International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

Radar Powered Perception: Multiple Approaches for Vehicle Classification

Deepak Misra[1], Brijesh Jha[2], Shyam Raman[3]

[1]Assistant Manager, Research & Advanced Engineering, Renault Nissan Technology and Business Center India Pvt Ltd
[2]Deputy General Manager, Research & Advanced Engineering, Renault Nissan Technology and Business Center India Pvt Ltd
[3]Senior Engineer, Research & Advanced Engineering, Renault Nissan Technology and Business Center India Pvt Ltd

Abstract
The proliferation of Advanced Driver-Assistance Systems (ADAS) and autonomous vehicles demands robust and reliable techniques for vehicle detection and classification. This paper investigates the potential of automotive radar systems for this critical task, focusing on classifying vehicles into three distinct categories (bus/truck, car/sedan/SUV, bike/bicycle). We explore multiple approaches that leverages rich set of features extracted from radar measurements. These features include target size, point density, and other signature characteristics derived from the radar return signal. By analyzing these features, we explain multiple approach aims to achieve vehicle classification. The effectiveness of the proposed method is evaluated using real-world data, with a particular focus on demonstrating its suitability for real-time applications in ADAS. The paper delves into the critical trade-offs that exist between classification accuracy, computational complexity, and inherent limitations of automotive radar sensors. We discuss the impact of sparse data, such as low point cloud data on classification accuracy. Additionally, we discuss on potential mitigation strategies to ensure reliability in challenging scenarios. Finally, the paper concludes by outlining potential future directions for enhancing vehicle classification using automotive radar systems. These advancements will pave the way for the development of more robust and reliable ADAS and autonomous vehicles, ultimately contributing to improved safety on our roads.

Keywords: Automotive Radar, Vehicle Detection, Vehicle Classification, ADAS, Autonomous Vehicles, Point Cloud Data, Machine Learning

1. Introduction

The rapid advancement of automotive technologies has led to the widespread adoption of Advanced Driver-Assistance Systems (ADAS) and a growing interest in autonomous vehicles. These technologies aim to enhance vehicle safety, improve driving comfort, and reduce the likelihood of accidents. Central to the effectiveness of ADAS and autonomous driving systems is the ability to accurately detect and classify vehicles in various driving conditions. Traditionally, vehicle detection and classification have relied heavily on optical sensors like cameras and LiDAR. However, these sensors have limitations, particularly in adverse weather conditions and scenarios with poor visibility. Automotive radar systems offer a promising alternative due to their robustness in diverse environmental conditions. Unlike optical sensors, radar can operate effectively in fog, rain, and darkness, making it a reliable tool for continuous vehicle detection and classification. Radars emit radio waves and analyze the reflected signals to determine the range, velocity, and characteristics of objects. This capability enables radar to provide essential information for distinguishing different types of vehicles. This paper explores the potential of automotive radar systems for vehicle classification, focusing on the categorization of vehicles into three distinct classes: bus, car, and bike. Our approach leverages a rich set of features extracted from radar measurements, such as target size, relative velocity, and other signature characteristics derived from the radar return signal. By analyzing these features, we aim to achieve accurate vehicle classification suitable for real-time applications in ADAS. Our study begins with a comprehensive literature review, highlighting the strengths and limitations of existing vehicle classification methods and the role of radar technology in automotive applications. We then present our research design, detailing the system overview, major components, and various classification approaches we investigated. These approaches include clustering algorithms, analysis of object dimensions, point count, density, accumulation of data over continuous frames, and machine learning techniques. We conducted extensive experiments using real-world data collected from a vehicle equipped with radar sensors. The data encompasses various types of vehicles, including buses, trucks, cars, sedans, SUVs, auto-rickshaws, bikes, and bicycles. We annotated this data into three categories for analysis: bus, car, and bike. The effectiveness of our proposed method is evaluated based on its classification accuracy and computational efficiency, with a particular focus on its suitability for real-time deployment. The results and discussion section provides a detailed analysis of the performance of each classification approach. We delve into the critical trade-offs between classification accuracy, computational complexity, and the inherent limitations of automotive radar sensors. We also examine the impact of different environmental conditions, such as adverse weather and occlusions caused by other vehicles, on the system's performance. Potential mitigation strategies are discussed to ensure robustness in challenging scenarios. In conclusion, this paper outlines the potential future directions for enhancing vehicle classification using automotive radar systems. These advancements will pave the way for developing more robust and reliable ADAS and autonomous vehicles, ultimately contributing to improved safety on our roads. Through this study, we aim to provide valuable insights into the application of radar technology for vehicle classification and highlight its critical role in the evolution of automotive safety systems.

2. Literature Review

The increasing interest in leveraging radar for object detection and classification in autonomous vehicles is reflected in many literatures. Several key themes emerge:

Radar Technology and Applications

Schneider and Wenger (2003) and Gao et al. (2022) delve into the utilization of radar for object detection and classification in automotive applications [10, 14]. Schneider and Wenger (2003) highlight the potential of high-resolution radar sensors for enhancing object detection and classification in traffic environments [14]. They emphasize its ability to provide detailed information about the environment, including the detection of hidden objects and road boundaries, and discuss the system design, hardware, and software implementation of a prototype high-resolution automotive radar system [14]. Gao et al. (2022) present their experiments with a millimeter-wave radar test-bed based on Texas Instrument's automotive chipset family [10]. They describe the test-bed components, summarize FMCW radar operational principles, and present preliminary results on object recognition using radar imaging algorithms [10].

Srivastav and Mandal (2023) provide a comprehensive review of deep learning methods and challenges in utilizing radar data for autonomous driving [1]. They discuss various data formats and representations for radar data, including range-Doppler tensors, range-azimuth heatmaps, point clouds, and micro-Doppler spectrograms, and address challenges such as low resolution, sparsity, clutter, high uncertainty, and lack of good datasets [1]. Schumann et al. (2021) introduce the RadarScenes dataset, a large-scale dataset with point-wise annotations of radar data collected from four series production automotive radar sensors mounted on a test vehicle [13]. The dataset is designed to enable the development of novel machine learning-based radar perception algorithms with a focus on moving road users [13]. Kraus et al. (2021) present the Radar Ghost Dataset, which includes detailed manual annotations for different kinds of ghost objects (multi-path reflections) in automotive radar data [12]. The dataset aims to facilitate research on multi-path object suppression and exploitation, and the authors evaluate two different approaches for identifying these ghost objects [12].

Schneider and Wenger (2003) discuss potential applications of high-resolution radar in Advanced Driver Assistance Systems (ADAS), such as adaptive cruise control, obstacle warning, emergency braking, collision avoidance, and pedestrian detection [14]. They also explore the possibility of using radar for parking space identification and curve warning [14].

Machine Learning and Clustering Algorithms

Ulrich et al. (2021) introduce DEEPREFLECS, a deep learning method for classifying objects (e.g., pedestrian, cyclist, car) using radar reflections [2]. Their approach utilizes a lightweight neural network that can handle varying numbers of radar reflections and combines the extraction of local and global features [2]. The method is shown to outperform existing methods of handcrafted or learned features in experiments with real data [2].

Bindra and Mishra (2017) present a detailed study of clustering algorithms, a fundamental task in data mining [11]. They review various clustering techniques, including partitional (e.g., K-means, K-medoids, CLARANS, ISODATA), hierarchical (e.g., BIRCH, CURE, ROCK, CHAMELEON), and density-based (e.g., DBSCAN, DENCLUE) clustering, and discuss their efficiency, advantages, and disadvantages in terms of scalability, handling of arbitrary shaped clusters, similarity/dissimilarity measures, and sensitivity to noise and outliers [11]. Han (2021) compares the performance of decision tree, support vector machine (SVM), and naive Bayes classifiers in a classification problem using the iris dataset [15]. The study finds that the decision tree model outperforms the other two models in terms of accuracy for this specific dataset [15]. The paper also discusses the theoretical principles and implementation details of each classifier, highlighting the differences between generative and discriminative models and their performance on small datasets [15].

Data Annotation and Labeling

Agrawal et al. (2023) and Wang et al. (2021) address the challenge of annotating sensor data for training deep learning models in the context of autonomous driving [3, 4]. Agrawal et al. (2023) propose a semi-automatic annotation methodology for RGB camera images and 3D radar point cloud data using a smart infrastructure-based sensor setup [3]. Their approach aims to reduce the manual effort required for data annotation by leveraging information from both sensors and introducing a new object category, "GROUP," for radar-based object detection of closely located vulnerable road users [3]. They also present a new method for 3D radar background subtraction to remove clutter [3]. Wang et al. (2021) focus on radar object detection (ROD) and propose a systematic annotation system that aligns camera and radar data to generate accurate 3D object labels in radar radio frequency (RF) images [4]. They introduce the CRUW dataset, a large-scale dataset with synchronized camera-radar frames collected in various driving scenarios, and propose evaluation metrics for ROD [4].

Sensor Fusion and Calibration

Wang et al. (2021) highlight the importance of camera-radar fusion for improving object detection in challenging driving scenarios [4]. Their annotation system leverages the strengths of both sensors, utilizing the rich semantic information from camera images and the accurate range and velocity information from radar data, to generate precise 3D object labels in radar RF images [4]. Agrawal et al. (2023) emphasize the importance of sensor calibration and time synchronization in multi-sensor setups for accurate data annotation and fusion [3]. They describe their smart infrastructure-based measurement setup, which includes an RGB mono camera, 3D automotive radar, and 360° automotive lidar sensor, and detail the calibration process to ensure accurate spatial alignment and temporal synchronization of the sensor data [3].

Other Automotive Technologies

Visconti et al. (2019) describe the design and testing of an electronic control system for a Formula SAE race car [6]. The system utilizes an STM32 Nucleo board to acquire data from various sensors installed on the vehicle, such as temperature sensors, a thermistor for cooling liquid temperature, Hall effect speed sensors, and potentiometers for suspension extension detection [6]. The acquired data is then transmitted wirelessly to a base station for real-time monitoring, enabling the technical staff to monitor the car's performance and ensure driver safety [6].

Margapuri and Neilsen (2019) present a prototype implementation of a temperature control system using Controller Area Network (CAN) communication and FreeRTOS on STM32F407 Discovery Boards [7]. The system utilizes low-cost components and open-source software, making it an ideal target for classroom use and demonstrating the practical implementation of a real-time control system [7]. The paper also discusses the concepts of CAN messaging and its parameters, providing a valuable resource for those interested in setting up their own CAN network [7].

Aiswarya and Prabhakar (2015) propose an efficient high-gain DC-DC converter for automotive applications, such as supplying high-intensity discharge (HID) lamps [8]. The converter is designed to provide a high output voltage from a 12V input supply and utilizes a clamp circuit to achieve soft switching, reducing switching stress and increasing efficiency [8].

Woo et al. (2019) introduce an angle sensor module for vehicle steering based on a multi-track impulse ring [9]. The module aims to improve angle detection accuracy and reduce errors compared to conventional torque angle sensor (TAS) modules [9]. The paper details the fabrication, testing, and evaluation of the module, demonstrating its excellent performance with an average deviation of 0.4 degrees and its applicability to actual vehicles by evaluating its electromagnetic interference (EMI) environmental reliability [9].

Ghael et al. (2020) provide a review of the Raspberry Pi, a versatile and affordable single-board computer [5]. They discuss its technical specifications, boot process, advantages, limitations, and applications in various projects, highlighting its significance in the field of embedded systems [5].

3. Research Design

This research aims to develop a robust radar-based vehicle classification system for real-time ADAS applications. The study encompasses data collection from real-world driving scenarios, annotation of point cloud data into bus, car, and bike categories, and the exploration of various classification approaches.

These approaches include clustering algorithms, analysis of object dimensions and densities, data accumulation over frames, formulation of relevant features, and machine learning techniques. Rigorous evaluation on real-world data focuses on classification accuracy, computational efficiency, and the impact of environmental factors. The research aims to identify the most effective classification methods and provide insights into the potential of automotive radar for accurate and real-time vehicle classification.

Other Automotive Technologies

The vehicle classification system leverages a multi-sensor architecture for robust and comprehensive data acquisition. Figure 1 shows the system architecture where the core of the system has three radar sensor boards, each responsible for capturing raw radar data from the surrounding environment. Dedicated DC-DC converters provide stable power to these sensor boards, ensuring their optimal performance. The system additionally incorporates an STM Nucleo module, which interfaces with the steering angle encoder. This encoder provides valuable contextual information about the vehicle's orientation and maneuvers, which can be crucial for accurate vehicle classification, particularly in dynamic traffic scenarios.

Figure 1: System Architecture

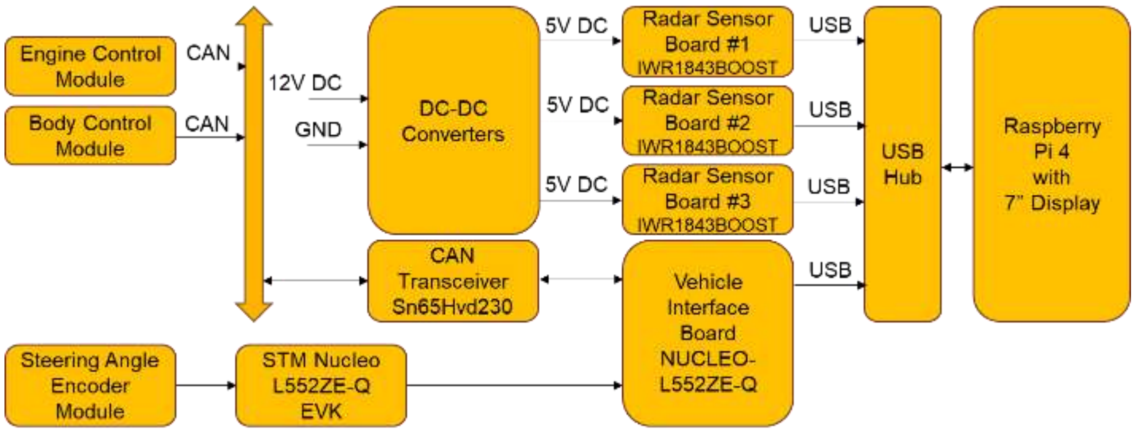

The data acquisition process is further enhanced by the vehicle interface board. This central hub consolidates data streams from the radar sensors and the steering angle encoder. It also facilitates communication with the vehicle's Controller Area Network (CAN) bus, enabling the system to interact with other onboard sensors and control units. Additionally, the vehicle interface board transmits the aggregated data to the Raspberry Pi 4 through a USB hub. The Raspberry Pi 4 serves as the brain of the system, responsible for real-time data processing, analysis, and classification. Equipped with a 7" display, it allows for system monitoring and visualization of classification results. The computational power of the Raspberry Pi 4 makes it suitable for implementing complex machine learning algorithms, enabling the system to achieve high levels of classification accuracy. This modular architecture offers several advantages. The use of multiple radar sensor boards provides a wider field of view and potentially improved detection accuracy, particularly for larger vehicles or those at oblique angles. The system's modularity also allows for future expansion or modifications, such as the integration of additional sensors or processing units, to enhance its capabilities.

Major Components

The vehicle classification system leverages the Texas Instruments (TI) IWR1843BOOST radar sensor board as its primary data acquisition unit. This board is a highly integrated single-chip millimeter-wave (mmWave) sensor front end, combining a 76-81 GHz frequency modulated continuous wave (FMCW) transceiver with a digital signal processor (DSP). The sensor's capabilities include range, velocity, and angle detection, making it ideal for precise object detection and classification in automotive scenarios. Three of these radar sensor boards are strategically positioned around the vehicle to provide a comprehensive 360-degree view of the surrounding environment. Each radar sensor board is powered by a dedicated DC-DC converter, transforming the vehicle's 12V DC power supply into the required 5V DC. This ensures stable operation of the radar sensors and other components in the system. To integrate the radar data with the vehicle's dynamics, an STM Nucleo L552ZE-Q microcontroller board is used to interface with the steering angle encoder. This module reads the steering angle information, providing crucial context for interpreting the radar data. The steering angle information is then transmitted to the vehicle interface board, a custom-designed board that consolidates data from all three radar sensor boards and the steering angle encoder. It also houses a CAN transceiver (Sn65HVD230) that enables seamless communication with the vehicle's CAN bus, allowing for the exchange of information with other onboard systems. A Raspberry Pi 4, equipped with a 7-inch display, serves as the central processing unit of the system. The aggregated data from the vehicle interface board is transferred to the Raspberry Pi via USB. The Raspberry Pi's powerful processing capabilities are leveraged to perform higher-level data processing, analysis, and the application of machine learning algorithms for vehicle classification. The display serves as a visual interface for real-time monitoring of the system's operation and the visualization of the classification results. A USB hub is employed to expand the Raspberry Pi's connectivity, accommodating the data streams from all three radar sensor boards simultaneously. This comprehensive system architecture integrates radar technology, microcontroller capabilities, and high-performance computing to create a robust and efficient platform for vehicle classification. The integration of the steering angle encoder allows for a more nuanced understanding of the radar data, potentially leading to improved classification accuracy. The modular design of the system ensures flexibility and scalability for future enhancements and modifications.

Classification Approaches

To achieve robust and accurate vehicle classification using radar data, this research investigates multiple approaches, each leveraging different aspects of the radar point cloud data and employing distinct methodologies.

Clustering Algorithms

Clustering algorithms, specifically DBScan and Agglomerative clustering, are applied to the radar point cloud data. These algorithms aim to group data points based on their spatial proximity and density, with the expectation that points belonging to the same vehicle will form distinct clusters. By analyzing the characteristics of these clusters, such as their size and shape, we aim to differentiate between different vehicle types.

Analysis of Object Dimension, Point Count, and Density

This approach focuses on analyzing the geometric properties of detected objects in the radar point cloud. The length, width, point count, and density of each object are extracted and analyzed. Statistical measures like mean, max, min, and percentiles are calculated to understand the distribution of these features across different vehicle types. Thresholds are then established based on these statistical measures to classify objects into the predefined categories (bus, car, bike).

Accumulation of Data Over Continuous Frames

Recognizing that a single frame of radar data may not provide sufficient information for accurate classification, this approach involves accumulating data over multiple consecutive frames. By integrating information from multiple frames, the point cloud representation of each object becomes denser and more informative. This accumulated data is then used for feature extraction and classification, with the expectation that the increased information content will lead to improved classification accuracy.

Formulation of Features

In this approach, a comprehensive set of features is derived from the radar point cloud data. These features go beyond basic object dimensions and include characteristics like diagonal length, maximum distances between points, perimeter, area, and standard deviations of point coordinates. The rationale behind this approach is to capture a wide range of geometric and statistical properties that can effectively differentiate between different vehicle types.

Machine Learning

Supervised machine learning algorithms are employed to learn patterns and relationships in the formulated features and classify vehicles accordingly. Various algorithms, including Random Forest, Support Vector Machines (SVM), Bayesian classifiers, linear classifiers, LightGBM, Gradient Boosting, Voting Classifier, and AdaBoost, are trained and evaluated. Hyperparameter tuning and cross-validation techniques are applied to optimize the performance of these models. Additionally, techniques like Synthetic Minority Over-sampling Technique (SMOTE) are used to address class imbalance in the dataset, ensuring that the models are not biased towards the majority class.

4. Results and Discussion

The evaluation of the different classification approaches reveals varying degrees of success and highlights the challenges and trade-offs inherent in radar-based vehicle classification.

Classification Approaches

The initial exploration of clustering algorithms, specifically DBSCAN and Agglomerative clustering, for vehicle classification using radar point cloud data did not yield satisfactory results. DBSCAN, a density-based algorithm, struggled to form distinct clusters due to the varying density of points within the radar data, particularly for vehicles located at greater distances from the sensor. This variability in point density hindered the algorithm's ability to accurately group data points belonging to the same vehicle. Table 1 shows the result of applying various clustering techniques. Agglomerative clustering, while able to form clusters, presented challenges in accurately estimating the dimensions of vehicles and consistently including all relevant data points. The algorithm's performance varied depending on the distance metric used. With Euclidean distance, the mean length and width for both car and bus categories were too similar, with overlapping 75th percentiles, making it difficult to distinguish between them. Specifically, buses had a mean width of 1.6m and a 75th percentile width of 1.8m, while cars had a mean width of 1.2m and a 75th percentile width of 1.5m. Similarly, the mean lengths were 2.5m for buses and 2.1m for cars, with 75th percentiles of 3.6m and 3.2m respectively. Using Manhattan distance did not significantly improve the results, as the mean lengths for both categories were identical at 3.9m, with mean widths of 2.2m for buses and 1.4m for cars. Furthermore, neither distance metric was able to form clusters for bikes, indicating a limitation in handling smaller objects with lower point densities. These findings suggest that clustering algorithms alone may not be sufficient for accurate vehicle classification using radar data, as they struggle to handle the inherent variability and noise in real-world radar point clouds. The overlapping cluster characteristics between different vehicle types, especially in terms of length and width, necessitate the exploration of additional features and classification techniques to achieve more accurate and reliable vehicle classification.

Table 1: Clustering Results

Clustering	Parameters	Category	Results
DBSCAN	NA	NA	Proper Clusters not formed
Agglomerative	Euclidean	Bus	Mean Width: 1.6 | 75th Percentile: 1.8 Mean Length: 2.5 | 75th Percentile: 3.6
Agglomerative	Euclidean	Car	Mean Width: 1.2 | 75th Percentile: 1.5 Mean Length: 2.1 | 75th Percentile: 3.2
Agglomerative	Euclidean	Bike	Clusters not formed
Agglomerative	Manhattan	Bus	Mean Width: 2.2 | 75th Percentile: 3.1 Mean Length: 3.9 | 75th Percentile: 5.1
Agglomerative	Manhattan	Car	Mean Width: 1.4 | 75th Percentile: 2.2 Mean Length: 3.9 | 75th Percentile: 4.2
Agglomerative	Manhattan	Bike	Clusters not formed

Analysis of Vehicle Dimension, Point Count & Density

In our quest to discern distinctive characteristics among vehicle types using radar data, we embarked on a meticulous analysis of object dimensions, point count, and density. Table 2 shows the impact on Point Count and Density based on vehicle's position from radar and it's distance. The analysis commenced with the collection of radar data for individual vehicles (bus, car, and bike) in isolated scenarios to minimize interference. To understand the impact of distance on object representation, we segmented the data into two ranges: less than 30 meters and greater than 30 meters from the radar sensor. Also, the detected position of the vehicle by the radar is segmented as Straight and Slant.

Table 2: Impact of vehicle position (Slant/Straight) and it's distance on Point Count and Density

Vehicle	Position behind Radar	Distance	No. of records	Point Count	Density
Bus	Straight	Less than 30 meters	1560	3-6	0.21
Bus	Straight	Greater than 30 meters	940	3-4	0.33
Bus	Slant	Less than 30 meters	1123	2-6	0.17
Bus	Slant	Greater than 30 meters	912	2-5	0.37
Car	Straight	Less than 30 meters	2344	1-5	0.39
Car	Straight	Greater than 30 meters	1360	0-2	0.35
Car	Slant	Less than 30 meters	740	1-4	0.24
Car	Slant	Greater than 30 meters	1007	0-2	0.33
Bike	Straight	Less than 30 meters	403	0-3	0.12
Bike	Straight	Greater than 30 meters	943	0-2	0.14
Bike	Slant	Less than 30 meters	252	0-2	0.07
Bike	Slant	Greater than 30 meters	278	0-1	0.08

For each vehicle category and distance range, we computed the mean, maximum, minimum, and percentile values for both length (ysize) and width (x-size). This statistical analysis unveiled that y-size exhibited a clear differentiation between vehicle categories, aligning with our expectations that larger vehicles like buses would generally have greater lengths compared to cars or bikes. However, the x-size demonstrated high variability within each category and substantial overlap between categories, rendering it less effective as an independent discriminator. Figure 2 shows how ysize varies with change in position of vehicle and increase/decrease of distance in a scatter plot.

Figure 2: Comparison of Length (ysize) for the vehicles (Bus, Car & Bike) over different range of distance and positions (Straight/Slant)

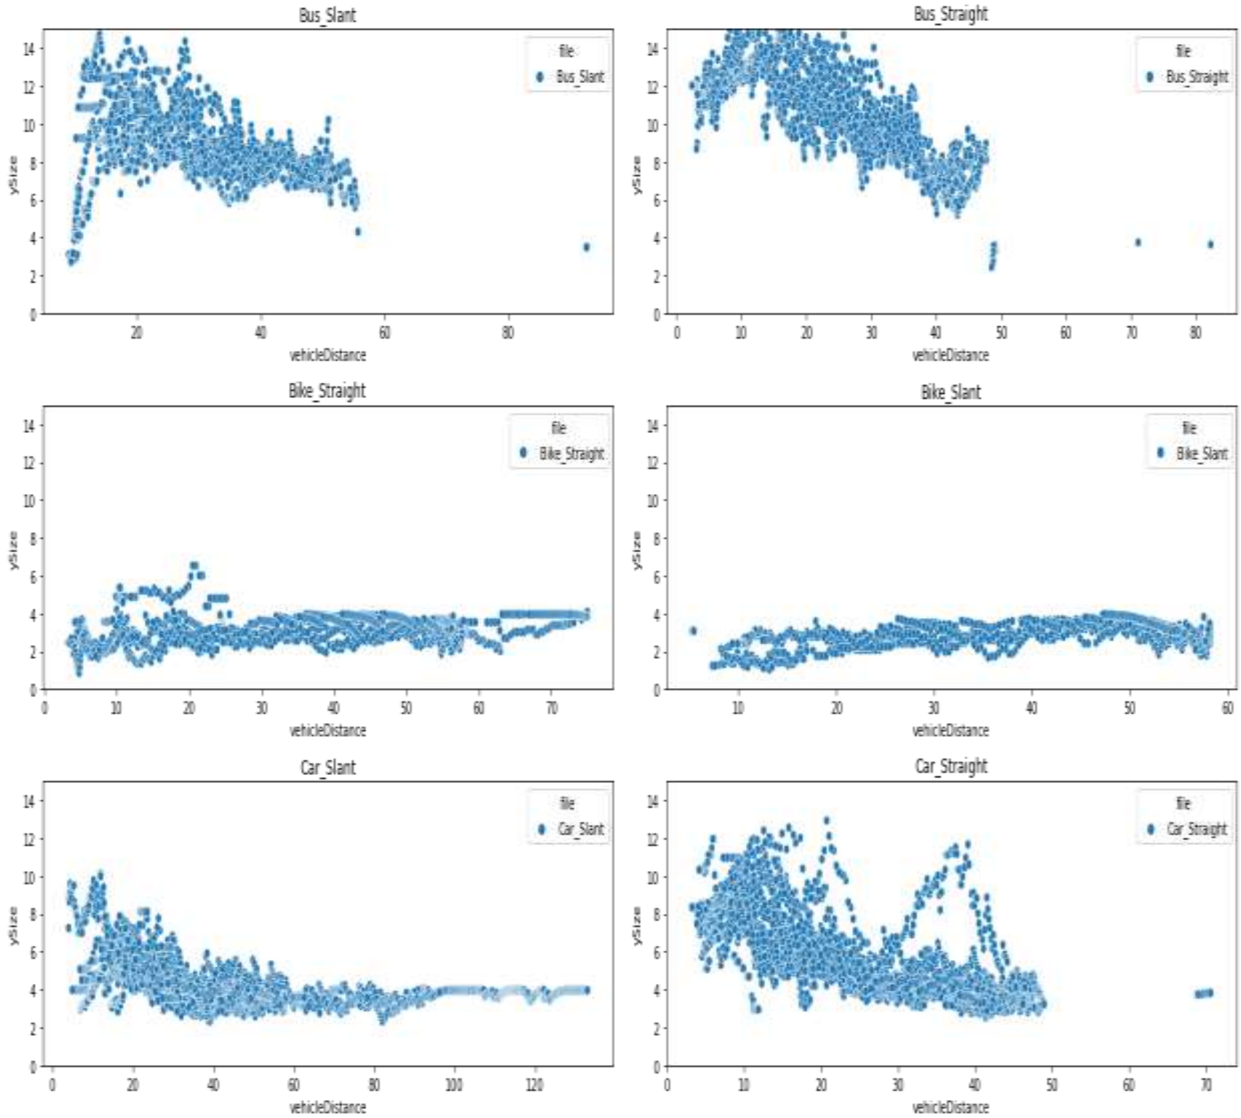

To further explore the potential of radar point cloud characteristics for classification, we developed an algorithm to calculate the point count (number of points within a detected object cluster) and density (point count divided by the object's area). We hypothesized that larger vehicles would generally exhibit higher point counts and potentially different density patterns compared to smaller vehicles. Our analysis of the calculated point count and density revealed several key insights. Firstly, the point count for all vehicle categories was relatively low, particularly for bikes, which often registered zero points. The results are shown in Figure 3.

Figure 3: Comparison of Point Count and Density for the vehicles (Bus, Car & Bike) over different range of distance and positions (Straight/Slant)

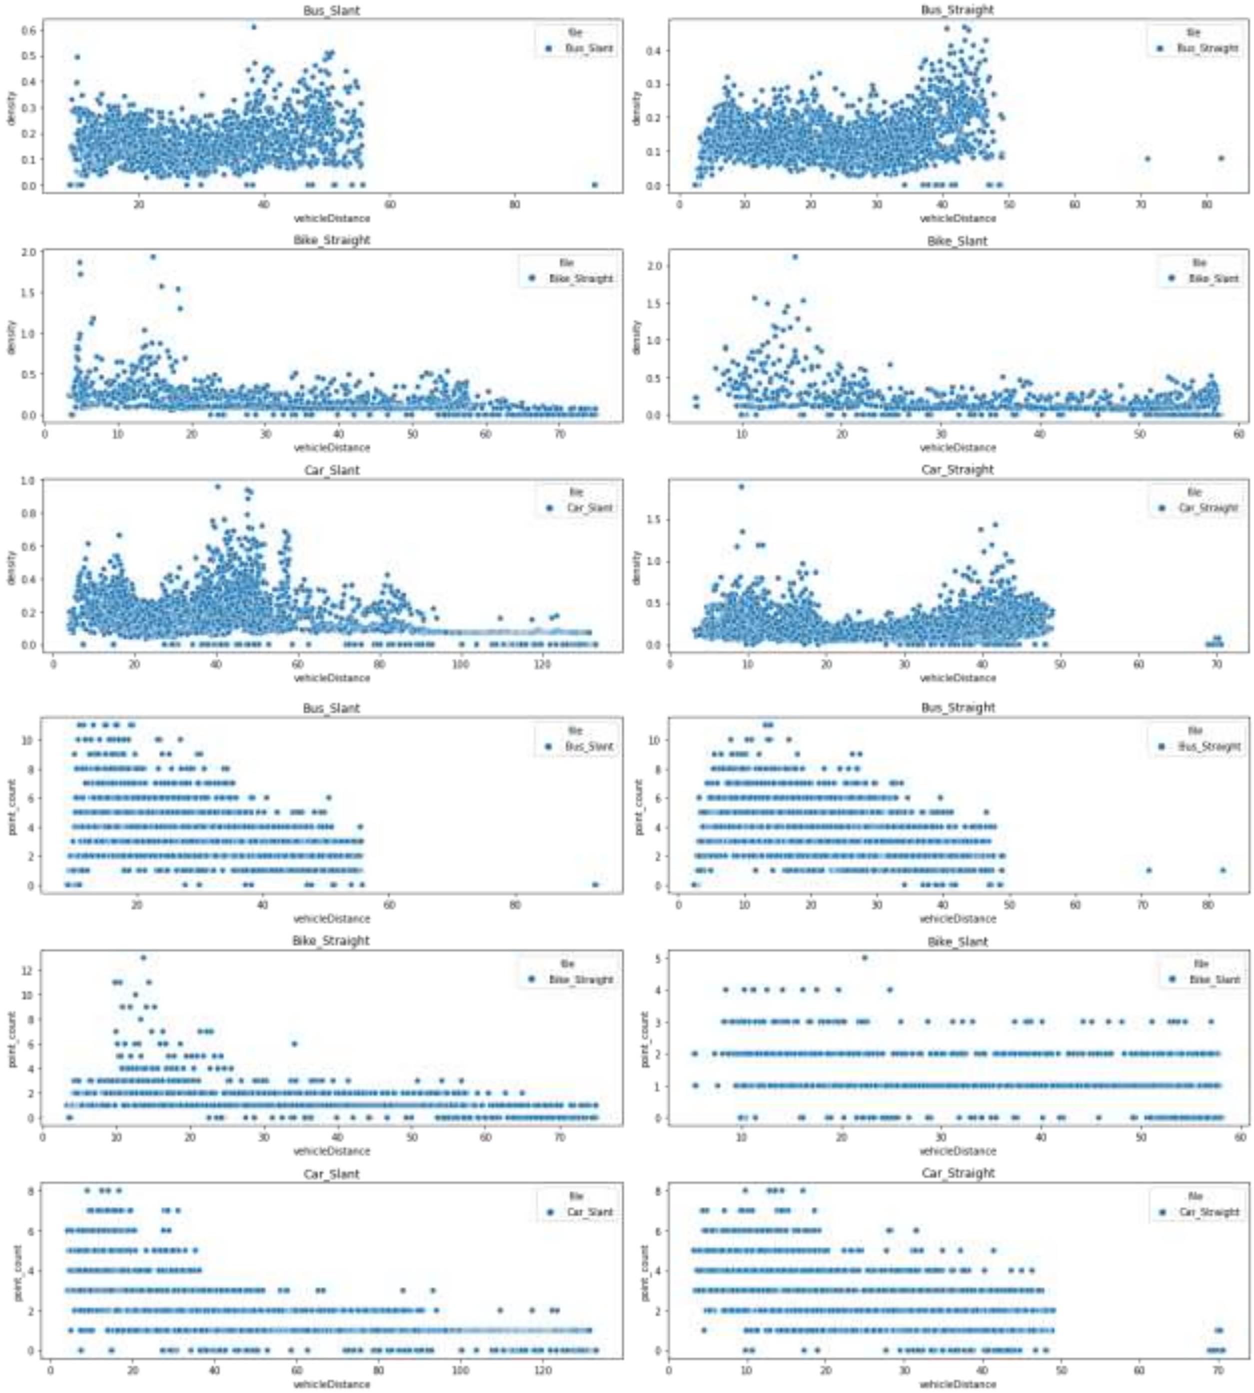

This observation underscores the challenge of capturing sufficient radar returns from smaller objects, especially at greater distances. Secondly, the density values showed considerable overlap between categories, making it difficult to establish distinct thresholds for classification. These findings suggest that while point count and density offer some discriminatory potential, they are not robust enough as standalone features for accurate vehicle classification. To assess the effectiveness of these features in a classification task, we calculated single-frame accuracy for each category using y-size, point count, and density individually and in combination shown in Figure 4. As expected, y-size alone yielded the highest accuracy, reinforcing its discriminatory power. While point count and density performed poorly individually due to the aforementioned limitations, combining all three features resulted in a notable improvement in accuracy compared to using point count or density alone. This indicates that, despite their individual weaknesses, these features can complement each other to enhance classification performance when used in conjunction. The results presented in Table 1 provide a comprehensive summary of the statistical analysis and classification accuracies obtained in this section. The table highlights the differentiability of y-size across vehicle categories and distances, while also illustrating the challenges posed by overlapping point count and density values. Overall, this analysis underscores the need for a more sophisticated approach that incorporates multiple features and potentially leverages machine learning techniques to achieve robust and accurate vehicle classification using radar data.

Figure 4: Accuracy comparison for combined, ysize, point count and density over straight and slanted position of the detected vehicles with distance less than 30 meters

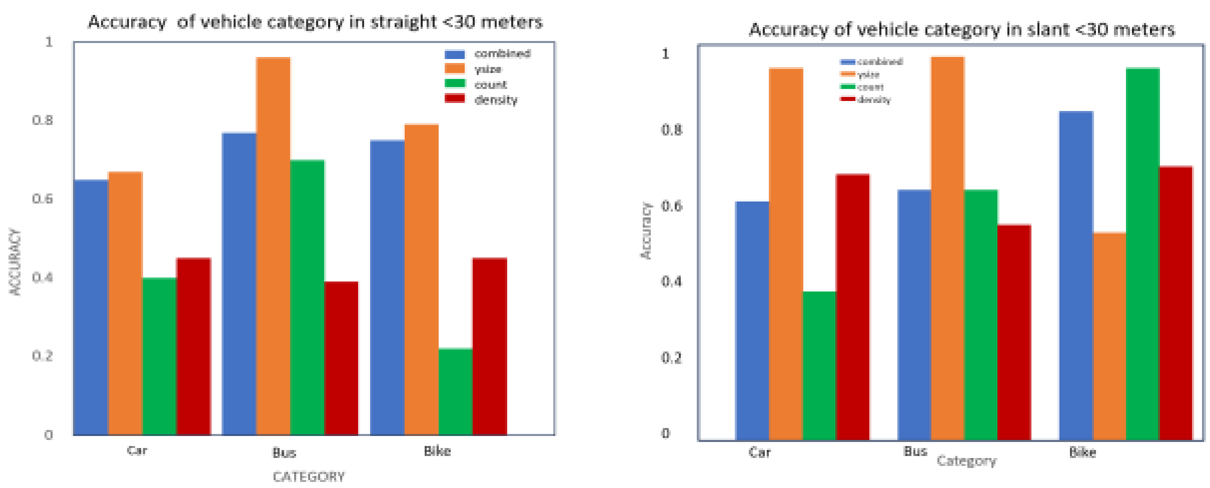

Accumulation of data over continuous frames

To enhance the richness of radar point cloud data and potentially improve classification accuracy, we investigated the accumulation of data over three consecutive frames. This approach aimed to address the issue of low point count observed in single-frame analysis, particularly for smaller objects or those at greater distances. By integrating information from multiple frames, we hypothesized that the resulting point clouds would become denser and more informative. Analysis of the accumulated data revealed notable improvements in point count for all vehicle categories, especially for those located within 30 meters of the radar sensor. This increase in point density enhanced the discriminative power of the y-size (length) feature, leading to improved accuracy in classifying vehicles based on their size alone. However, challenges remained in achieving consistently high accuracy for all classes using point count, indicating that this feature alone may not be sufficient for robust classification. Density calculations on the accumulated data also presented some limitations. In certain scenarios, such as bus detection in straight positions, bike detection in slanted positions, and car detection in slanted positions within 30 meters, density-based classification yielded lower accuracy. This suggests that density, while informative, may not be a universally reliable feature for distinguishing between all vehicle types across different orientations and distances. Furthermore, the manual determination of thresholds for y-size, point count, and density proved to be a cumbersome and potentially inaccurate process.

Figure 5: Accuracy comparison of 3 Frame data accumulation for combined, ysize, point count and density of the detected vehicles with distance less than 30 meters and greater than 30 meters

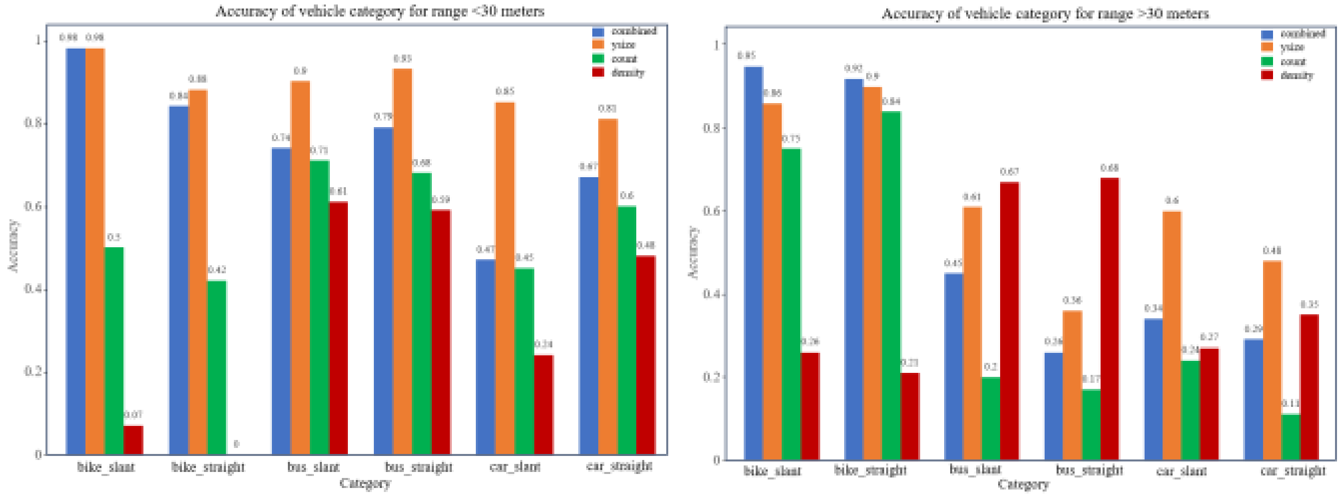

In Figure 5, we observed variations in accuracy for vehicles detected in straight and slanted positions relative to the radar, indicating that a single set of thresholds may not be optimal for all scenarios. This limitation highlights the need for a more adaptive and data-driven approach to threshold selection or, alternatively, the exploration of classification methods that do not rely on manual thresholding. Overall, the accumulation of data over continuous frames demonstrated potential for enhancing radar-based vehicle classification, particularly by increasing point count and improving the discriminative power of y-size. However, the challenges encountered in achieving consistent accuracy across all classes and scenarios using point count and density underscore the need for further refinement of feature selection and classification methodologies. These results pave the way for the exploration of more sophisticated approaches, such as machine learning, that can leverage a wider range of features and adapt to varying environmental conditions.

Formation of Features

Given the limitations of using basic object dimensions, point count, and density for accurate vehicle classification, a more comprehensive feature extraction process was undertaken. To capture the nuanced characteristics of different vehicle types, we collected radar data under diverse scenarios, including both approaching and receding trajectories in the same and parallel lanes. This data, focusing on individual vehicles (bus, car, and bike) in isolation, was then meticulously analysed to identify potential features that could enhance classification accuracy. A wide array of features was considered, encompassing both geometric and statistical properties of the radar point clouds. These features included: Point count, Y-size (length), X-size (width), Density of points in cluster, Diagonal length of cluster, Perimeter of cluster, Area of cluster. A custom Python algorithm was developed to extract these features from both single-frame and accumulated (three-frame) radar data. The resulting feature values were then visualized and analyzed to assess their discriminatory power. This analysis revealed that two features density of points in cluster as shown in Figure 6 and point count as shown in Figure 7 was not particularly useful for distinguishing between vehicle types. These features exhibited high variability within each category and significant overlap between categories, rendering them ineffective as discriminators.

International Journal for Multidisciplinary Research (IJFMR)

E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

Figure 6: Density of the vehicle categories in different positions (Straight/Slant) for range less than 30 meters

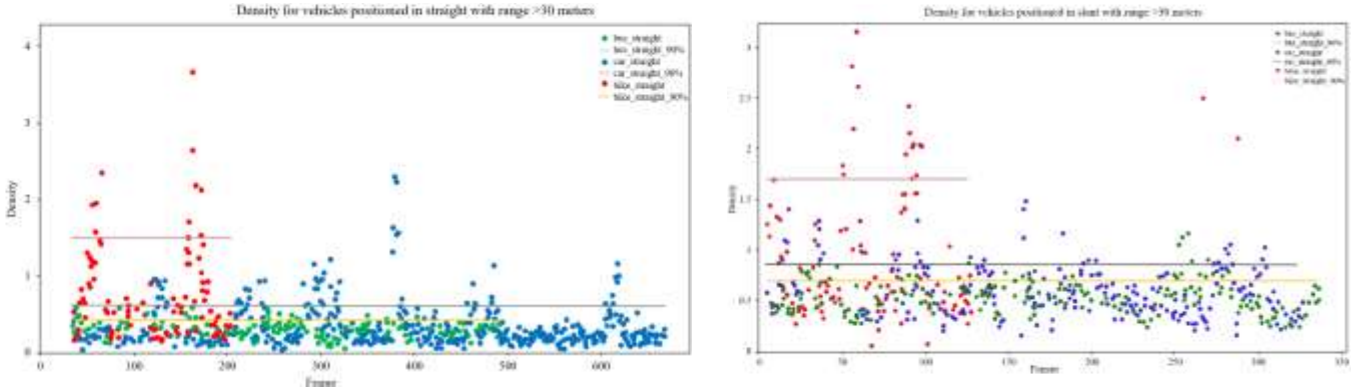

Figure 7: Point Count of the vehicle categories in different positions (Straight/Slant) for range less than 30 meters

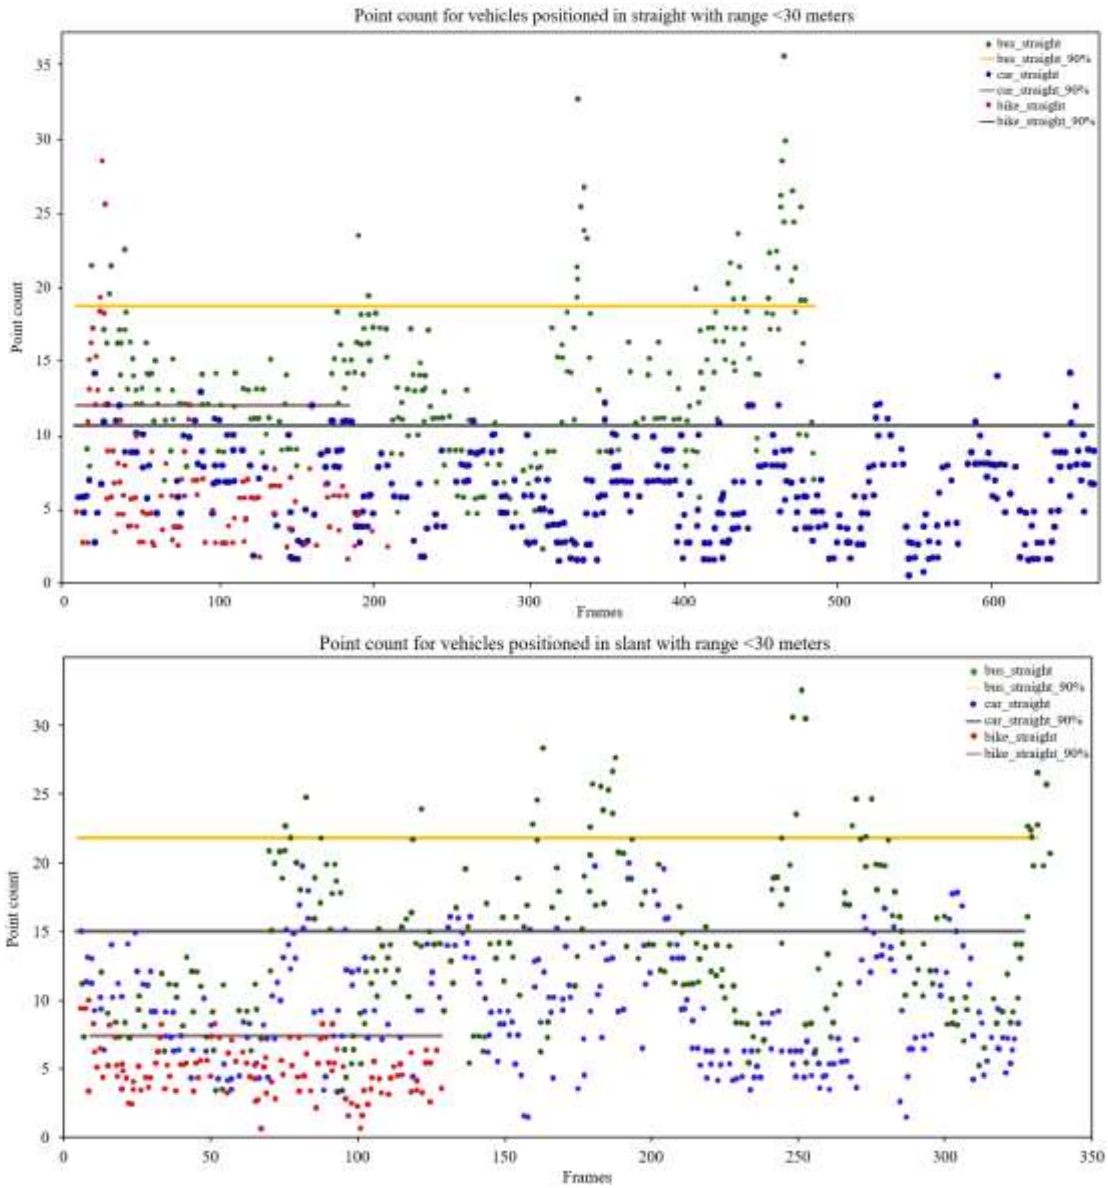

However, the remaining features, including y-size (length) as shown in Figure 8, diagonal length as shown in Figure 11, perimeter of cluster as shown in Figure 10 and area of cluster as shown in Figure 9 demonstrated greater potential for differentiation. These features captured distinct characteristics of different vehicle types, such as their overall size, shape, and distribution of points within the cluster.

Figure 8: YSize of the vehicle categories in different positions (Straight/Slant) for range less than 30 meters

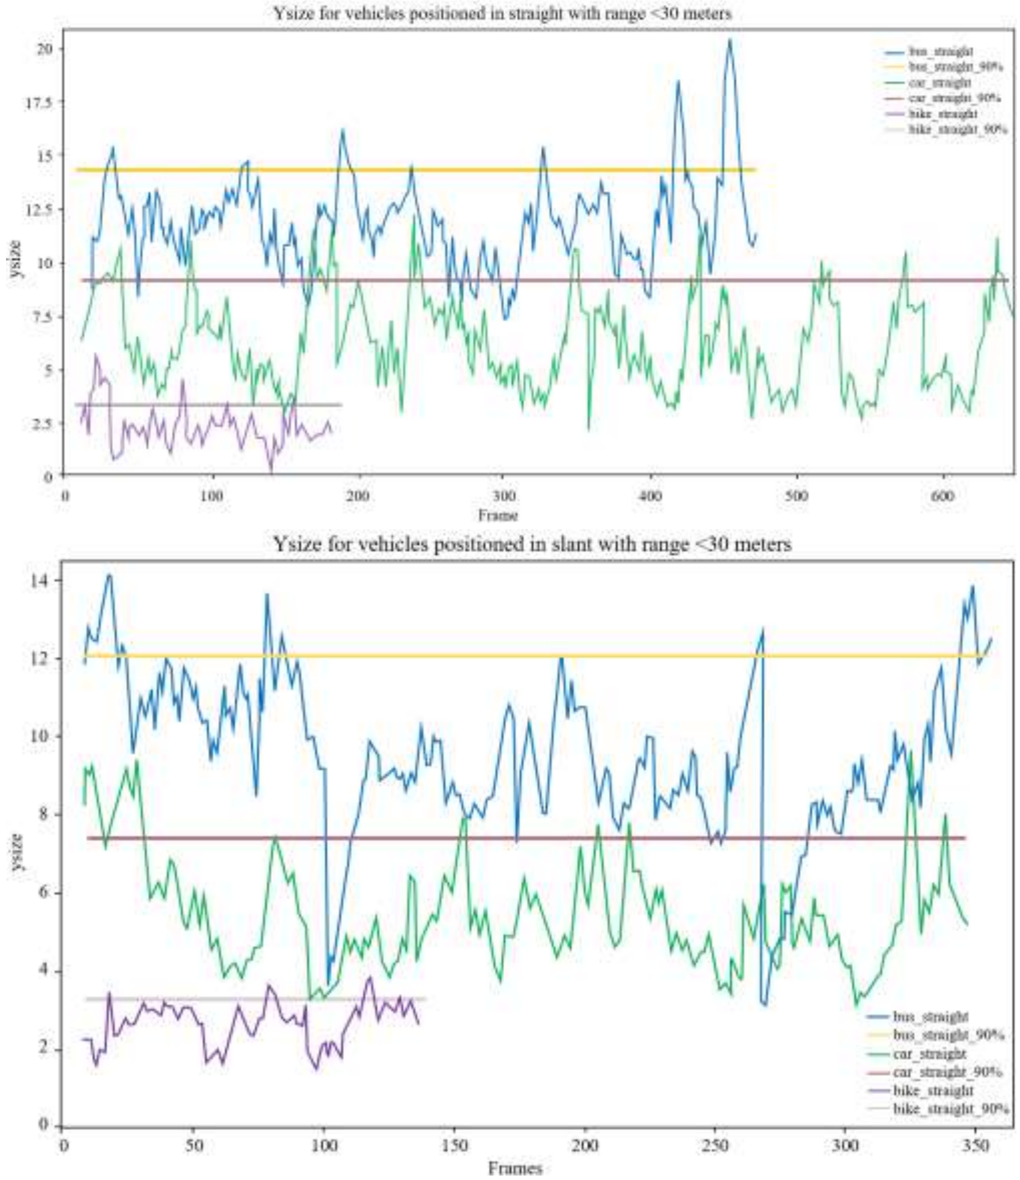

E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

Figure 9: Area of the vehicle categories in different positions (Straight/Slant)

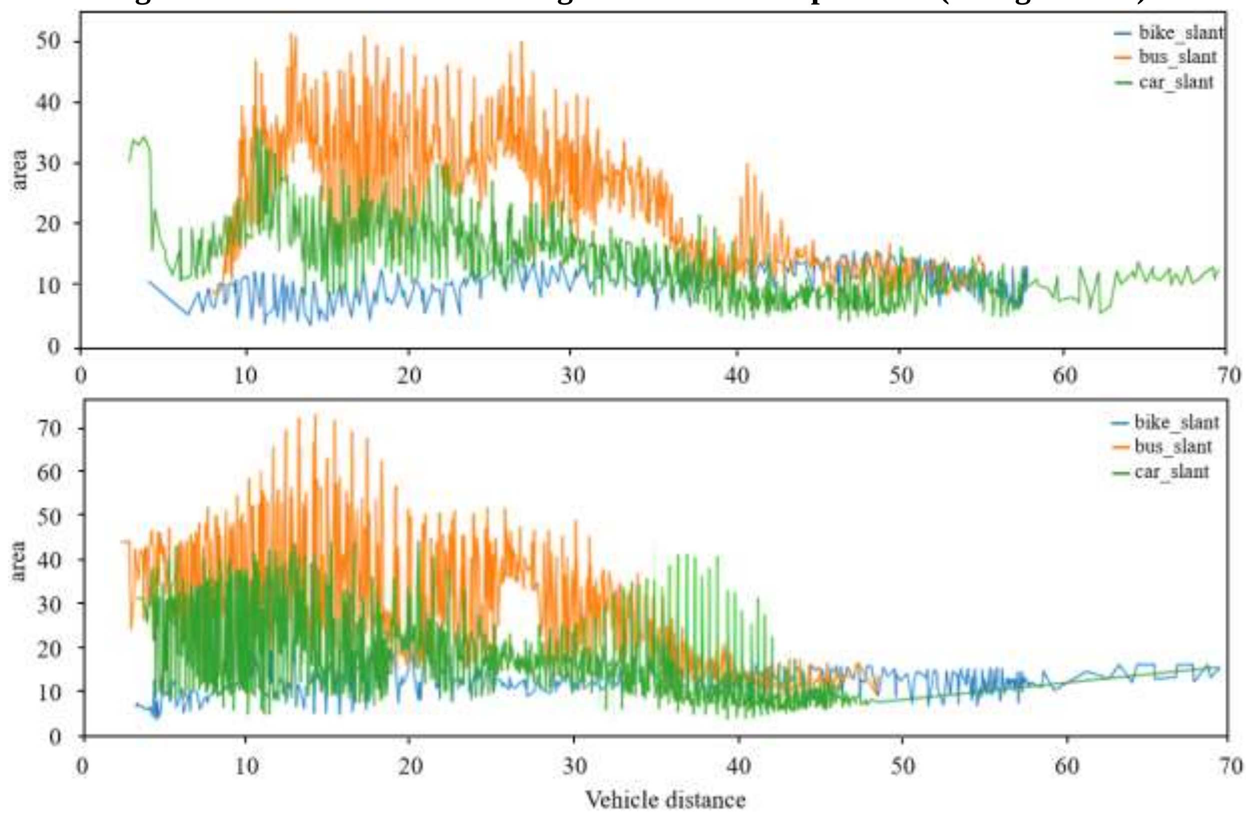

Figure 10: Perimeter of the vehicle categories in different positions (Straight/Slant)

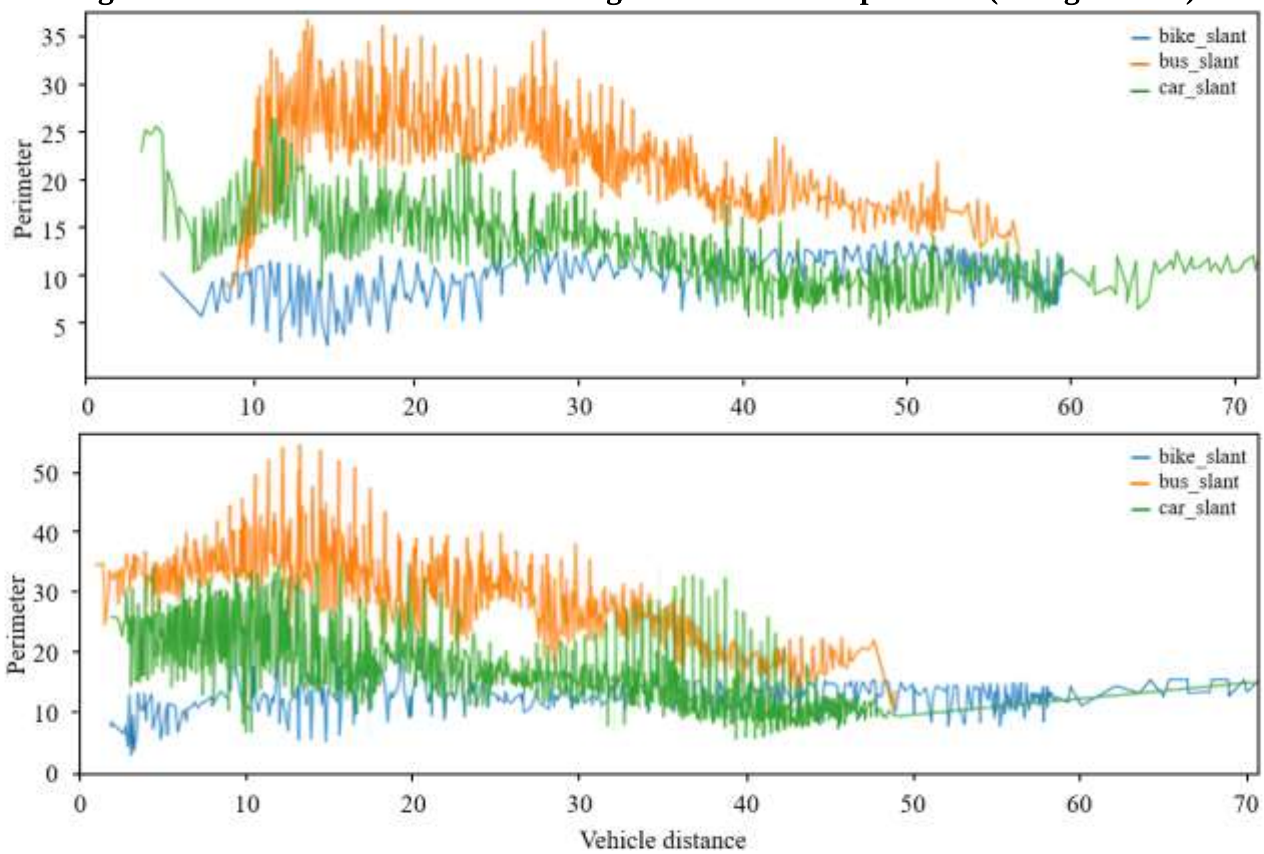

Figure 11: Diagonal Length of the vehicle categories in different positions (Straight/Slant)

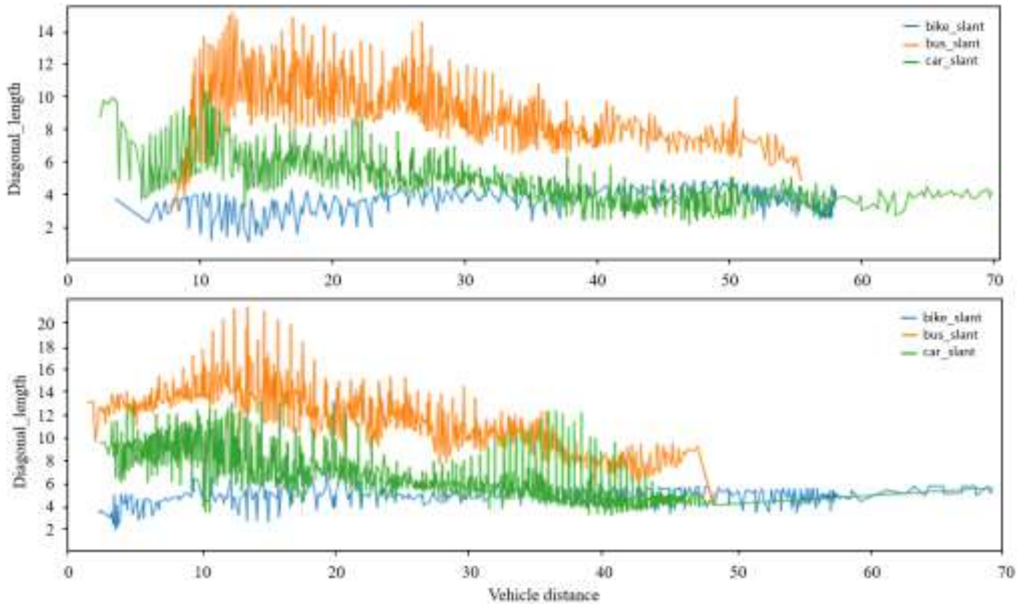

Notably, the features calculated from accumulated frames were found to be more informative than those from single frames as shown in Figure 12, highlighting the importance of incorporating temporal information into the feature extraction process. This comprehensive feature formulation process allowed us to identify a set of key features that could be leveraged by machine learning algorithms to achieve more accurate and robust vehicle classification. The emphasis on accumulated frames and the selection of features that capture both geometric and statistical properties of the radar point clouds laid the groundwork for the subsequent machine learning phase of the research.

Figure 12: Manual threshold accuracy for different vehicle categories for range less than 30 meters and greater than 30 meters

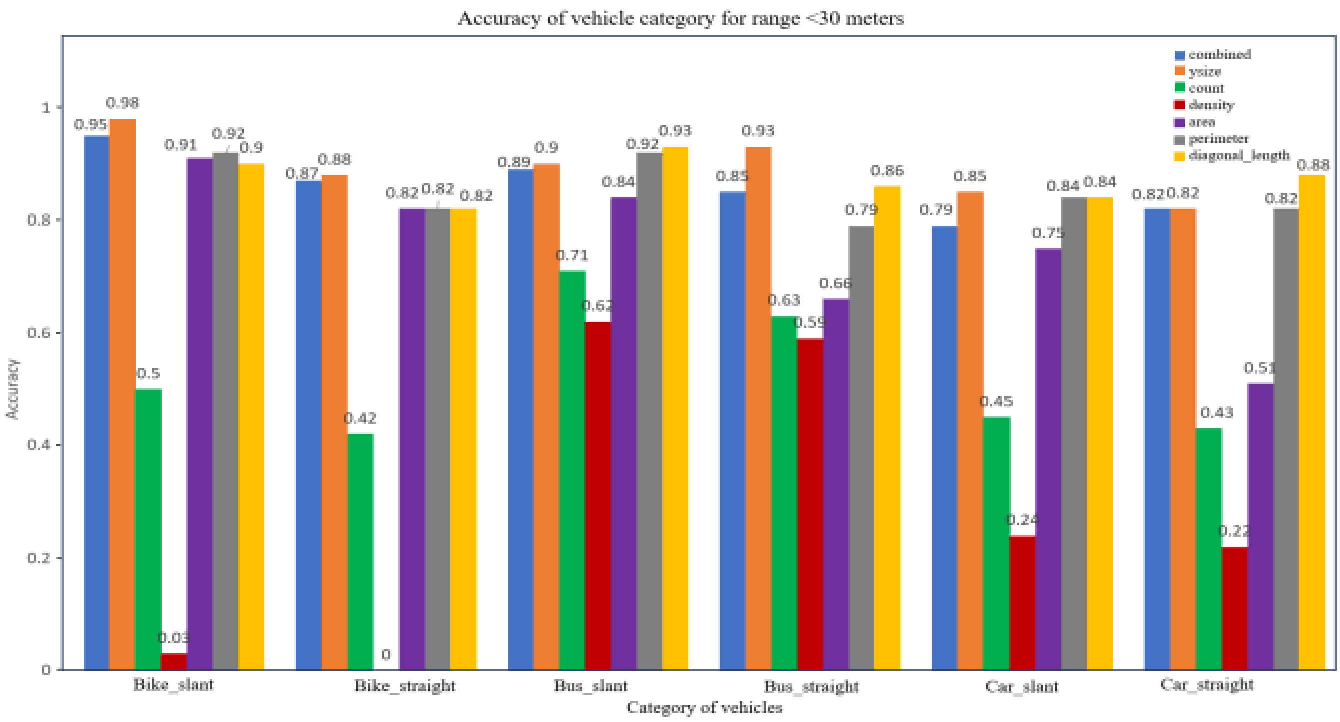

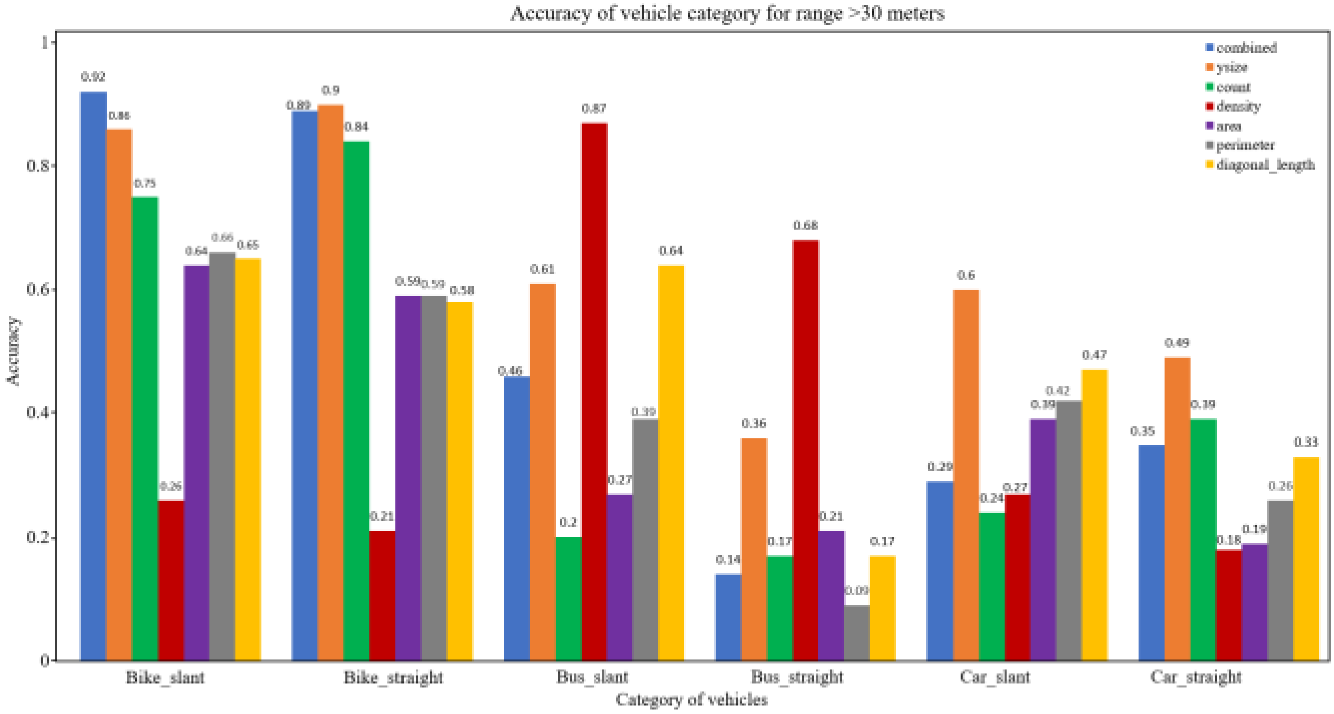

Machine Learning

To overcome the limitations of manual thresholding and leverage the formulated features effectively, we employed supervised machine learning algorithms for vehicle classification. In the initial phase, we trained and cross-validated various ML algorithms (Decision Tree, Random Forest, LightGBM, SVM_SVC) on both single-frame and accumulated-frame data, ensuring each frame contained only one vehicle. To mitigate class imbalance, we applied the Synthetic Minority Over-sampling Technique (SMOTE) and normalized the data using MinMax scaling. The initial results showed promising cross-validation accuracies of 0.89 for single-frame and 0.91 for accumulated-frame data. LightGBM emerged as a top performer, achieving 0.89 accuracy at the 50th percentile for single-frame and 0.94 for accumulated-frame data. Conversely, Decision Tree and SVM_SVC models exhibited the lowest cross-validation accuracies for single-frame (0.86) and accumulated-frame (0.91) data, respectively. However, these results were obtained in a controlled environment and required validation in real-world scenarios. The precision, recall, f1-score and accuracy results for both single and accumulated frames are shown in Table 3, the cross-validation score for single frame and accumulated frame is shown in Figure 13.

Table 3: Precision, Recall, F1-score and Accuracy results for test performed in controlled environment

	Frame	Precision	Recall	F1-score	Accuracy
Bike	Single	0.90	0.95	0.92	0.89
Bus	Single	0.91	0.92	0.91	
Car	Single	0.88	0.82	0.85	
Bike	3 Frame	0.86	0.86	0.86	0.91
Bus	3 Frame	0.92	0.82	0.87	
Car	3 Frame	0.91	0.94	0.93	

Figure 13: Single Frame (Left Plot) & Accumulated Frame (Right Plot) Cross-validation accuracy results on different ML Models

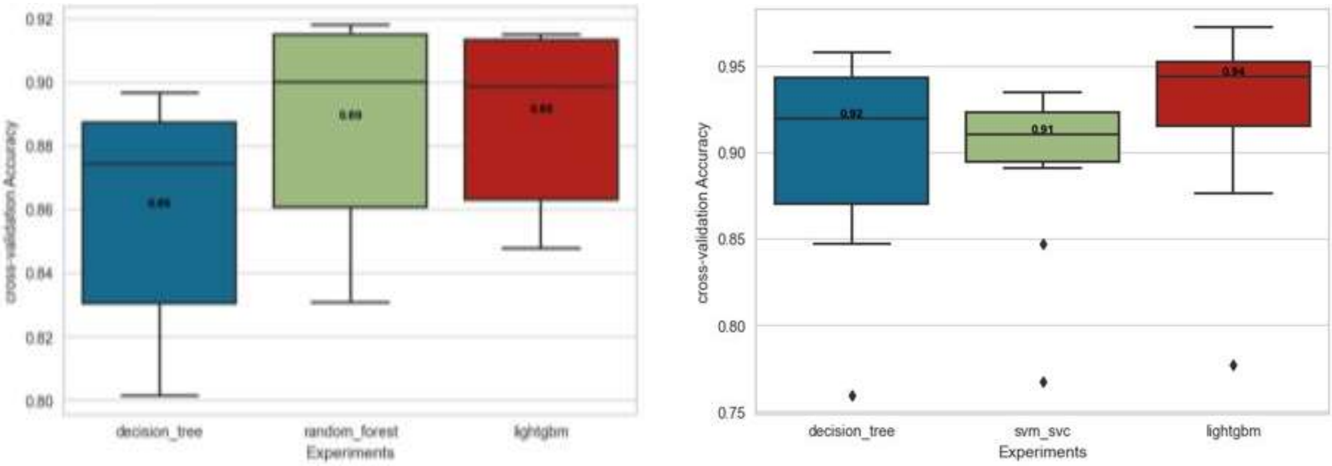

Table 4 shows the results for real-world testing on city and highway roads with multiple vehicles in the scene revealed a significant drop in accuracy to 0.59 when compared to manual annotations. This discrepancy highlighted the need for a more extensive and diverse training dataset to improve the model's generalization capabilities.

Table 4: Precision, Recall, F1-score and Accuracy results for test performed in real-world scenarios

	Frame	Precision	Recall	F1-score	Accuracy
Bike	3 Frame	0.63	0.50	0.56	0.59
Bus	3 Frame	0.16	0.51	0.24	
Car	3 Frame	0.65	0.66	0.66	

To address this, we merged multiple iterations of live data, incorporating a wider range of scenarios and vehicle types. We then retrained the models using Gradient Boosting and a Voting Classifier (combining predictions from multiple classifiers). Hyperparameter tuning was performed using GridSearch, and k-fold cross-validation with 10 folds was employed to assess model robustness. The re-trained models demonstrated improved performance, achieving an average cross-validation accuracy of 0.74 for single-frame and 0.72 for accumulated-frame data as shown in Table 5. Gradient Boosting consistently outperformed other models in both single-frame and accumulated-frame scenarios as shown in Figure 14, exhibiting the least interquartile range in cross-validation scores and the highest 50th percentile accuracy. These results underscore the effectiveness of machine learning in leveraging the formulated features for vehicle classification and the importance of a comprehensive and diverse training dataset. While the accuracy improvement is notable, further enhancements can be achieved by expanding the training data and exploring additional machine learning techniques.

Table 5: Re-trained results of Precision, Recall, F1-score and Accuracy for test performed in real-world scenarios

	Frame	Precision	Recall	F1-score	Accuracy
Bike	3 Frame	0.83	0.79	0.81	0.77
Bus	3 Frame	0.76	0.78	0.77	
Car	3 Frame	0.72	0.74	0.73	

Figure 14: Accumulated Frame Cross-validation accuracy results on different ML Models re-trained with real-world data

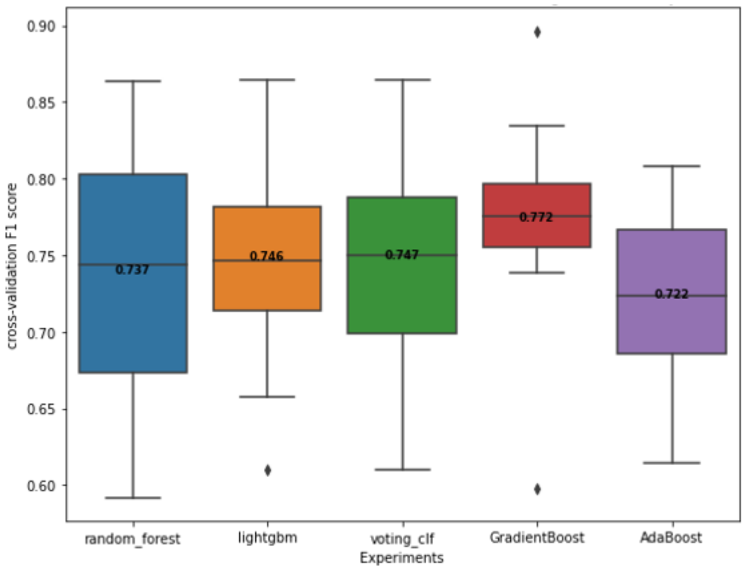

5. Conclusion

In conclusion, this research delves into the potential of automotive radar for accurate and real-time vehicle classification, a critical task for ADAS and autonomous vehicles. Through a comprehensive exploration of various classification approaches, we have demonstrated the feasibility and effectiveness of utilizing radar data for this purpose. While traditional clustering algorithms like DBSCAN and Agglomerative clustering faced challenges due to the inherent variability and noise in radar point clouds, the formulation of a diverse set of features, including object dimensions, point count, density, and geometric properties, significantly enhanced classification accuracy. The accumulation of radar data over continuous frames further improved the discriminative power of these features, particularly for smaller objects and those at greater distances. Machine learning algorithms, especially Gradient Boosting and Voting Classifier, proved to be highly effective in leveraging these features for accurate vehicle classification. The iterative process of data collection, annotation, model training, and evaluation highlighted the importance of a comprehensive and diverse training dataset for achieving robust performance in real-world scenarios. While this research demonstrates the potential of radar-based vehicle classification, it also acknowledges the limitations and challenges that remain. Environmental factors, occlusions, and the inherent characteristics of radar data can still impact classification accuracy. Future research directions include exploring advanced signal processing techniques, sensor fusion with other modalities like cameras and LiDAR, and the development of more sophisticated machine learning models to further enhance the robustness and reliability of radar-based vehicle classification systems. Overall, this study contributes valuable insights into the application of radar technology for vehicle classification and paves the way for the development of more advanced and reliable ADAS and autonomous driving systems. By harnessing the power of radar data and machine learning, we can strive towards a future where vehicles are equipped with robust perception capabilities, leading to safer and more efficient transportation systems.

6. References

1. Srivastav and S. Mandal, "Radars for autonomous driving: A review of deep learning methods and challenges," arXiv preprint arXiv:2306.09304, 2023.
2. M. Ulrich, C. Glaser, and F. Timm, "Deepreflects: Deep learning for automotive object classification with radar reflections," Robert Bosch GmbH, Stuttgart, Germany.
3. S. Agrawal, S. Bhanderi, and G. Elger, "Semi-automatic annotation of 3D radar and camera for smart infrastructure-based perception," IEEE Access, 2023.
4. Y. Wang, G. Wang, H.-M. Hsu, H. Liu, and J.-N. Hwang, "Rethinking of radar's role: A camera-radar dataset and systematic annotator via coordinate alignment," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, 2021.
5. H. D. Ghael, L. Solanki, and G. Sahu, "A review paper on raspberry pi and its applications," International Journal of Advances in Engineering and Management (IJAEM), vol. 2, no. 12, pp. 225–227, 2020.
6. P. Visconti, B. Sbarro, P. Primiceri, R. de Fazio, and A. Lay-Ekuakille, "Design and testing of electronic control system based on STM X-Nucleo board for detection and wireless transmission of sensors data applied to a single-seat Formula SAE car," International Journal of Electronics and Telecommunications, vol. 65, no. 4, pp. 671–678, 2019.
7. V. Margapuri and M. Neilsen, "Prototype implementation of temperature control system with CAN and FreeRTOS on STM32F407 discovery boards," in Proceedings of 32nd International Conference on Computer Applications in Industry and Engineering, 2019, pp. 130–139.
8. T. M. and M. Prabhakar, "An efficient high gain DC-DC converter for automotive applications," International Journal of Power Electronics and Drive Systems (IJPEDS), vol. 6, no. 2, pp. 242–252, 2015.
9. S. T. Woo, Y. B. Park, J. H. Lee, C. S. Han, S. Na, and J. Y. Kim, "Angle sensor module for vehicle steering device based on multi-track impulse ring," Sensors, vol. 19, no. 3, p. 526, 2019.
10. X. Gao, G. Xing, S. Roy, and H. Liu, "Experiments with mmwave automotive radar test-bed," arXiv preprint arXiv:1912.12566, 2022.
11. K. Bindra and A. Mishra, "A detailed study of clustering algorithms," in 2017 6th International Conference on Reliability, Infocom Technologies and Optimization (ICRITO) (Trends and Future Directions), 2017, pp. 371–376.
12. F. Kraus, N. Scheiner, W. Ritter, and K. Dietmayer, "The radar ghost dataset - an evaluation of ghost objects in automotive radar data," in 2021 IEEE Intelligent Vehicles Symposium (IV), 2021, pp. 1–8.
13. O. Schumann, M. Hahn, N. Scheiner, F. Weishaupt, J. F. Tilly, J. Dickmann, and C. Wohler, "Radarscenes: A real-world radar point cloud data set for automotive applications," arXiv preprint arXiv:2104.02493, 2021.
14. R. Schneider and J. Wenger, "High resolution radar for automobile applications," Advances in Radio Science, vol. 1, pp. 105–111, 2003.

15. B. Han, "Comparison of different machine learning algorithms in classification," Journal of Physics: Conference Series, vol. 2037, no. 1, p. 012064, 2021.