2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
October 14-18, 2024. Abu Dhabi, UAE

Perception for Connected Autonomous Vehicles under Adverse Weather Conditions

Dimitra Tsakmakopoulou, Konstantinos Moustakas, Senior Member, IEEE
Electrical and Computer Engineering Department, University of Patras, Greece

Abstract— **Autonomous Vehicles (AVs) have recently attracted considerable attention due to their potential to significantly reduce road accidents and improve people's lives. However, they rely solely on the data collected by their mounted sensors to make predictions, which can lead to inaccurate results if a sensor becomes occluded or damaged. This issue can be addressed by employing Vehicle-to-Vehicle communication, which allows a Connected Autonomous Vehicle (CAV) to interact with other CAVs within its field of view and exchange information about their surrounding objects. Existing research on cooperative perception has primarily focused on clear weather scenarios, with limited exploration into adverse weather conditions. This paper demonstrates the necessity of Vehicle-to-Vehicle communication by showcasing its benefits in maintaining high accuracy under adverse weather conditions. A collaborative perception system is introduced and its performance in foggy weather scenarios is assessed to further improve adverse weather perception. The pipeline of the network combines state-of-the-art methods for accurate object detection. Specifically, with PointPillars as the backbone, the Spatial-wise Adaptive Feature Fusion method is used to aggregate information from different vehicles. The model is trained on the large-scale dataset OPV2V and evaluated on modified data to simulate fog. The experiments show that cooperative perception can maintain high detection accuracy even in challenging weather conditions. Finally, a comparative analysis of LiDAR detectors for cooperative perception in bad weather conditions is presented.**

I. Introduction

Autonomous Vehicles (AVs) require an accurate perception of their 3D environment, as false detections can lead to accidents with potentially severe consequences for human safety. Despite recent advancements in machine learning algorithms, the challenge of inaccurate results persists, particularly when surrounding objects are occluded or reduced visibility due to adverse weather degrades sensor recordings. Connected Autonomous Vehicles (CAVs) offer a promising solution to enhance dynamic environmental perception. By facilitating the exchange of data from different angles for the same road scene, CAVs contribute to a more comprehensive understanding of the objects in the environment.

To ensure safety, it is essential to evaluate the performance of AVs under diverse environmental and weather conditions, especially given the recent increase in extreme weather changes. Most existing research has focused on single-vehicle detection in adverse weather conditions. Considering the significant contribution of cooperative perception in clear weather scenarios, this paper aims to evaluate and compare the performance of state-of-the-art (SOTA) LiDAR detection algorithms for cooperative perception in foggy weather. This is achieved by applying a fog simulation method to the first published large-scale connected AVs dataset. In addition, a pipeline is proposed that outperforms other SOTA methods in terms of accuracy. The pipeline uses the S-AdaFusion [1] method for data fusion from different CAVs, which is applied at an earlier stage according to the overall network architecture introduced in [2]. While this pipeline demonstrates higher accuracy, which is critical for when the visibility is reduced and human safety is the primary concern, the research also examines the potential drawbacks that may be introduced of prioritising higher accuracy.

Section II describes current SOTA cooperative perception methods and the impact of adverse weather on AVs. Section III provides a detailed explanation of the proposed pipeline and a brief overview of the selected fog simulation method. The experimental results and comparative analysis are presented in section IV. Section V concludes the paper.

II. Related Work

This section begins with an overview of existing work on LiDAR detectors for cooperative perception. It then examines the impact of adverse weather conditions on point cloud data, highlighting the research gap that this study aims to address.

A. Cooperative Perception

In cooperative perception, vehicles share data with other vehicles in their vicinity. This allows each vehicle to fuse and process data from its onboard sensors with information from other vehicles, extending the individual vehicle's perception range. Data fusion methods can be divided into three categories depending on the type of data being shared: early, late and intermediate fusion. Early fusion methods [3], involve sharing raw data, which the ego vehicle later fuses and processes. However, these methods face challenges in transmitting real-time data over limited network bandwidth. On the contrary, in late fusion methods [4] each agent transmits its proposals which are then fused into a final prediction by the ego vehicle. While this approach reduces the required bandwidth, it does not significantly enhance accuracy, as it depends on each agent's individual performance. Intermediate fusion methods have proven to strike a balance, maintaining low bandwidth requirements while achieving high detection accuracy. In these methods, the ego vehicle fuses feature maps obtained from surrounding vehicles.

For instance, F-Cooper [5] employs max pooling to aggregate the most important features from feature maps of various

CAVs. This straightforward floating-point operation introduces minimal computational overhead. Other operations like summation and average pooling are viable alternatives, although max pooling tends to produce better results. The above methods do not extract features taking into account the correlations between neighbouring features in the feature maps. V2VNet [6] proposes a framework that utilizes a Graph Neural Network (GNN) to create a map of the vehicles based on their geographical positions. In this graph, each vehicle is represented as a node, with each node containing information about the vehicle's feature maps, position, time delay and other relevant data. This method effectively aggregates data from each vehicle and keeps the graph nodes updated. Xu et al. [2] suggest the use of a self-attention mechanism that pays attention to important observations, filters out potential sensor noise, and captures interactions between features of neighbouring connected vehicles. [1] focuses on trainable neural networks capable of selecting features more effectively than the plain reduction operators. They propose Spatial-wise Adaptive feature fusion method that outperforms all other state-of-the-art approaches.

B. Object Detection Frameworks in Autonomous Vehicles

In recent years, numerous methods employing convolutional neural networks (CNNs) for object detection in 3D space have emerged, each utilizing different data sources. Some solely rely on camera data, monocular or stereo images, while others combine them with RGB-D sensors to leverage depth information. Despite their performance, especially in autonomous driving scenarios, it appears that data from LiDAR (Light Detection And Ranging) sensors is preferred due to its higher accuracy [7], which is critical for safety. One of the first notable works in object detection from point clouds for autonomous vehicles is PIXOR [8]. In this approach, the authors project the point cloud data onto a plane, transforming it into a 2D representation known as the bird's eye view (BEV). Another significant contribution, PointNet [9], proposes a method for learning from unordered point clouds. In VoxelNet [10], the three-dimensional space is divided into uniform voxels, and the points are grouped based on the voxel in which they reside. Subsequently, a PointNet is applied to each voxel, supplemented by convolutional middle layers. Despite its impressive performance, VoxelNet is too slow for real-time applications. The SECOND approach [11] addresses this challenge by utilizing sparse convolutional networks to reduce inference time and memory requirements. Nevertheless, the use of 3D convolutions remains a limitation. In [12], the authors managed to overcome this problem by proposing the PointPillars method, which converts the point cloud into a pseudo-image and applies a 2D convolutional backbone. This approach significantly reduces inference time compared to previous methods.

C. Autonomous Vehicles on Adverse Weather Conditions

Adverse weather conditions have a significant impact on data collected by LiDAR sensors, as atmospheric particles such as raindrops, fog and snowflakes can lead to unwanted reflections, refractions and absorptions of laser beams. This results in a point cloud with missing or deformed points. Unlike human drivers who rely on instinct and experience when navigating complex weather scenarios, autonomous vehicles depend solely on sensor data, making them susceptible to both false positive and false negative object detections.

In [13] Peynot et al. conducted a pioneering study, collecting a multi-modal dataset that included dust, snow and rain. Their analysis revealed significant attenuation of lidar sensors caused by dust particles, at times leading to objects disappearing behind airborne dust. A performance evaluation of popular automotive lidar systems under challenging weather conditions (fog and rain) was performed in [14] and [15]. These experiments were conducted in a controlled climate chamber, capable of generating stable fog with a certain meteorological visual range in two different fog droplet distributions. Moreover it provided a rain simulator with a stabilized rainfall rate. The comprehensive analysis of object perception in these conditions showed a notable reduction in the number of points per object and decreased variance of the measured intensity. In [16] the authors introduced the first dataset containing data collected during winter driving conditions in Canada. Additionally, in [17] a dataset collected in northern Europe, capturing rare adverse weather situations, was published.

Recent research has focused on generating synthetic data [18] to simulate adverse weather conditions, as collecting and annotating real data in such conditions is time-consuming and costly. In [19] a physically valid fog simulation method was developed that is applicable to any lidar dataset. Another approach for simulating rain was presented in [20]. In [21] an efficient physics-based simulation method was introduced for augmenting lidar data collected under clear weather conditions, exploring fog, rain and snow. The work of synthesis winter scenes was further extended in [22], where snowfall and wet surfaces were modeled with precision and physical accuracy. However, despite the direct proven impact of point clouds, only limited work has been published on simulating adverse weather conditions on a Vehicle-to-Vehicle perception dataset.

III. Overview of the Proposed Framework

This section provides a detailed examination of the proposed network architecture used for training and evaluation, followed by a brief description of the fog simulation method chosen.

A. Cooperative Perception

For evaluation in adverse weather conditions, the developed network must be robust to varying levels of noise. While the use of cooperative perception alone can significantly contribute to maintaining high accuracy, a pipeline that emphasizes the extraction and fusion of the most important underlying data from each CAV's field of view is essential for robust detection. The pipeline consists of five primary stages, as illustrated in Fig. 1.

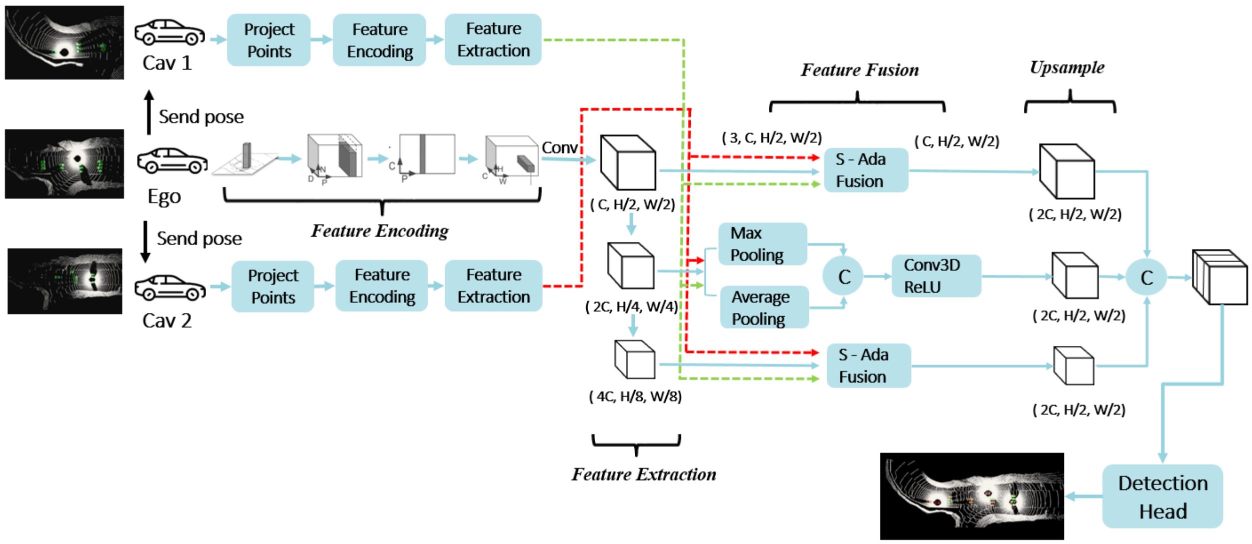

Fig. 1: The architecture of the proposed framework consists of five primary stages: 1) metadata sharing and point cloud projection, 2) feature encoding, 3) feature extraction, 4) feature map fusion, and 5) generation of predictions. A comprehensive representation of the actual dimensions of each involved tensor is depicted.

In the first phase, one of the vehicles is designated as the 'ego vehicle', which then transmits data regarding its transformation to its adjacent vehicles. Upon receiving this information, the remaining vehicles project their point clouds onto the coordinate system of the ego vehicle before proceeding with further data processing. The feature encoding and extraction follow the PointPillars [12] method. This choice is based on a comparative analysis published in [2], highlighting that PointPillars, along with VoxelNet, outperforms the other algorithms in terms of accuracy. However, PointPillars was selected for implementation because of its faster inference time, making it ideal for use in real-time applications. It also requires less computational resources and optimises memory usage.

A point cloud consists of a number of points, each defined by its (x, y, z) in three-dimensional space and intensity, i. According to the PointPillars method, the feature encoding process begins by dividing the space into a uniform grid within the x-y plane, resulting in the formation of a set of vertical columns, denoted as pillars, P. The points within each pillar are enhanced with parameters (x_c, y_c, z_c, x_p, y_p), representing the offset from the arithmetic mean of all points within the pillar and the offset from the pillar's center, respectively. Consequently, the point cloud can be described by the tensor (D, P, N), where D signifies the length of the vector describing each point, P indicates the number of pillars, and N represents the maximum number of points within each pillar. Subsequently, PointNet [9] is applied to each point within each pillar, followed by a max operation, resulting in a new tensor with dimensions (C, P). After feature encoding, the features are scattered back to their original pillar locations, constructing a pseudo-image of size (C, H, W).

For feature extraction, the backbone architecture consists of two sub-networks. The first subnetwork is responsible for generating features at progressively lower spatial resolutions and is composed of a sequence of blocks, each containing 2D convolutional layers, followed by batch normalization and ReLU activation functions. Conversely, the second subnetwork focuses on upsampling the features derived from the first subnetwork, which are then concatenated.

In the proposed pipeline, the fusion of feature maps from different vehicles is performed using the Spatial-wise adaptive feature fusion (S-AdaFusion) method [1] in between the two subnetworks. Let us consider the input tensor F, where F $\in R^{nxCxHxW}$ with n representing the number of different vehicles whose feature maps are to be fused. S-AdaFusion, employs max pooling and average pooling to extract spatial features from the input tensor F. This results in the generation of two new tensors, $S_{max} \in R^{1xCxHxW}$ and $S_{avg} \in R^{1xCxHxW}$, which are concatenated to form a new 4D tensor, $F_{spatial} \in R^{2xCxHxW}$. A 3D convolutional layer is then applied to this tensor, followed by a ReLU activation function, to facilitate further feature selection and dimensionality reduction. This process leads to an outpout tensor, $F_{fusion} \in R^{CxHxW}$. The implementation of the fusion method at this stage of the pipeline allows the extraction of finer details from the deeper layers of the network and, to some extent, reduces the loss of accuracy due to additional noise from sensors observations in challenging weather conditions.

The fused features are fed into the prediction header, which is responsible for predicting, for each anchor box, both a confidence score and the offsets (x, y, z, w, l, h, ϑ) pertaining to the ground truth bounding boxes.

B. Fog simulation model

To effectively simulate fog within point clouds, several critical factors must be considered to achieve a high level of realism. It is vital to acknowledge that fog can interact with laser beams by absorbing, scattering, and even backscattering them through fog particles. These interactions result in a reduction in the number of points within a point cloud compared to clear weather conditions, as well as the introduction of false points. Additionally, the strength of the returned laser pulse for each generated point, also known as intensity, is influenced by weather conditions. Specifically, fog increases humidity, making surfaces wet and reducing their reflectivity, resulting in decreased intensity.

The fog simulation approach employed is based on a probabilistic model for creating realistic fog, assuming a uniform fog distribution. This model was introduced in [23] and can be applied to any LiDAR dataset collected under clear weather conditions. The algorithm initiates by selecting the points in the initial point cloud that are modified by the presence of fog. The selection of a point that is modified is determined by a probability that depends on both the distance of the point from the sensor and the visibility. The greater the distance and the lower the visibility, the more likely it is that the point will be modified. The introduction of the visibility into the model's equations, enables the evaluation of how different fog intensities affect the accuracy of a prediction model. The points selected for modification are then either deleted or moved closer to the sensor. Examples of fog simulation are shown in Fig. 2.

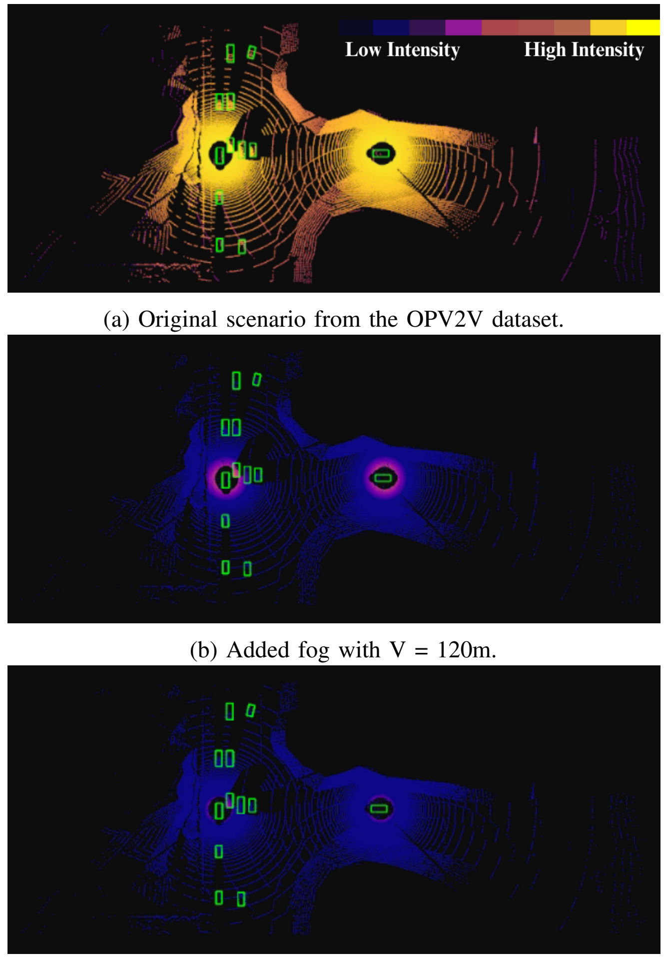

(a) Original scenario from the OPV2V dataset.

(b) Added fog with V = 120m.

(c) Added fog with V = 80m.

Fig. 2: A comparison between the same scenario in clear weather and under varying levels of fog. The colour change of the points indicates the reduced intensity due to the fog. The blue colour indicates low intensity.

IV. Experiments

A. Experiment Details

The experiments were conducted on the OPV2V dataset [2], which is the first large-scale open simulated dataset for Vehicle-to-Vehicle perception. This dataset was collected using a cooperative driving co-simulation framework called OpenCDA [24], in conjunction with the CARLA simulator [25]. The data were gathered from 8 different cities, resulting in a total of 73 scenarios, each with a duration of approximately 25 seconds. To enhance the diversity and realism of the dataset, the scenarios include 6 different types of road. Each scenario features on average 3 CAVs, with a minimum of 2 and a maximum of 7 CAVs. The dataset is divided into three separate sets for training, validation and testing purposes. Although the dataset includes data from both cameras and LiDAR sensors, this study focuses specifically on utilizing the LiDAR-generated point clouds.

The network architecture, as presented in Section II, is trained on the OPV2V dataset, which contains data collected under clear weather conditions. For training, the ranges of x, y, z are set to [(-140.8, 140.8), (-40, 40), (-3, 1)] meters, respectively, following the configuration in [2]. Given the significant amount of data within each scenario, a batch size of 2 is employed. The x and y resolutions of the pillar are set to 0.2 meters, as this value allows the model to capture features in detail without excessive memory usage or significant slowdown in the detection process. For the development of the model the Pytorch framework is utilized. Moreover, the Adam Optimizer [26] is employed with an ε of 10^{-10} and a weight decay of 0.0001. The model is trained for 20 epochs with a learning rate starting at 0.0005. Additionally, the Multisteplr scheduler is employed to adjust the learning rate when the training reaches 10 and 15 epochs.

Performance assessment is based on the Average Precision (AP) metric, using IoU threshold 0.7. The proposed model achieves an accuracy of **0.88** for AP@0.7. All training and evaluation procedures were performed on an NVIDIA GeForce RTX 3090 GPU.

B. Results

To evaluate the impact of fog on LiDAR-based object detection, the fog simulation model presented in Section II is applied to the testing data from the OPV2V dataset, considering varying levels of visibility. Specifically, experiments are conducted on two distinct datasets provided by OPV2V: Default Towns, collected from the CARLA simulator, and the Culver City Digital Town. The latter dataset features scenarios that more closely resemble challenging real-world driving environments, providing insights into how a CAV would react in such conditions.

A comparative analysis of the results is presented in Table I for both of these different testing sets. When visibility drops below 140 meters, a thick fog becomes noticeable, to the extent that even an experienced driver would find it difficult to drive in. Under these conditions,there is a slight reduction in the accuracy of the model, but overall the model performs well. For visibility values of 50 meters or less, severe disruption to transport occurs. Consequently, a significant reduction in the model's accuracy is observed. Only vehicles in close proximity to the ego vehicle are reliably detected, while further away, some false positive detections might occur. In the case of the Culver City scenarios, detection accuracy is lower due to the more realistic traffic conditions, making the detection task more challenging.

TABLE I: Performance of the proposed pipeline in foggy conditions for different visibility values.

Visibility [m]	No Fog	140	120	100	80	60	40
Default Towns	0.88	0.74	0.73	0.72	0.70	0.66	0.61
Culver City	0.81	0.63	0.61	0.59	0.55	0.49	0.41

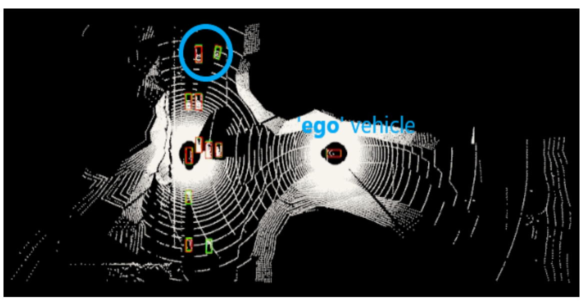

(a) Original point cloud from the OPV2V dataset.

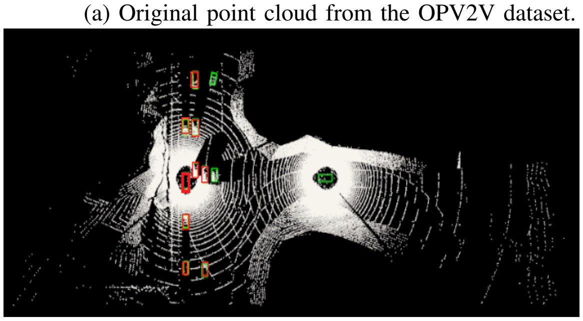

(b) Detected vehicles in fog with V = 60m.

Fig. 3: Examples of the proposed pipeline on the OPV2V dataset. The ground truth bounding boxes are shown in green, while the predicted bounding boxes are shown in red. The blue circle indicates vehicles that are occluded and would not be detected in single vehicle perception by the 'ego' vehicle.

Fig. 4 demonstrates the drop in accuracy under foggy conditions for various fusion methods and single vehicle perception (no fusion). It is observed that the accuracy of the proposed framework significantly surpasses that of other state-of-the-art methods, including early, late, intermediate and no fusion methods. This promising result can be attributed to several key factors. Firstly, the fusion method,

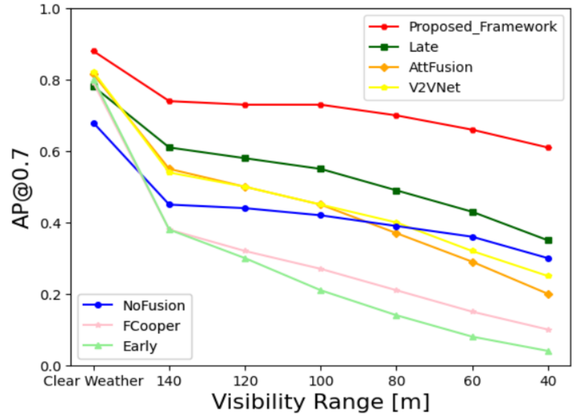

(a) Results on the Default Town.

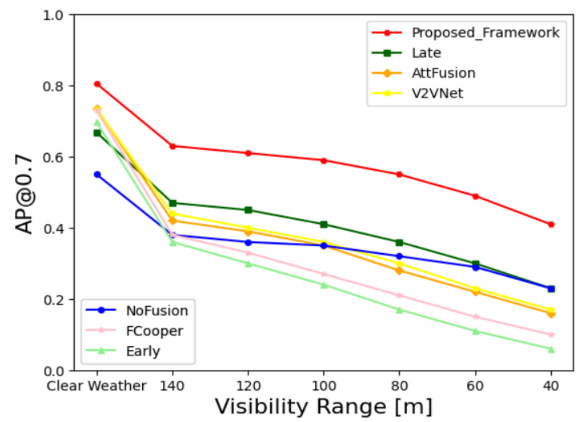

(b) Results on Culver City.

Fig. 4: Results of different methods of cooperative perception and single vehicle perception under fog simulation.

S-AdaFusion, makes use of CNNs with larger kernels which focus on more than one cells of the feature maps and extract spatial features from the neighbouring cells. Additionally, the fusion process takes place between the feature extraction and the upsampling of the feature maps, allowing the selection of more detailed features from the deeper layers of the network. This approach leads to an understanding of the data at different levels of abstraction, which is particularly important in the context of fog, where noise is introduced into the data. Furthermore, the choice of using smaller-sized pillars creates feature maps with larger dimensions, which contain more spatial information and can capture finer details, resulting in a more detailed representation of the underlying data.

While employing a model that achieves higher accuracy is desirable, it is essential to consider potential drawbacks. Therefore, the time complexity, critical for real-time applications, is examined. The time delay of each model per sample is shown in Table II. The delay of the proposed pipeline is comparable to that of other methods, considering the fact that accuracy is maintained regardless of the noise introduced by the fog. Ultimately, given that hardware limitations persist, it is necessary to strike a balance between time delay and accuracy to find the most appropriate solution.

TABLE II: Time delay per sample for SOTA methods of vehicle perception.

Method	Delay [sec]	Method	Delay [sec]
NoFusion [12]	0.09	Early Fusion [3]	0.084
Late Fusion [4]	0.138	F-Cooper [5]	0.09
AttFusion [2]	0.09	V2VNet [6]	0.144
Proposed Framework	0.157		

V. Conclusions

This study addresses some of the limitations of AVs, such as object occlusion, by incorporating vehicle-to-vehicle perception to increase the visibility range of each vehicle. Taking this a step further, CAVs are tested in challenging weather conditions by applying a fog simulation model to the first large-scale open simulated dataset for Vehicle-to-Vehicle perception. On the modified dataset a network trained in clear weather conditions is evaluated. The proposed pipeline combines PointPillars, a leading object detection method for autonomous vehicles, with Spatial-wise Adaptive feature fusion method, which effectively extracts the most significant features using neural networks. The results indicate highly promising outcomes and no significant degradation is observed even in the presence of strong fog.

In conclusion, the results of testing connected autonomous vehicles in diverse environments highlights the substantial contribution of cooperative perception and enhances the trust in autonomous vehicles.

Acknowledgement

This work was supported by the European Union's Horizon Europe Research and Innovation Program under Grant 101092875—DIDYMOS-XR: Digital DynaMic and respOnsible twinS for XR.

References

[1] D. Qiao and F. Zulkernine, "Adaptive feature fusion for cooperative perception using lidar point clouds," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2023, pp. 1186–1195.

[2] R. Xu, H. Xiang, X. Xia, X. Han, J. Li, and J. Ma, "Opv2v: An open benchmark dataset and fusion pipeline for perception with vehicle-to-vehicle communication," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 2583–2589.

[3] Q. Chen, S. Tang, Q. Yang, and S. Fu, "Cooper: Cooperative perception for connected autonomous vehicles based on 3d point clouds," in *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 2019, pp. 514–524.

[4] A. Rauch, F. Klanner, R. Rasshofer, and K. Dietmayer, "Car2x-based perception in a high-level fusion architecture for cooperative perception systems," in *2012 IEEE Intelligent Vehicles Symposium*. IEEE, 2012, pp. 270–275.

[5] Q. Chen, X. Ma, S. Tang, J. Guo, Q. Yang, and S. Fu, "F-cooper: Feature based cooperative perception for autonomous vehicle edge computing system using 3d point clouds," in *Proceedings of the 4th ACM/IEEE Symposium on Edge Computing*, 2019, pp. 88–100.

[6] T.-H. Wang, S. Manivasagam, M. Liang, B. Yang, W. Zeng, and R. Urtasun, "V2vnet: Vehicle-to-vehicle communication for joint perception and prediction," in *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*. Springer, 2020, pp. 605–621.

[7] A. Gkillas, G. Arvanitis, A. S. Lalos, and K. Moustakas, "Federated learning for lidar super resolution on automotive scenes," in *2023 24th International Conference on Digital Signal Processing (DSP)*. IEEE, 2023, pp. 1–5.

[8] B. Yang, W. Luo, and R. Urtasun, "Pixor: Real-time 3d object detection from point clouds," in *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2018, pp. 7652–7660.

[9] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 652–660.

[10] Y. Zhou and O. Tuzel, "Voxelnet: End-to-end learning for point cloud based 3d object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 4490–4499.

[11] Y. Yan, Y. Mao, and B. Li, "Second: Sparsely embedded convolutional detection," *Sensors*, vol. 18, no. 10, p. 3337, 2018.

[12] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, "Pointpillars: Fast encoders for object detection from point clouds," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 12 697–12 705.

[13] T. Peynot, J. Underwood, and S. Scheding, "Towards reliable perception for unmanned ground vehicles in challenging conditions," in *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2009, pp. 1170–1176.

[14] R. Heinzler, P. Schindler, J. Seekircher, W. Ritter, and W. Stork, "Weather influence and classification with automotive lidar sensors," in *2019 IEEE intelligent vehicles symposium (IV)*. IEEE, 2019, pp. 1527–1534.

[15] M. Bijelic, T. Gruber, and W. Ritter, "A benchmark for lidar sensors in fog: Is detection breaking down?" in *2018 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2018, pp. 760–767.

[16] M. Pitropov, D. E. Garcia, J. Rebello, M. Smart, C. Wang, K. Czarnecki, and S. Waslander, "Canadian adverse driving conditions dataset," *The International Journal of Robotics Research*, vol. 40, no. 4-5, pp. 681–690, 2021.

[17] M. Bijelic, T. Gruber, F. Mannan, F. Kraus, W. Ritter, K. Dietmayer, and F. Heide, "Seeing through fog without seeing fog: Deep multimodal sensor fusion in unseen adverse weather," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11 682–11 692.

[18] A. Kloukiniotis, A. Papandreou, C. Anagnostopoulos, A. Lalos, P. Kapsalas, D.-V. Nguyen, and K. Moustakas, "Carlascenes: A synthetic dataset for odometry in autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 4520–4528.

[19] M. Hahner, C. Sakaridis, D. Dai, and L. Van Gool, "Fog simulation on real lidar point clouds for 3d object detection in adverse weather," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 15 283–15 292.

[20] C. Goodin, D. Carruth, M. Doude, and C. Hudson, "Predicting the influence of rain on lidar in adas," *Electronics*, vol. 8, no. 1, p. 89, 2019.

[21] V. Kilic, D. Hegde, V. Sindagi, A. B. Cooper, M. A. Foster, and V. M. Patel, "Lidar light scattering augmentation (lisa): Physics-based simulation of adverse weather conditions for 3d object detection," *arXiv preprint arXiv:2107.07004*, 2021.

[22] M. Hahner, C. Sakaridis, M. Bijelic, F. Heide, F. Yu, D. Dai, and L. Van Gool, "Lidar snowfall simulation for robust 3d object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 16 364–16 374.

[23] S. Teufel, G. Volk, A. Von Bernuth, and O. Bringmann, "Simulating realistic rain, snow, and fog variations for comprehensive performance characterization of lidar perception," in *2022 IEEE 95th Vehicular Technology Conference:(VTC2022-Spring)*. IEEE, 2022, pp. 1–7.

[24] R. Xu, Y. Guo, X. Han, X. Xia, H. Xiang, and J. Ma, "Opencda: an open cooperative driving automation framework integrated with co-simulation," in *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2021, pp. 1155–1162.

[25] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," in *Conference on robot learning*. PMLR, 2017, pp. 1–16.

[26] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.