Overcoming Catastrophic Forgetting in Radar and LiDAR Object Detection in Rain via Layer Freezing and Data Augmentation

Richard Capraru[1,2], Jia-Yu Wu[1], Jian-Gang Wang[2], Matthew Ritchie[3], Emil C. Lupu[4] and Boon Hee Soong[1]

***Abstract*— Advanced Driver-Assistance Systems (ADAS) use sensors like radar, LiDAR, and cameras for reliable vehicle perception in different weather conditions. While LiDAR and cameras offer high-resolution perception in clear weather, radar excels in adverse conditions such as low light, fog, or rain. Adapting systems trained on clear-weather data to cope with adverse weather often causes catastrophic forgetting, significantly reducing their initial performance after re-training. Unsupervised domain adaptation (UDA) techniques aim to address this but are complex. In this paper, we examine catastrophic forgetting effects on radar and LiDAR, proposing methods to reduce it: model freezing, pre-training with mixed data, and adding simulated data. Our experiments on the well-established RADIATE dataset show these methods improve clear-weather retention and rain detection, with radar showing a 6.59% reduction in forgetting and a 17.19% rain detection gain, and LiDAR a 13.62% reduction in forgetting and 24% improvement with simulations.**

I. Introduction

Autonomous vehicles rely on sensors such as radar, LiDAR, and cameras for perception and navigation. Among these sensors, radar and LiDAR are crucial for depth perception.

Recent advances in radar and LiDAR-based object detectors introduce powerful detection strategies [1]–[4]. Whilst effective, these architectures are specific to one modality, and their applicability to other sensors is limited. However, by transforming LiDAR data into bird's-eye view (BEV) representations and radar data into Cartesian images, both data types can be processed with image detection models such as Faster R-CNN [5], yielding robust detection capabilities and enabling direct comparisons between sensors.

However, adverse weather conditions, such as rain, pose significant challenges for these sensing systems. For instance, Vargas et al. [6] showed that LiDAR signals are impacted by rain droplets through Mie scattering, especially for wavelengths commonly used in autonomous vehicles (e.g., 905 nm and 1550 nm). Goodin et al. [7] developed a rain impact model validated through climate chamber experiments conducted by BMW [8], which shows that rain substantially increases data sparsity and diminishes long-distance detection accuracy. Radar, on the other hand, is more resilient to weather effects. [9]. However, using radar alone is not sufficient. Its low spatial resolution and lack of elevation data compromise the ability of radar to capture fine object details and accurately detect small objects such as pedestrians [10].

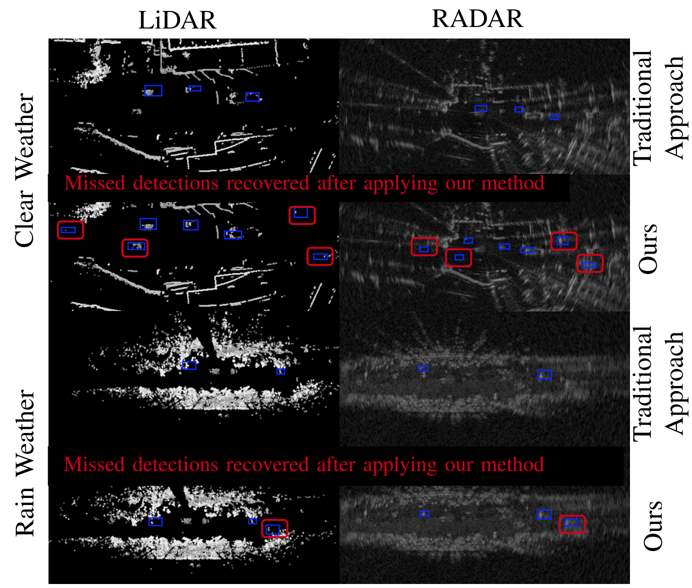

Fig. 1: Comparison of the results between traditional and our approach. In contrast to traditional methods, which pre-train on clear-weather data and fine-tune on rain data, our approach combines layer freezing with mixed data training. This avoids long-range detection failures due to forgetting.

Most radar and LiDAR detectors are evaluated in clear weather, requiring incremental learning for adverse conditions. This risks catastrophic forgetting, where fine-tuning on new data (e.g., rain) degrades performance on prior data (e.g., clear weather). While extensively studied in image processing [11], language models [12], and recently LiDAR segmentation [13], its impact on radar and LiDAR object detection remains underexplored. This paper addresses catastrophic forgetting in radar and LiDAR detection, proposing methods to mitigate it and enhance detection in rain. A qualitative example is shown in Figure 1.

The main contributions of this paper are:

- **New Knowledge.** To the best of our knowledge, this work is the first to study the impact of catastrophic forgetting on radar/LiDAR object detectors under rainy conditions.
- **New fine-tuning strategy.** We propose a fine-tuning strategy using selective layer freezing to retain clear-weather knowledge and reduce catastrophic forgetting. Our approach improves detection for networks trained on both clear-weather and mixed-weather data.
- **Data augmentation.** We explore the use of simulated data from both clear and rainy conditions to improve

[1] School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore 639798
[2] Institute for Infocomm Research (I^2R), A*STAR, Singapore 138632
[3] Electronic and Electrical Engineering Department, University College London, UK WC1E 7JE
[4] Department of Computer Science, Imperial College London, UK SW7 2AZ

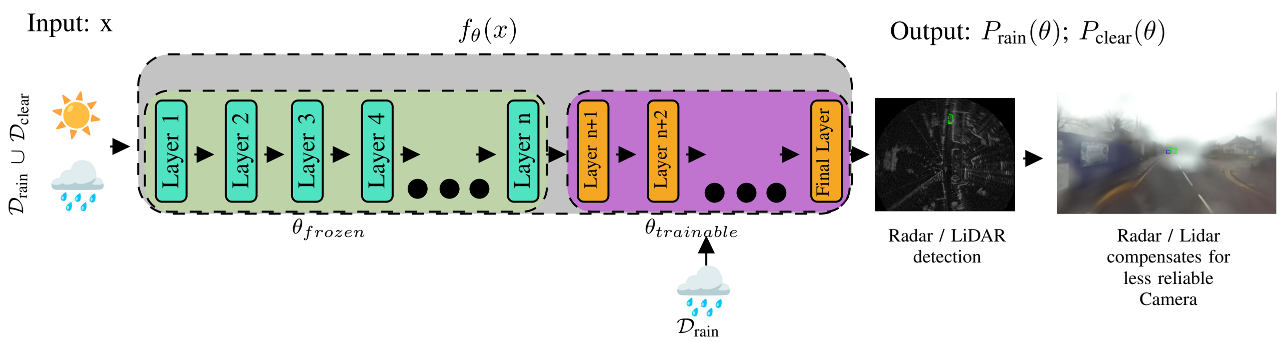

Fig. 2: Overview of our Methodology. The model is first trained on a combined dataset of clear and rainy weather ($\mathcal{D}_{\text{clear}} \cup \mathcal{D}_{\text{rain}}$) to learn robust shared features. During adaptation to rain, the initial layers (θ_{frozen}) are kept fixed while only the later layers ($\theta_{\text{trainable}}$) are fine-tuned with rain-specific data ($\mathcal{D}_{\text{rain}}$). This strategy preserves clear-weather performance while improving detection in rain. Simulated weather data can optionally supplement real samples. The resulting predictions $P_{\text{rain}}(\theta)$ and $P_{\text{clear}}(\theta)$ are used to evaluate detector robustness. This method allows radar/LiDAR to remain effective in rain, where cameras typically degrade.

LiDAR model robustness and generalization.

- **Performance Enhancement.** We demonstrate how combining layer freezing and data augmentation leads

to improvements in both clear and rainy weather object detection on a well-established dataset, RADIATE [14].

The paper is structured as follows: Section II covers the problem formulation and methodology, while Section III details the experimental setup, dataset, and evaluation metrics. Section IV examines rain's impact and forgetting in RADAR/LiDAR detection, followed by a performance comparison of mitigation methods in Section V. Finally, Section VI presents the conclusion.

II. Methodology

To address catastrophic forgetting in radar and LiDAR detection, we explore fine-tuning on rain data, selective layer freezing, and incorporating simulated data at different training stages. These methods balance performance across weather conditions while reducing forgetting. Figure 2 provides an overview of our approach.

A. Problem formulation

1) 2D Object Detection: Consider a neural network $f_\theta(x)$ with parameters θ that takes as input an image $x \in \mathbb{R}^{H \times W \times C}$, where H and W represent the height and width of the image, and C is the number of channels (e.g., RGB channels). The model is trained to predict a bounding box and label, $y \in \mathcal{Y}$, where $\mathcal{Y}$ denotes the set of possible target outputs such as classes and locations.

In our context, we define.

- $\mathcal{D}_{\text{clear}} = \{(x_i, y_i)\}_{i=1}^{N_{\text{clear}}}$ denotes the dataset of images with annotations representing bounding box locations and labels of targets in clear weather conditions.
- $\mathcal{D}_{\text{rain}} = \{(x_i, y_i)\}_{i=1}^{N_{\text{rain}}}$ denotes a dataset of images with annotations representing bounding box locations and labels of targets in rainy conditions.

The parameter vector θ represents all trainable weights in the network, and $f_\theta(x)$ represents the network's prediction for a given input x. Our goal is to optimize θ such that the network achieves high performance on both clear and rainy weather data, even as the network is successively fine-tuned on data for new weather conditions.

2) Catastrophic Forgetting: When a model trained on one dataset (e.g., clear weather) is fine-tuned on another dataset (e.g., rainy weather), it often experiences a phenomenon called *catastrophic forgetting*. When this happens, the model's performance on the initial dataset deteriorates after fine-tuning on the new dataset, as the model's parameters θ adapt specifically to the new data. To quantify catastrophic forgetting, let us define:

- $P_{\text{clear}}(\theta)$: Average precision (Ap) of $f_\theta(x)$ on clear $\mathcal{D}_{\text{clear}}$.
- $P_{\text{rain}}(\theta)$: Average precision (Ap) of $f_\theta(x)$ on rain $\mathcal{D}_{\text{rain}}$.

After the initial training on $\mathcal{D}_{\text{clear}}$, the initial performance on clear data, is denoted as $P_{\text{clear}}^{(\text{initial})} = P_{\text{clear}}(\theta_{\text{clear}})$. When fine-tuning the model on $\mathcal{D}_{\text{rain}}$, the parameters become $\theta_{\text{fine-tune}}$, leading to a different performance on the respective datasets. The forgetting can then be defined as follows:

$$\Delta P_{\text{clear}} = P_{\text{clear}}^{(\text{initial})} - P_{\text{clear}}(\theta_{\text{fine-tune}}) \tag{1}$$

where ΔP_{clear} quantifies the loss in performance on clear data due to fine-tuning on rain data. Our goal is to minimize ΔP_{clear}, thus controlling catastrophic forgetting, while still improving $P_{\text{rain}}(\theta)$ to adapt effectively to rainy conditions.

3) Our Multi-Objective Optimization Problem: involves the following objectives:

1) Maximize performance on rainy data, $P_{\text{rain}}(\theta)$.
2) Minimize the degradation in Ap on clear data, ΔP_{clear}.

Our task is a multi-objective optimization problem:

$$\max_{\theta} \left(P_{\text{rain}}(\theta), P_{\text{clear}}(\theta), -\Delta P_{\text{clear}}(\theta)\right) \tag{2}$$

where maximizing $P_{\text{rain}}(\theta)$ improves performance on rain data, maximizing $P_{\text{clear}}(\theta)$ maintains a robust performance on clear data, and minimizing ΔP_{clear} mitigates catastrophic forgetting on clear data.

B. Our approach

To reduce forgetting for both radar and LiDAR, we use two strategies: pre-training on combined data and layer freezing during fine-tuning. In addition, we explore the use of simulated data to augment the LiDAR dataset.

TABLE I: Sensor Specifications, number of images used during our experiments and the specific data sets from RADATE used.

	Sensor	RADAR	LiDAR
	Model	Navtech CTS350-X	Velodyne HDL-32e
	Coverage	360° Horizontal	360° Horizontal 41° Vertical
	Range	100 m	100 m
	Range resolution	17.5 cm	<20mm
Pre-training Data	Clear Weather	1432 images (e.g. City_1_1,City_6_0, City_1_3)	
	Clear+Rain Weather	1432 images (e.g. City_1_1,City_6_0,Rain_4_1)	
Fine-tuning Data	Rain Weather	719 images (e.g. Rain_4_0)	
Test Data	Clear Weather	719 images (e.g. City_3_7)	
	Rain Weather	718 images (e.g. Rain_3_0)	

1) Mixed Data training: To establish a balanced feature representation across clear and rainy conditions, we initialize the model parameters θ by pre-training on a combined dataset $\mathcal{D}_{\text{combined}} = \mathcal{D}_{\text{clear}} \cup \mathcal{D}_{\text{rain}}$. This yields an initial parameter setting $\theta^*_{\text{combined}}$, given by:

$$\theta^*_{\text{combined}} = \arg\max_{\theta} \left(\alpha P_{\text{clear}}(\theta) + (1-\alpha) P_{\text{rain}}(\theta)\right) \quad (3)$$

where $\alpha \in [0, 1]$ balances the model's performance on clear and rain data during pre-training.

2) Layer Freezing: Following pre-training, we fine-tune only the final layers of the network on $\mathcal{D}_{\text{rain}}$, freezing the initial layers. This strategy prevents substantial modification to the network's representations learned for clear data, helping to retain performance on clear-weather data. We denote the parameters as $\theta = (\theta_{\text{frozen}}, \theta_{\text{trainable}})$ where:

- θ_{frozen}: are the parameters of the frozen layers and fixed during fine-tuning.
- $\theta_{\text{trainable}}$: are the parameters of the layers allowed to update during fine-tuning.

Thus, we solve the following optimization problem:

$$\max_{\theta_{\text{trainable}}} P_{\text{rain}}(\theta_{\text{frozen}}, \theta_{\text{trainable}}) \quad (4)$$

where the fixed θ_{frozen} layers retain the features learned for clear weather, retaining performance on clear data.

3) Incorporating Simulated Data: In addition to pre-training and layer freezing, we introduce a third approach for LiDAR: incorporating simulated data from clear or rainy weather conditions into the pre-training or fine-tuning stages. The simulated data, derived from either $\mathcal{D}_{\text{clear}}$ or $\mathcal{D}_{\text{rain}}$, serves to reinforce the network's learning on specific weather.

We systematically test various combinations of simulated data, adding it to either:

- **Pre-training stage**: Augmenting the combined dataset $\mathcal{D}_{\text{combined}}$ with simulated samples to enhance feature diversity for both clear and rainy conditions.
- **Fine-tuning stage**: Adding simulated data to $\mathcal{D}_{\text{rain}}$ during fine-tuning, enabling the network to generalize better to rain while preserving clear-weather features.

Through empirical evaluation, we assess which combination of simulated data additions optimally balances improvements in P_{rain} with minimal degradation in P_{clear}.

III. EXPERIMENTAL SETUP AND DATASET

A. Data

1) RADIATE Dataset: RADIATE [14] is an automotive dataset developed by Heriot-Watt University, designed to capture diverse weather conditions. It includes data collected across various scenarios such as clear weather and adverse weather (e.g. rain, fog etc). The dataset contains radar images, LiDAR point clouds, and camera images, with annotations for various road objects such as vehicles (cars, vans, trucks, buses, motorbikes, bicycles) and pedestrians. For our analysis, we focus on the rain and clear weather (sunny/cloudy in the city) vehicle detection, leveraging the LiDAR and radar systems for their robustness under varying lighting and weather conditions. Details about the sensor models and the number of images used in our training and fine-tuning experiments are provided in Table I.

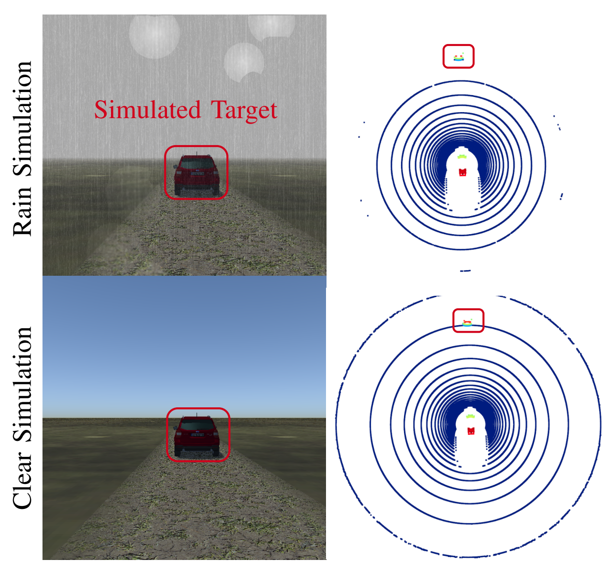

Fig. 3: Example simulated target in rainy and clear weather

2) Simulation Data: To explore using simulated LIDAR data to reduce catastrophic forgetting, we generated target cars at various angles and distances using the MAVS simulator [7]. A Velodyne HDL-32E sensor was positioned 2.01 meters above the host vehicle (AV) to match the RADIATE dataset setup precisely. We simulated a car in front of the AV, aligned parallel, at distances from 20 to 60 meters with increments of 0.5 or 1 meter, producing two sets of 80

TABLE II: Validating our study on Layer Freezing and Mixed Data Training for reducing forgetting for Radar and LiDAR; Exp. No. = Experiment Number

Exp. No.	Sensor	Training Data	Use Fine-Tuning?	Freezing layers?	Validation			
					$P_{\text{clear}}(\theta)$	$P_{\text{rain}}(\theta)$	ΔP_{clear}	Improvement in ΔP_{clear} compared to traditional fine-tuning
(1)	**RADAR**	Clear Weather	-	-	28.54%	14.33%	-	-
(2)			✓	✓	13.52%	20.21%	15.02%	3.19%
(3)		Clear+Rain Weather	-	-	36.71%	20.84%	-	-
(4)			✓	✓	25.09%	31.52%	11.62%	**4.17%**
(5)	**LiDAR**	Clear Weather	-	-	67.64%	43.19%	-	-
(6)			✓	✓	40.35%	56.69%	27.29%	**12.27%**
(7)		Clear+Rain Weather	-	-	76.44%	48.18%	-	-
(8)			✓	✓	45.03%	62.54%	31.41%	5.87%

simulations each (40 in clear weather and 40 with a rain rate of 10 mm/h), one for pre-training and one for fine-tuning, respectively. Figure: 3 illustrates an example of a simulated car in clear and rain conditions. In the final experiments, we expanded the dataset by introducing additional angle offsets (e.g., 10° and 350°) and rain rates (e.g., 2, 5, and 15 mm/h), creating a total of 720 simulated samples.

Rain Model. The theoretical model used in this simulator can be found in [7]. The model can be summarized as the equation of the relative intensity returned by the LiDAR as a function of rainfall rate and is defined as follows:

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

where z is the range and R is the rainfall rate in mm/h. The returned points with a power/intensity less than the value defined by (5) are removed to simulate rain data sparsity.

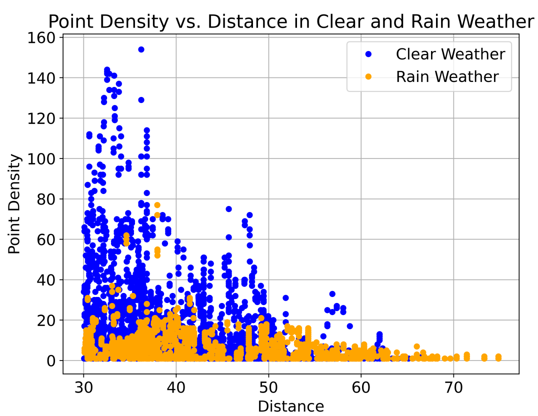

Fig. 4: LiDAR Point Density vs. Distance (meters) in Clear and Rain Weather for cars

B. Detector

Faster R-CNN [5] is an advanced object detection model that combines a Region Proposal Network (RPN) for generating region proposals and a Fast R-CNN detector [15] for classifying and refining these proposals. For this paper, we used Faster-RCNN with a ResNet50 [16] backbone.

1) LiDAR pre-processing: To fairly compare radar and LiDAR, we converted LiDAR point clouds into BEV images to match radar's Cartesian format. However, LiDAR data became increasingly sparse with range, making long-range radar annotations unusable without introducing false positives. We analyzed LiDAR point density in relation to range under both clear and rainy weather conditions. As shown in Figure 4, rain significantly reduces point density with distance, with target points nearly absent beyond 60 meters. Consequently, we filtered out empty annotations and those beyond 60 meters to prevent training and validation on false positives.

C. Evaluation Metrics

We evaluate each model on the validation sets of clear and rain datasets using Average Precision (Ap), a standard metric for object detection. Specifically, we use:

- $P_{\text{rain}}(\theta)$: Ap on the rainy validation set.
- $P_{\text{clear}}(\theta)$: Ap on the clear validation set.
- ΔP_{clear}: forgetting measure, calculated as the difference between the Ap of the initial model and the model fine-tuned on rain data.

IV. Impact of Rain and Forgetting on Radar/LiDAR Object Detection

In order to define the problem, we aimed to answer the question: *What is the impact of forgetting on radar and LiDAR?*. To investigate, we first trained the radar and LiDAR detectors on clear-weather data, then fine-tuned them with data from adverse weather conditions. Throughout training and fine-tuning, we monitor the loss to prevent overfitting and select the epoch with the best test performance. As shown in Table II, experiments (1)–(4) introduce radar experiments: (1) is trained only on clear-weather data; (2) is trained on clear-weather and then fine-tuned with layer freezing on rain; (3) is trained on mixed clear and rain data; (4) is trained on mixed data and then fine-tuned with layer freezing on rain. Experiments (5)–(8) repeat this design for LiDAR. To quantify the impact, the original detection accuracy in clear-weather was 28.54% for radar (Exp. (1)) and 67.64% for LiDAR (Exp. (5)). After standard fine-tuning, detection in clear weather dropped to 10.33% for radar ($\Delta P_{\text{clear}} = 18.21\%$) and to 27.98% for LiDAR ($\Delta P_{\text{clear}} = 39.66\%$). These results underscore the severity of the forgetting issue, particularly for the LiDAR system.

V. Results and Discussion

After identifying the problem, we focused on the key question: *How can we reduce forgetting and maintain robust performance in clear weather during fine-tuning on rain?*

A. Mixed Data Training and Layer Freezing

We first investigate the optimal number of layers to freeze, denoted as θ_{frozen}. In this experiment, we train the network using either clear weather data or a combination of clear

TABLE III: Study on using simulations for LiDAR object detection; CW = Clear Weather; RW = Rain Weather; Exp. No. = Experiment Number

Exp. No.	**Training**				**Fine-Tuning**					**Validation**		
	Real Data	**Simulated Data**			**Use Fine-tuning?**	**Real Data**	**Simulated Data**					
		None	**CW 40 images**	**RW 40 images**		**RW 719 images**	**None**	**CW 40 images**	**RW 40 images**	$P_{\text{clear}}(\theta)$	$P_{\text{rain}}(\theta)$	ΔP_{clear}
(7)	**CW+RW 1432 images**	✓	-	-	-	-	-	-	-	76.44%	48.18%	-
(9)		-	✓	-	-	-	-	-	-	79.48%	46.25%	-
(10)		-	-	✓	-	-	-	-	-	75.26%	48.37%	-
(11)		-	✓	✓	-	-	-	-	-	**80.19%**	49.91%	-
(8)		✓	-	-	✓	✓	✓	-	-	45.03%	62.54%	31.41%
(12)		-	✓	-	✓	✓	✓	-	-	48.98%	62.46%	30.50%
(13)		-	-	✓	✓	✓	✓	-	-	44.17%	66.52%	31.09%
(14)		-	✓	✓	✓	✓	✓	-	-	49.32%	65.79%	30.87%
(15)		✓	-	-	✓	✓	-	✓	-	47.11%	61.52%	29.33%
(16)		✓	-	-	✓	✓	-	-	✓	45.69%	66.58%	30.75%
(17)		✓	-	-	✓	✓	-	✓	✓	49.38%	66.64%	27.06%
(18)		-	✓	✓	✓	✓	✓	-	-	49.32%	65.79%	30.87%
(19)		-	✓	✓	✓	✓		✓	-	53.47%	62.64%	26.72%
(20)		-	✓	✓	✓	✓		-	✓	50.81%	67.59%	29.38%
(21)		-	✓	✓	✓	✓		✓	✓	54.25%	68.72%	**25.94%**
(22)		-	✓	✓	✓	✓	Adding more rain rates and angles			60.62%	**72.19%**	-

and rain weather data, while progressively freezing the first θ_{frozen} layers. As shown in Figure 5, we find that $\theta_{\text{frozen}} = 15$ strikes the best balance, improving average precision (Ap) for rain conditions while maintaining reasonable performance in clear weather for both radar and LiDAR systems. This result holds for both networks trained on domain-specific data (clear weather) and those trained on mixed (clear and rain) weather data highlighting the improvements layer freezing offers. This improvement is attributed to the fact that catastrophic forgetting primarily affects low-level features, and freezing layers helps preserve these features. Additionally, pre-training on both source and target domains (clear and rain weather) enables the network to learn generalized features, mitigating catastrophic forgetting and enhancing continuous learning.

After identifying the optimal number of frozen layers ($\theta_{\text{frozen}} = 15$, in our experiments), we evaluated object detection performance for both clear and rainy weather conditions, comparing networks pre-trained with either domain-specific data (clear weather only) or mixed data (clear and rain). As introduced earlier, the results for radar and LiDAR are summarized in Table II.

Radar: With mixed-data pre-training (Exp. (4)), our layer-freezing strategy reduced forgetting, achieving a clear-weather performance drop of $\Delta P_{\text{clear}} = 11.62\%$, compared to 18.21% with traditional fine-tuning (no freezing, $P_{\text{clear}} = 20.92\%$). This constitutes a 6.59% improvement. Additionally, detection performance in rainy conditions improved by 17.19% compared to a baseline trained only on clear weather data (Exp. (1)).

LiDAR: Layer freezing further reduced forgetting, with a clear-weather performance drop of $\Delta P_{\text{clear}} = 27.29\%$ (Exp. (6)) compared to 39.56% using traditional fine-tuning without freezing ($P_{\text{clear}} = 28.08\%$), yielding a 12.27% improvement. The strongest reduction in forgetting was observed when pre-training exclusively with clear-weather data, likely because initial performance on clear weather was comparatively lower, thus amplifying the benefits of freezing. Importantly, LiDAR performance in rainy conditions improved by 19.35% when pre-trained with mixed data (Exp. (8)) compared to clear-weather-only training (Exp. (5)).

Based on these results, we selected the LiDAR architecture pre-trained on mixed data (clear and rain) for subsequent simulated experiments. This choice provides the highest overall accuracy: specifically, 4.68% higher in clear weather and 5.85% higher in rainy conditions compared to clear-weather-only pre-training (Exp. (8) versus Exp. (6)).

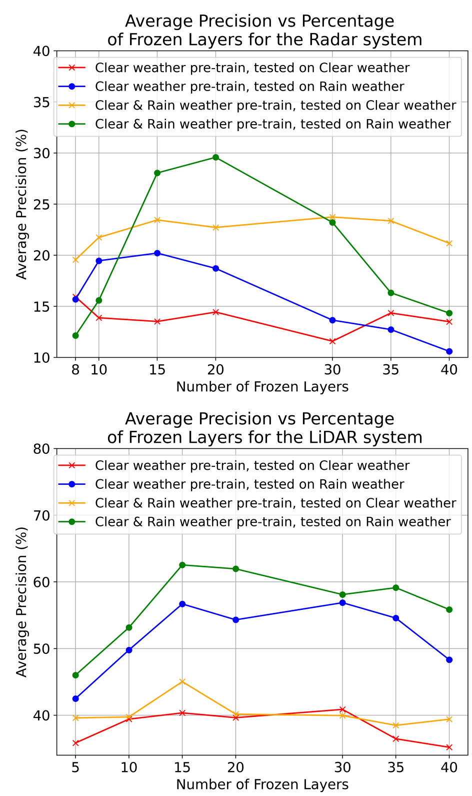

Fig. 5: Average Precision on Clear and Rain Weather vs Number of Frozen Layers for (up) the Radar system and (down) the LiDAR System

B. Data Augmentation

LiDAR object detection faces challenges due to data sparsity for critical targets and difficulty in capturing reliable long-distance data. To mitigate these, we used a simulated dataset from MAVS [17] to generate targets at 20–60 meters in both clear and rainy conditions. We systematically investigated the effects of incorporating small volumes of simulated data during pre-training, fine-tuning, or both stages (Experiments (9)–(22)), summarized in Table III). Specifically: Experiments (9)–(11) incorporate simulated clear, rain, or both data during pre-training; (12)–(14) add simulated data at pre-training followed by real rain data for fine-tuning; (15)–(17) introduce simulated data only during fine-tuning, with pre-training on real data; (18)–(20) include simulated data in both pre-training and fine-tuning; (21) employs both simulated clear and rain data throughout all stages; and (22) assesses the effect of simulated data spanning a wide range of rain intensities and incidence angles.

Results have shown that using simulated data at both pre-training and fine-tuning stages (Exp. (21)) notably reduced forgetting, lowering the clear-weather performance drop (ΔP_{clear}) from 39.56% (traditional fine-tuning) to 25.94%—an improvement of 13.62%. Furthermore, augmenting simulations across diverse rain rates and angles (Exp. (22)) significantly boosted rainy-weather detection by 24% compared to the baseline (Exp. (7)), with a clear-weather performance loss of 15%, half of that observed without simulations (performance drop in clear weather in Exp. (8) versus Exp. (7)).

Even minimal use of simulated data (less than 6% of the total training set), clear-weather and rainy-weather (Exp. (11)), yielded clear-weather gains of 3.75% and rainy-weather gains of 1.73% compared to the baseline (Exp. (7)), highlighting the benefit of increased data variation in enhancing model robustness.

Distance Performance Improvements. To deepen our analysis, we evaluated radar and LiDAR performance improvements following our optimal methods: layer freezing combined with mixed-data training for radar, and a combination of layer freezing, mixed-data, and simulation training for LiDAR. These results were compared against the network trained solely on clear-weather data. Our experiments reveal an improvement for detection in rain, with a 29% gain for LiDAR and a 28.86% gain for radar. For most distances, the radar system showed the highest accuracy gains after applying our method, reaching a peak improvement of 31.08% at 30–40 meters. In comparison, LiDAR achieved its best improvement of 20.36% at 20–30 meters. This difference may stem from the greater domain shift in radar data relative to LiDAR BEV images, as LiDAR loses significant spatial and intensity details when converted into images.

Limitation. While these comparisons provide valuable insights, a direct comparison between LiDAR and radar systems is still limited. Due to inherent data sparsity in LiDAR, we filtered empty annotations for training and validation—annotations still present in radar analysis—which may lead to the perception that LiDAR performs better under rainy conditions. Furthermore, the MAVS rain simulator, validated in prior work [18], requires further real-world testing to assess its accuracy.

VI. Conclusions and Future work

One of the core challenges in deep learning-based object detection is catastrophic forgetting, which hinders the model's ability to retain performance across sequential tasks or datasets. In this paper, we propose an approach to improve object detection in rainy conditions while preserving performance in clear-weather. We achieve this through model freezing, pre-training with mixed data (clear and rain), and adding simulated data. This helps maintain detectability for clear data when adapting the detector to rain. Experimental results on a publicly available dataset validate our approach. In the future, we could compare catastrophic forgetting and rain detection across sensor modalities and test our methods on other networks and datasets.

Acknowledgements

The authors extend their deepest gratitude to Dr. Teoh Eam Khwang for his insightful suggestions during the writing of this paper. Assoc. Prof. B.H. Soong gratefully acknowledges the support of MOE Grant RG66/22. We also wish to thank the anonymous reviewers for their valuable feedback.

References

[1] Y. Zhou, L. Liu, H. Zhao, M. López-Benítez, L. Yu, and Y. Yue, "Towards deep radar perception for autonomous driving: Datasets, methods, and challenges," *Sensors*, vol. 22, no. 11, 2022.

[2] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, "Pointpillars: Fast encoders for object detection from point clouds," in *CVPR*, 2019, pp. 12 689–12 697.

[3] S. Shi, X. Wang, and H. Li, "Pointrcnn: 3d object proposal generation and detection from point cloud," in *CVPR*, 2019, pp. 770–779.

[4] Y. Chen, S. Liu, X. Shen, and J. Jia, "Fast point r-cnn," in *ICCV*, 2019, pp. 9774–9783.

[5] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: towards real-time object detection with region proposal networks," in *NeurIPS*, ser. NIPS'15. Cambridge, MA, USA: MIT Press, 2015, p. 91–99.

[6] J. Vargas, S. Alsweiss, O. Toker, R. Razdan, and J. Santos, "An overview of autonomous vehicles sensors and their vulnerability to weather conditions," *Sensors (Basel, Switzerland)*, vol. 21, 08 2021.

[7] 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.

[8] R. Rasshofer, M. Spies, and H. Spies, "Influences of weather phenomena on automotive laser radar systems," *Advances in Radio Science*, vol. 9, 07 2011.

[9] V. Sharma and S. Sergeyev, "Range detection assessment of photonic radar under adverse weather perceptions," *Optics Communications*, vol. 472, p. 125891, 2020.

[10] T. Akita and S. Mita, "Object tracking and classification using millimeter-wave radar based on lstm," in *ITSC*, 2019, pp. 1110–1115.

[11] T. Kalb and J. Beyerer, "Principles of forgetting in domain-incremental semantic segmentation in adverse weather conditions," in *CVPR*, 2023, pp. 19 508–19 518.

[12] Y. Yu, C.-H. H. Yang, J. Kolehmainen, P. G. Shivakumar, Y. Gu, S. R. R. Ren, Q. Luo, A. Gourav, I.-F. Chen, Y.-C. Liu, T. Dinh, A. G. D. Filimonov, S. Ghosh, A. Stolcke, A. Rastow, and I. Bulyko, "Low-rank adaptation of large language model rescoring for parameter-efficient speech recognition," in *IEEE ASRU*, 2023, pp. 1–8.

[13] J. Zhou, J. Mei, P. Wu, L. Chen, F. Zhao, X. Zhao, and Y. Hu, "Teff: Tracking-enhanced forgetting-free few-shot 3d lidar semantic segmentation," in *2024 IROS*, 2024, pp. 14 103–14 110.

[14] M. Sheeny, E. De Pellegrin, S. Mukherjee, A. Ahrabian, S. Wang, and A. Wallace, "Radiate: A radar dataset for automotive perception in bad weather," in *ICRA*. IEEE Press, 2021, p. 1–7.

[15] R. Girshick, "Fast r-cnn," in *ICCV*, 2015, pp. 1440–1448.

[16] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, 2016, pp. 770–778.

[17] "MSU autonomous vehicle simulator," accessed: 2023-01-29. [Online]. Available: https://www.cavs.msstate.edu/capabilities/mavs.php

[18] R. Capraru, J.-G. Wang, and B. H. Soong, "Upsampling data challenge: Object-aware approach for 3d object detection in rain," in *International Conference on Advanced Concepts for Intelligent Vision Systems*, 2023, pp. 148–159.