back into the point cloud.

Input order invariance is an important constraint for the NARF. An input transform block, as shown in Figure 4a, was used to transform the input points to a canonical form. To test the robustness of the algorithm to changes in input point order, a random point cloud in the test dataset was shuffled to four progressively more extensive degrees: 25%, 50%, 75%, and 100% of points. For each of these scenarios, and the non-shuffled case, the accuracy of the NARF was calculated as reported in Table II. The accuracy dips slightly as the extent of shuffling increases, and the accuracy falls by 3.07% on average when all points are shuffled, showing that the algorithm is fairly robust to changes in input point order.

A. Performance Metrics

Three preliminary performance metrics, four key performance metrics, and a time complexity metric are used to compare the different de-snowing methods. The preliminary performance metrics used are the True Negative Rate (TNR), False Positive Rate (FPR), and False Negative Rate (FNR).


Fig. 5: Accuracy of RestoreNet for output bins

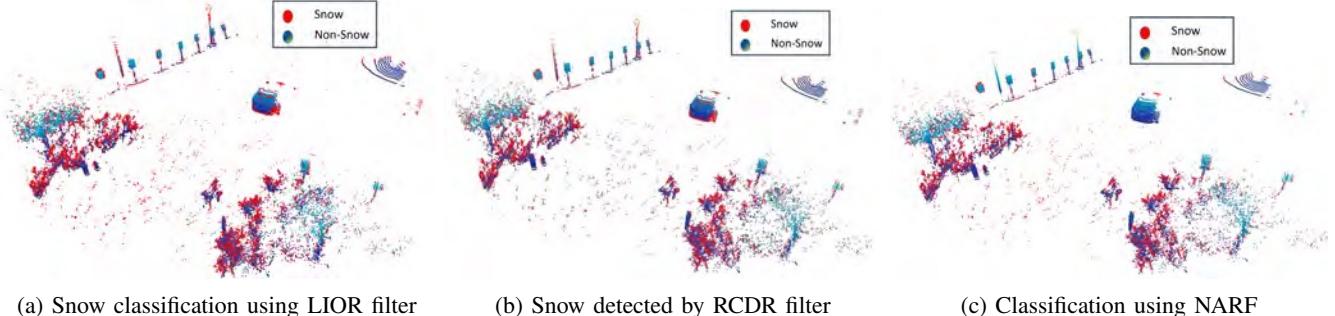


Fig. 6: Comparison of snow particle classification using different methods

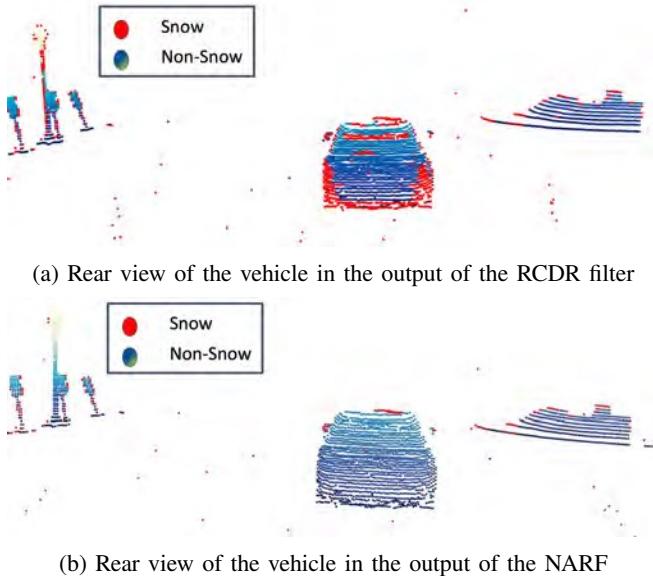


Fig. 7: Visualization of the capability of the RestoreNet in restoring false positives back into the point cloud

The key performance metrics used for comparison are accuracy, precision, recall, and intersection-over-union (IoU). The formulations for these metrics can be found in [26].

Additionally, to compare the time complexity of the filtering methodologies, the Mean Filtering Time (MFT, in seconds) is used. MFT is the average execution time for the filtering process of the point clouds using each of the filtering methods. The compared filtering methodologies are implemented on the whole dataset.

B. Experimental Results

The NARF, the baseline RCDR filter, the LIOR filter [12], and the DROR filter [10] were compared with each other. Table I shows the performance metrics for each filter. The DROR and LIOR filters were tuned using the grid search technique. The standalone RCDR filter and the NARF outperformed the DROR and LIOR filters across multiple key metrics. The DROR filter had the lowest FPR and highest TNR, yielding the highest precision by significantly reducing FPs. However, DROR struggled to correctly detect snow particles, exhibiting low recall and a high FNR. Comparing the overall accuracy, the RCDR filter (95.17%) outperformed

the LIOR filter (94.4%) and the DROR filter (69.10%), while the NARF (98.79%) fared the best. As expected, the NARF had a lower FPR and a higher TNR than the RCDR filter due to the corrective capabilities of RestoreNet. Consequently, it exhibits a 10.47% and 7.09% increase in precision from the RCDR and DROR filters, respectively. Finally, the NARF demonstrates a significantly lower MFT compared to DROR and LIOR. The RCDR has a significantly lower MFT than the NARF as RCDR filtering is the first step of the NARF. This reduction in processing time underscores the high-speed capabilities of the NARF, which is essential to ensure real-time functionality of the filter is possible in autonomous vehicle applications.

C. Bridging Physics and Learning-Based Filters

In this work, we focus primarily on physics-based or physics-informed methods for snow detection due to their high explainability. Black-box learning-based methods have consistently been shown to outperform more explainable methods, but concerns about the generalizability of black-box methods to real-world adverse weather scenarios are worthy of serious consideration. Indeed, in preliminary testing of RestoreNet, when trained directly on the limited dataset considered in this work (without input from the RCDR filter), its accuracy marginally outperformed the NARF. However, physics-informed methods like the RCDR filter maintain important advantages: adaptability to new circumstances, reduced training data requirements, and greater explainability.

The NARF method seeks to leverage the advantages of the physics-based RCDR filter with the high predictive performance and flexibility of RestoreNet. The objective is not for RestoreNet to overpower the influence of the RCDR filter, but rather to augment its weaknesses. For example, reflective surfaces are often falsely classified as snow by the RCDR filter, and the optimal RCDR threshold may change for different types of snow (e.g., particle size). RestoreNet or a similar learning-based method may be able to correct for such factors without overconfidently neglecting the strengths of the RCDR filter as a physically sound baseline for identifying snow particles.

D. Limitations and Prospects for Further Study

Concerns about NARF's generalization to different adverse weather conditions, such as rain or fog, are mitigated by its

reliance on Mie theory, which applies Maxwell's equations assuming spherical particle shapes. Hence, rain, fog, and snow can be treated similarly, except that the back-scatter coefficient β_b might be lower for fog due to smaller particle sizes. Nevertheless, the main limitation of this work is the lack of evidence for such generalization. This is attributed to the limited dataset used in this study and the lack of a diverse open-source dataset. As long-wavelength LiDAR systems gain traction, we encourage the research community to collect and publish more comprehensive datasets, similar to the Winter Adverse Driving dataSet (WADS) [27], using these systems in diverse road environments and driving scenarios. Such data is needed to fully evaluate the novel NARF and RCDR filters proposed in this work against state-of-the-art learning-based methods.

If RestoreNet (or similar learning-based technique) is capable of real-time application in autonomous vehicles, then the overall NARF should also be suitably fast. Although the reported MFTs show NARF to be faster than LIOR, vehicle-in-the-loop tests should be used to validate these expectations and ensure that filter latency falls within acceptable bounds to maintain safe operation.

VI. CONCLUSIONS

In this work, we presented a novel reflectance-based (RCDR) snow particle filter as part of a two-step method for filtering snow from point clouds collected using long-wavelength LiDAR. Experimental results showed the RCDR filter had high recall, while RestoreNet, a new CNN-based model inspired by WeatherNet [14], was used to effectively identify and correct false positives. When tested on a limited experimental dataset, the proposed NARF method filtered snow while conserving points from solid objects, and improved upon existing physics-informed methods both in terms of accuracy and mean filtering time. More comprehensive datasets from long-wavelength LiDAR systems are needed to compare the NARF with state-of-the-art learning-based methods, but similar models could be constructed for short-wavelength LiDAR by replacing the RCDR filter with an appropriate alternative.

ACKNOWLEDGEMENT

This research was funded by the Minnesota Department of Transportation (MnDOT) under project number 00081748 with the significant collaboration of Luminar Technologies Inc. and Vision System Intelligence (VSI) labs.

REFERENCES

- [1] Y. Li and J. Ibanez-Guzman, "Lidar for autonomous driving: The principles, challenges, and trends for automotive lidar and perception systems," *IEEE Signal Processing Magazine*, vol. 37, no. 4, pp. 50–61, 2020.
- [2] S. Warren, "Optical properties of ice and snow," *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences*, vol. 377, 04 2019.
- [3] "Light scattering by small particles. by h. c. van de hulst. new york (john wiley and sons), london (chapman and hall), 1957. pp. xiii, 470; 103 figs.; 46 tables. 96s," *Quarterly Journal of the Royal Meteorological Society*, vol. 84, no. 360, pp. 198–199, 1958.
- [4] M. Byeon and S. W. Yoon, "Analysis of automotive lidar sensor model considering scattering effects in regional rain environments," *IEEE Access*, vol. 8, pp. 102669–102679, 2020.
- [5] M. Berk, M. Dura, J. V. Rivero, O. Schubert, H.-M. Kroll, B. Buschardt, and D. Straub, "A stochastic physical simulation framework to quantify the effect of rainfall on automotive lidar," *SAE International Journal of Advances and Current Practices in Mobility*, vol. 1, no. 2019-01-0134, pp. 531–538, 2019.
- [6] S. Hasirlioglu, I. Doric, C. Lauerer, and T. Brandmeier, "Modeling and simulation of rain for the test of automotive sensor systems," in *IEEE Intelligent Vehicles Symposium (IV)*, pp. 286–291, IEEE, 2016.
- [7] I. I. Kim, B. McArthur, and E. J. Korevaar, "Comparison of laser beam propagation at 785 nm and 1550 nm in fog and haze for optical wireless communications," vol. 4214, pp. 26–37, 2001.
- [8] T. G. Phillips, N. Guenther, and P. R. McAree, "When the dust settles: The four behaviors of lidar in the presence of fine airborne particulates," *Journal of field robotics*, vol. 34, no. 5, pp. 985–1009, 2017.
- [9] R. B. Rusu and S. Cousins, "3D is here: Point Cloud Library (PCL)," in *IEEE International Conference on Robotics and Automation (ICRA)*, (Shanghai, China), IEEE, May 9–13 2011.
- [10] 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)*, pp. 254–261, 2018.
- [11] Y. Duan, C. Yang, and H. Li, "Low-complexity adaptive radius outlier removal filter based on PCA for lidar point cloud denoising," *Appl. Opt.*, vol. 60, pp. E1–E7, Jul 2021.
- [12] J.-I. Park, J. Park, and K.-S. Kim, "Fast and accurate desnowing algorithm for lidar point clouds," *IEEE Access*, vol. 8, pp. 160202–160212, 2020.
- [13] B. Li, J. Li, G. Chen, H. Wu, and K. Huang, "De-snowing LiDAR Point Clouds With Intensity and Spatial-Temporal Features," in *2022 International Conference on Robotics and Automation (ICRA)*, (Philadelphia, PA, USA), pp. 2359–2365, IEEE, May 2022.
- [14] 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, pp. 2514–2521, apr 2020.
- [15] M.-Y. Yu, R. Vasudevan, and M. Johnson-Roberson, "Lisnownet: Real-time snow removal for lidar point cloud," 2022.
- [16] A. Seppanen, R. Ojala, and K. Tammi, "4denoisenet: Adverse weather denoising from adjacent point clouds," *IEEE Robotics and Automation Letters*, vol. 8, p. 456–463, Jan. 2023.
- [17] G. Bae, B. Kim, S. Ahn, J. Min, and I. Shim, "Slide: Self-supervised lidar de-snowing through reconstruction difficulty," 2022.
- [18] L. Zhan and W. F. Northrop, "Impact of fog particles on 1.55 μm automotive lidar sensor performance: An experimental study in an enclosed chamber," in *SAE technical paper series.*, vol. 1 of *SAE WCX Digital Summit*, (Warrendale, PA, USA), Society of Automotive Engineers.
- [19] C. Weitkamp, "Lidar, range-resolved optical remote sensing of the atmosphere," Berlin: Springer, 2005., vol. 102, 01 2005.
- [20] P. Laven, "Mieplot (a computer program for scattering of light from a sphere using mie theory & the debye series)," <http://www.philiplaven.com/mieplot.htm>, 2011.
- [21] G. Koh, "Physical and optical properties of falling snow," 1989.
- [22] R. S. Sekhon and R. C. Srivastava, "Snow size spectra and radar reflectivity," *Journal of Atmospheric Sciences*, vol. 27, no. 2, pp. 299–307, 1970.
- [23] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," 2017.
- [24] F. Piewak, P. Pinggera, M. Schäfer, D. Peter, B. Schwarz, N. Schneider, M. Enzweiler, D. Pfeiffer, and J. Zöllner, "Boosting LiDAR-Based Semantic Labeling by Cross-modal Training Data Generation," Munich, Germany, 2018, Proceedings, Part VI, pp. 497–513, 01 2019.
- [25] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2017.
- [26] A. Diab, R. Kashef, and A. Shaker, "Deep learning for lidar point cloud classification in remote sensing," *Sensors*, vol. 22, no. 20, p. 7868, 2022.
- [27] A. Kurup and J. Bos, "The Winter Adverse Driving dataSet (WADS)," 2021. Retrieved from: <https://digitalcommons.mtu.edu/wads/>.