2023 IEEE International Conference on Robotics and Automation (ICRA 2023)
May 29 - June 2, 2023. London, UK

Towards Unsupervised Filtering of Millimetre-Wave Radar Returns for Autonomous Vehicle Road Following

Dean Sacoransky, Joshua A. Marshall, and Keyvan Hashtrudi-Zaad

***Abstract*— Path planning and localization in low-light and inclement weather conditions are critical problems facing autonomous vehicle systems. Our proposed method applies a single modality, millimetre-wave radar perception system for the detection of roadside retro-reflectors. Radar-based perception tasks can be challenging to perform due to the sparse and noisy nature of radar data. We propose the use of an unsupervised learning approach for filtering radar point clouds through Density-Based Spatial Clustering of Applications with Noise (DBSCAN). The DBSCAN algorithm segments retro-reflector points from noise points, thus providing the autonomous vehicle with a predicted path for the road ahead. We tested the approach via indoor experiments that make use of Continental's ARS 408 radar, a mobile Husky A2000 robot, and a Vicon motion capture system for ground truth validation. The experimental results of the proposed system demonstrated a classification accuracy of 84.13 % and F1 score of 83.71 %.**

I. INTRODUCTION

Current autonomous vehicle (AV) perception systems fail to perform in low-light and poor-weather conditions due to the limitations of cameras and LiDAR [1]. Radar has emerged as an essential component in the AV sensor suite due to its low cost, ability to perceive the environment in inclement weather conditions, and its advantage in direct measurement of relative velocity [2]. The detection of physical infrastructure (i.e., passive retro-reflectors) could be the key to AV localization and road-following challenges under extreme weather conditions, such as fog, rain, and snow.

Automotive radar sensors are currently deployed in many modern vehicles for Advanced Driver Assistance Systems (ADAS). ADAS functionalities include automatic cruise control, blind spot detection, emergency braking, and lane departure warnings. With the emergence of improved radar technology, greater computation speed, and more robust artificial intelligence systems, AV researchers are shifting their focus from ADAS to fully-automated driving [1].

To deploy reliable radar-based autonomous systems, a significant performance increase is required for radar sensors because of the noise and sparsity of radar point clouds. However, the signal processing pipeline for radar has remained the same for years, relying on classical approaches such as the Constant False Alarm Rate (CFAR) algorithm [3].

This work was supported in part by Natural Science and Engineering Research Council of Canada (NSERC) under Discovery grant numbers RGPIN-05609 and RGPIN-2015-04025, as well as by the NSERC CREATE program for Building Trust in Connected Autonomous Vehicles (TrustCAV) under grant number 542999-2020.

D. Sacoransky, J. .A. Marshall, and K. Hashtrudi-Zaad are with the Department of Electrical & Computer Engineering and the Ingenuity Labs Research Institute, Queen's University, Kingston, ON K7L 3N6 Canada, {17drs3,joshua.marshall,khz}@queensu.ca

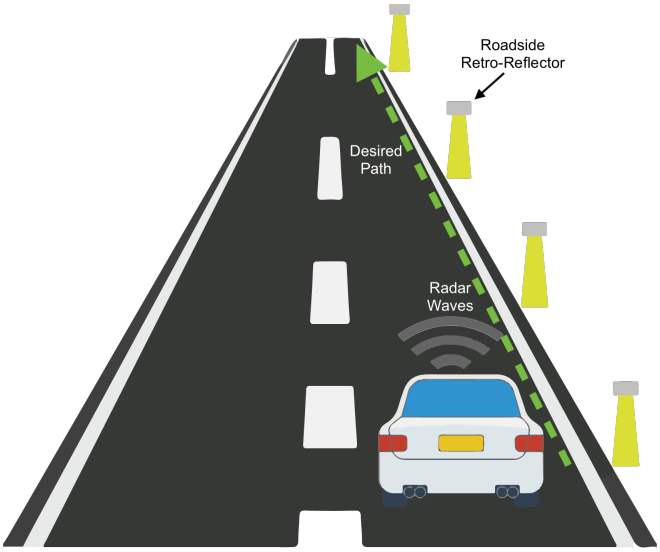

Fig. 1. Conceptual diagram illustrating an autonomous vehicle road following system, where reflective markers are installed on the side of the road to enhance the radar-based landmark detection and path prediction.

Researchers have very recently begun exploring the applications of deep learning to enhance radar signal processing for perception tasks such as object detection and segmentation, as explained in section I-A [4], [5], [6]. The Oxford Radar RoboCart data set [7] and nuScenes data set [8] are often used as benchmarks to evaluate the performance of radar-based perception systems. However, there is no publicly available data set that includes retro-reflector information, and thus, to our knowledge, no previous work has used radar sensors and unsupervised learning to identify retro-reflectors.

We suggest that, by installing reflective roadside infrastructure, one might use radar to robustly predict the path of a road despite the inherent noise and sparsity of the radar's acquired returns. A conceptual illustration of the proposed setup is provided in Fig. 1, where a radar is mounted to the front of a moving vehicle and roadside retro-reflectors are used to aid with landmark identification for path planning and road following.

In this work, we make use of unsupervised clustering to exploit the radar cross section (RCS) and spatial pattern of retro-reflectors in the vehicle's environment. Through the application of unsupervised learning techniques, our proposed method can deduce patterns from the collected radar data without the need for a training data set and without the concern of over fitting a model to that data set.

A. Related Work

This section summarizes related work, from radar data processing to point cloud segmentation, unsupervised learning, and the use of reflective lane markers with radar.

1) Radar Data Processing Techniques: The signal processing pipeline for radar has recently gained attention in the literature. Radar images are often noisy and require processing techniques to filter out clutter and isolate objects of interest. Feature detectors, such as CFAR, are used to identify key points from radar returns [3]. CFAR filters out radar detections whose returns are below a predefined power threshold. Feature descriptors, such as Oriented FAST and rotated BRIEF (ORB), have been used to obtain information on landmarks such as shape, orientation, location, and motion [9]. The Fast Fourier Transform (FFT) has been used as a classical radar processing method for transforming data to the range-Doppler domain to extract relative range and velocity information [10].

Modern radar processing techniques that use deep learning have emerged in recent years. Brodeski *et al.* [11] developed a data-driven deep learning approach to radar signal processing, thus eliminating the need for parameter selection in methods such as CFAR. Adolfsson *et al.* [3] extended CFAR to CFEAR (Conservative Filtering for Efficient and Accurate Radar Odometry) for accurate incremental odometry estimation and surface estimation of unknown environments. CFEAR is a conservative filtering approach that distinguishes signals from noise through k-strongest filtering. This method reduces false positive radar detections at the expense of missing landmarks with weaker returns [3]. Barnes *et al.* [12] used deep learning-based methods in an attempt to find optimal radar key points rather than using off-the-shelf key point extractors such as ORB or SURF. These methods, however, do not apply unsupervised learning for filtering noise points from objects of interest.

2) Segmentation of Radar Point Cloud Data: Many of the current radar segmentation techniques employ supervised learning on radar data, which is restricted to segmenting a finite set of classes that exist within radar data sets (e.g., cars, trucks pedestrians, and busses) [4], [5], [6]. Shah *et al.* [1] fused radar data with LiDAR and HD maps for effective trajectory prediction. Radar points were clustered to reduce noise and multiple radar sweeps were aggregated to provide temporal context and recover full velocity information [1]. Danzer *et al.* [5] explored learning-based 3D object detection on automotive radar data. Ouaknine *et al.* [6] proposed novel architectures for the semantic segmentation of radar data using the CARRADA dataset.

3) Unsupervised Learning Applied to Radar Data: In the context of automotive radar data, clustering algorithms can be used to identify groups of data points that correspond to different objects or obstacles in the environment. Jin *et al.* [4] used a Gaussian mixture model to perform semantic segmentation for pedestrians and cars. Kellner *et al.* [13] presented a density-based algorithm to cluster pedestrians, vehicles, barriers, and clutter from high resolution radar data with improved execution time over the classic DBSCAN.

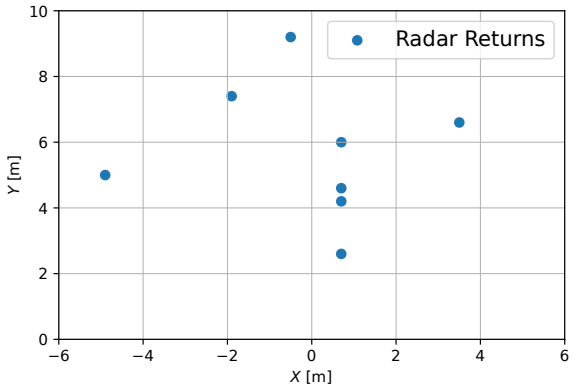

Fig. 2. 2D radar point cloud representation of the utilized indoor testing space, corresponding to that of Fig. 3. The four collinear points, which represent retro-reflectors, are equidistant in truth, but appear non-equidistant in the radar output. This figure illustrates the challenge of noise and sparsity associated with radar data.

Scheiner *et al.* [14] applied a two stage filtering approach to reduce background noise in radar data. Specifically, they proposed a modified version of the DBSCAN algorithm with specialized distance metrics. These methods do not consider the classification or segmentation of retro-reflectors from radar data.

4) Reflective Lane Markers: Greisman *et al.* [10] proposed conductive paint as a reflective lane marker for lateral vehicle localization. They conducted experiments to detect the lane markers through snow coverage using automotive radar. We build on the idea of using reflective infrastructure to aid radar perception systems for localization and road following purposes.

II. Motivation

Perhaps the biggest challenge facing radar-based perception systems is the sparsity and noise of the data, which is illustrated, for example, in Fig. 2. Given these issues, the question arises whether radar sensors can be used to reliably understand driving environments. We suggest that by installing reflective roadside infrastructure and using unsupervised clustering techniques, the path of a road can be predicted, despite the inherent noise and sparsity of the radar's acquired returns.

Retro-reflectors are designed to reflect radar waves directly back towards the radar antenna for stable and powerful returns. The RCS of an object depends on factors such as material, shape, orientation, and surface smoothness [15]. The retro-reflectors used in this research were metallic, diamond reflectors consisting of eight mutually perpendicular, intersecting flat surfaces, as shown in Fig. 3. We hypothesize that RCS information serves as a valuable feature to distinguish retro-reflectors (i.e., the object of interest) from non-reflector points (i.e., what we call noise). These retro-reflectors are used to effectively localize a vehicle within a "lane" and create a path for the vehicle to follow. This sensor-infrastructure setup has the potential to be robust to

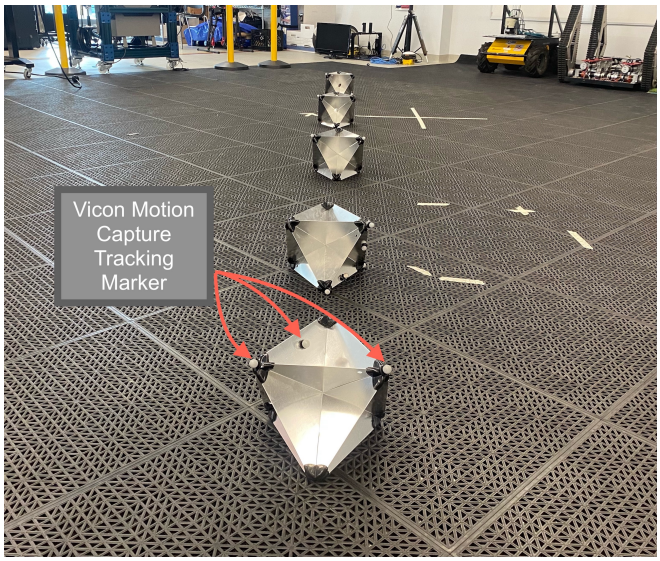

Fig. 3. Brewers' Aluminum, 12-in. radar reflectors with plastic corners. Five reflectors are set up to form a straight-line pathway, with Vicon Motion Capture markers installed on each reflector for ground truth data collection.

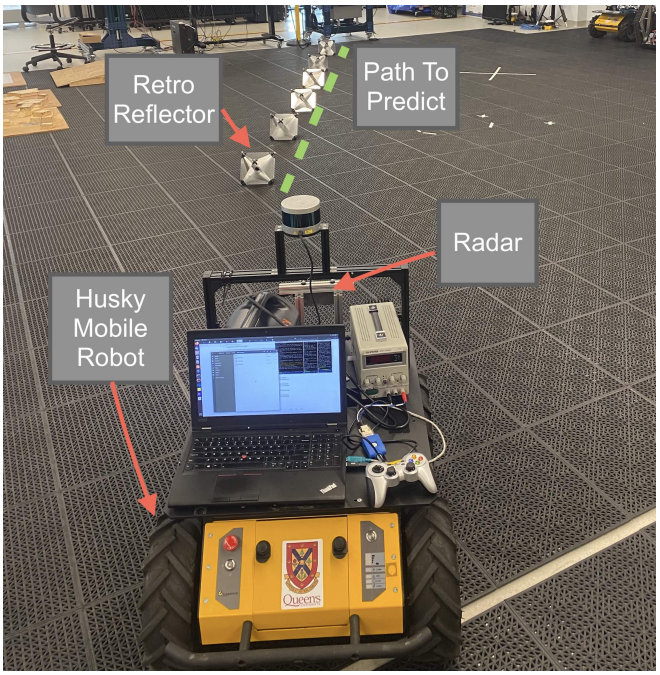

Fig. 4. Indoor testing setup, where the radar was mounted on a Husky mobile robot. Five retro-reflectors were organized in the environment, as shown. The clustering algorithm aims to filter out returns that do not correspond to retro-reflectors, which mark the roadway, as seen in Fig. 5.

adverse weather and lighting conditions due to the lower frequency operation (77 GHz) of automotive radar sensors [2]. Therefore, we hypothesize that radar sensors will be able to detect such roadside reflectors in darkness and low-visibility weather conditions, such as fog and heavy snow. Cameras, for example, are unable to perceive objects under these conditions due to their operation in the visual spectrum [1]. A similar problem exists for LiDAR in inclement weather, which operates in the infrared spectrum.

III. Experimental Setup

The research presented by this paper used a Continental ARS 408 radar unit to detect retro-reflector pathways and filter noise points with unsupervised clustering. This radar unit operates at 77 GHz and uses a pulse compression radar modulation scheme as basic principle for its measurements with a field of view of ± 60 degrees. Radar data was communicated over a CAN interface at a rate of 500 kbit/s. We interfaced with the radar unit through a Robot Operating System (ROS) Melodic driver in C++. The ROS system provided a radar feature vector which was newly evaluated every time step. Radar data was collected as `rosbag` data files and processed offline.

The radar was configured to infer a maximum of eight (8) objects per cycle by using Continental's proprietary cluster and object filter (*FilterCfg*) [16]. This filter reduced unwanted noise returns and promoted the detection of objects with stable returns (i.e., including retro-reflectors). Through trial and error, eight was found to be a reasonable number for this filter because it allowed for retro-reflector returns to be included, while also minimizing noisy returns.

The test vehicle used in this experiment was a Clearpath Husky A200 mobile robot with a mounted radar unit as shown in Fig. 4. Five retro-reflectors, separated by 1 m each, formed a straight line in the indoor testing environment to resemble a straight road scenario. For each experimental run, the Husky robot drove in a straight line at 0.5 m/s. We started recording data when the Husky begun driving forwards. Data collection stopped right before the Husky passed the first reflector, because the system, as configured, was intended to perceive the reflector pathway as the vehicle approaches from a distance. The Husky travelled approximately 1.7 m for two trial runs lasting approximately 4.7 s each. Illustrations of DBSCAN clustering on sequential radar data from our experiments are shown in Fig. 5.

Testing was conducted indoors to make use of a Vicon motion capture apparatus for performance validation. This provided millimetre-level precision motion capture and positional data for the Husky and for each retro-reflector. We were able to calculate the relative positions of the Husky to each retro-reflector to be used as ground truth values. Small cylindrical reflective markers were placed on the mobile robot and each retro-reflector, as seen in Fig. 3, to enable positional data collection at 100 Hz.

IV. Proposed Method

This section provides a detailed description of our approach, starting with the data pipeline in III-A and an explanation of the unsupervised clustering approach in III-B.

A. Data Collection and Standardization

Raw radar signals are automatically converted to features through Continental's proprietary on-device processing [16]. We collect a feature vector comprised of 2D Cartesian spatial

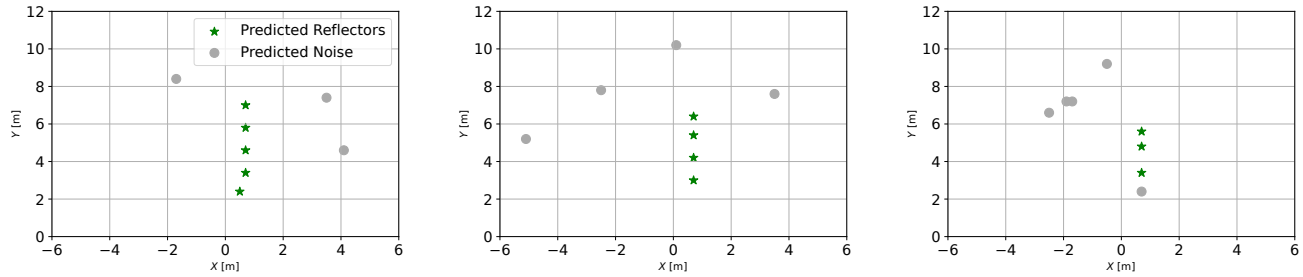

Fig. 5. Sequence of ordered DBSCAN clustering predictions 1 s apart (from left to right), where coordinate $(0, 0)$ is the origin of the radar frame. DBSCAN clustering provided effective segmentation of retro-reflector points from noise points and the linear pathway of the retro-reflectors was successfully predicted by the model. However, in the rightmost figure the model falsely predicts one reflector point as noise.

coordinates and RCS values, denoted α, of inferred objects. In general, the feature vector $\mathbf{f}_i$ is given by

$$\mathbf{f}_i = \begin{bmatrix} x_{i0} & y_{i0} & \alpha_{i0} & \dots & x_{in} & y_{in} & \alpha_{in} \end{bmatrix} \quad (1)$$

where $i = 0, 1, \dots, m$ is the time index and $n \geq 0$ is the number of inferred objects. RCS values were derived from the classic radar equation, which states [15] that

$$P_r = \frac{P_t G^2 \lambda^2 \alpha}{R^4 (4\pi)^3}, \quad (2)$$

where R is the range from the radar to the target, P_t is the transmit power, G is the antenna gain, λ is the radar wavelength, and P_r is the power of the radar return. We rearrange (2) to solve for α

$$\alpha = \frac{P_r R^4 (4\pi)^3}{P_t G^2 \lambda^2}. \quad (3)$$

Each element f_{ij}, $i = 0, 1, \dots, m$ and $j = 0, 1, \dots, n$, of the feature vector $\mathbf{f}_i$ is normalized by

$$z_{ij} = \left(f_{ij} - \bar{f}_i\right) / s_i \quad (4)$$

where $\bar{f}_i$ is the mean and s_i is the standard deviation over the n inferred objects, resulting in a normalized feature vector $\mathbf{z}_i$. This process normalizes input features by removing the mean and scaling to unit variance. Learned models often behave poorly if the input features are not normalized [17]. The high level system overview is summarized in Fig. 6.

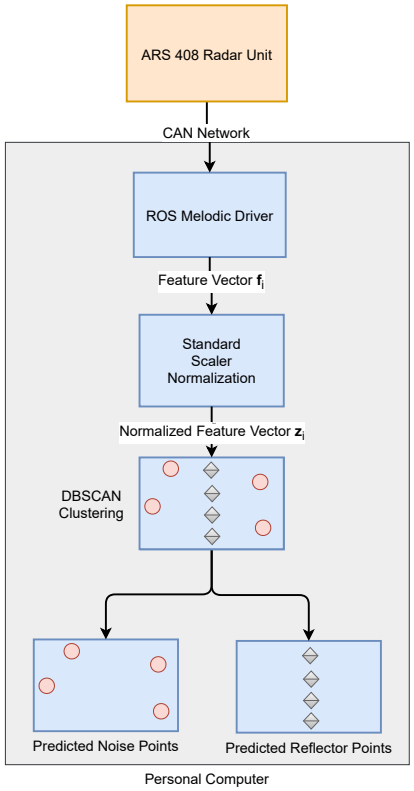

Fig. 6. High level overview of the radar data pipeline at a single time step. The raw radar data from the sensor is parsed to form feature vector $\mathbf{f}_i$ through a ROS Melodic driver. Feature vector $\mathbf{f}_i$ is normalized using a standard scaler and then inputted into the DBSCAN model. The final output classifies each data point as either noise or a retro-reflector.

B. DBSCAN Noise Filtering

This work aims to develop a solution to radar point cloud segmentation of noisy and sparse data. DBSCAN is designed to discover the clusters and noise in a spatial data set by using a fixed density threshold for the creation of a cluster. The computational complexity of DBSCAN is $O(n \log n)$ and the memory complexity is $O(n)$ [18]. Thus, the real-time implementation of this system incurs a low computational cost. The DBSCAN algorithm was implemented with *Python* 3.9.12 in a *Jupyter Notebook* using the *Sklearn* package.

Two parameters, ε and `MinPoints`, must be chosen carefully to ensure effective use of DBSCAN. The parameter ε defines the radii of a circle around any point to create an ε-neighbourhood [13]. `MinPoints` is the number of samples in a neighborhood for a point to be considered a core point [13]. This means that each observation is examined if there are at least `MinPoints` observations within a search radius ε. We ran the DBSCAN algorithm with different parameter values in order to discover the relationship between the parameters and the overall system accuracy. For this application, we tuned the parameters to promote high precision and overall accuracy. The DBSCAN algorithm yielded the highest accuracy when `MinPoints` was set to three (3) and ε was set to 1.35. As we increased the `MinPoints` and ε parameters, the system generated many false positive retro-reflector predictions (low precision), which is undesirable for

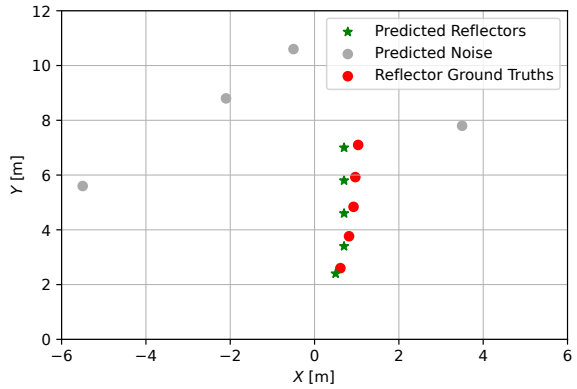

Fig. 7. Comparison of DBSCAN model predictions to the ground truth retro-reflector points. The model accurately clusters the noise points from the retro-reflector points with effective positional accuracy.

TABLE I
RADAR FEATURE VECTOR SAMPLED AT A SINGLE TIME STEP

Radar Return	X [m]	Y [m]	RCS [dBm^2]	Class
Detection 1	0.7	3.0	5.5	Retro-reflector
Detection 2	0.7	4.2	8.5	Retro-reflector
Detection 3	0.7	5.4	4.5	Retro-reflector
Detection 4	0.7	6.4	0.5	Retro-reflector
Detection 5	-5.1	5.2	24.0	Noise
Detection 6	-2.5	7.8	14.0	Noise
Detection 7	3.5	7.6	13.0	Noise
Detection 8	0.1	10.2	18.5	Noise

this system. As we decreased the parameters, the system had difficulty clustering the data and thus falsely predicted the retro-reflector points to be noise.

Although the overall radar point clouds were sparse, the retro-reflectors were installed with similar lateral positions in space and they possess low, but similar RCS values compared to noise points. Thus, we have intentionally created a dense region in the data space to complement the density-based clustering algorithm. DBSCAN is an effective clustering method when there are dense regions in a data space, i.e. a clear underlying pattern in the data [19]. The fundamental structure in the data can be seen in Table I, which presents a sample of feature vector $\mathbf{f}_i$ at approximately 2.08 s in trial run 1. We hypothesize that the high RCS value of noise points is due to the radar detecting and meshing together large indoor background objects. The DBSCAN model predictions were compared to the ground truth values from the Vicon system; a single time step, for example, is displayed in Fig. 7. This example shows an accurate classification (i.e., separation) of noise points and retro-reflector points. Based on the system's predictions, in this case, an AV could conceivably estimate a road's pathway using only one radar sensor.

V. RESULTS AND ANALYSIS

This section describes the results of two experimental trials. For each trial, exactly 63 sweeps of radar data were collected. At each time step, a maximum of eight objects were inferred by the radar, as set through the ARS 408 ROS configuration software for a total of 1008 inferred objects. We analyze the overall accuracy of the classification system and examine standard performance metrics.

The proposed clustering system predicts whether each radar return is a reflector point or a non-reflector point; i.e., what we call a noise point. One way to study the performance of this system is via a confusion matrix, as provided in Table II. Additionally, standard performance metrics such as accuracy, precision, recall, and $F1$ score are presented in Table III to describe the results from applying DBSCAN to 1008 radar returns. The confusion matrix was produced by manually comparing the DBSCAN output plots at each time step to the ground truth retro-reflector positions from the Vicon system. A radar return is considered to "match" a retro-reflector if its inferred position is within 0.5 m of a reflector's ground truth position. A 0.5 m error threshold was deemed reasonable given the radar's range resolution is 0.39 m [20]. A correctly predicted retro-reflector is considered a true positive. A correctly predicted noise point is considered a true negative. A falsely predicted retro-reflector is considered a false positive, and a falsely predicted noise point is considered a false negative.

We derived standard and widely used performance metrics from the confusion matrix, namely

$$\text{accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{5}$$

$$\text{precision} = \frac{TP}{TP + FP} \tag{6}$$

$$\text{recall} = \frac{TP}{TP + FN} \tag{7}$$

$$F1 = \frac{2 \times \text{precision} \times \text{recall}}{\text{precision} + \text{recall}} \tag{8}$$

where TP is true positive, TN is true negative, FP is false positive, and FN is false negative. The system accuracy was computed to be 84.13 %, which is an effective measurement of system performance because the data set is almost symmetric (i.e., the total number of reflector points (530) is nearly balanced with the total number of noise points (478)). The $F1$ score, which represents the weighted average of precision and recall, was 83.71 %. The system precision and recall were 90.93 % and 77.55 %, respectively. This indicates a conservative prediction approach with a low false positive rate. False positives have a high cost as they correspond to a false understanding of the road's boundaries which poses a safety risk to passengers. Two time steps that illustrate the impact of a FP and FN detection can be seen in Fig. 8 and Fig. 9, respectively. For the application of lane identification and vehicle localization, we would rather a system with high precision and few false positive predictions.

VI. CONCLUSIONS AND FUTURE WORK

Experimental analyses were conducted to determine the potential for using a millimetre-wave radar sensor for the localization and road following of an AV. This paper presented

TABLE II

CONFUSION MATRIX FOR THE CLASSIFICATION OF RETRO-REFLECTORS AND NOISE FROM RADAR RETURNS

		Prediction		
		Reflector	Noise	Total
Ground Truth	Reflector	411	119	530
	Noise	41	437	478
	Total	452	556	

TABLE III

EXPERIMENTAL PERFORMANCE METRICS OF THE SYSTEM

Test	**Accuracy (%)**	**Precision (%)**	**Recall (%)**	$F1$ **(%)**
1	84.33	93.52	75.66	83.64
2	83.93	88.56	79.47	83.77
Total	**84.13**	**90.93**	**77.55**	**83.71**

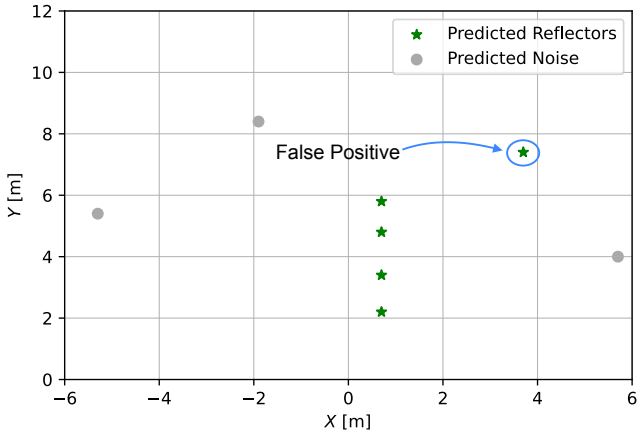

Fig. 8. DBSCAN filtering output where a FP is circled in blue. A FP detection incurs a high cost of danger as it results in an incorrect path estimation for the AV. Our system was designed to promote a low FP rate, and thus has a high precision value of 90.93 %.

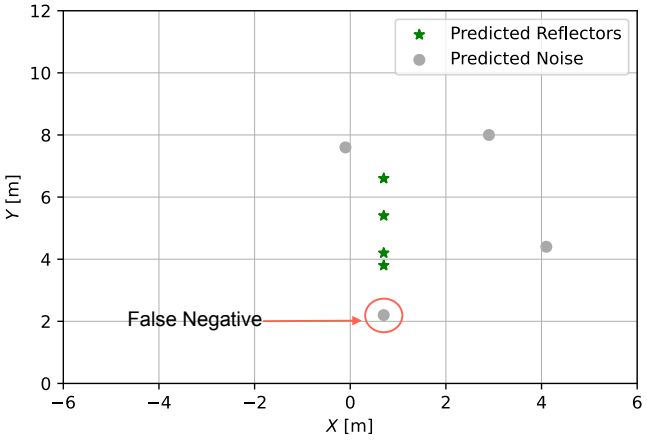

Fig. 9. DBSCAN filtering output where a FN is circled in red. A FN detection has a lower cost than a FP because the straight line path estimation is still maintained despite the missing retro-reflector detection.

a clustering solution for segmenting automotive radar data into two classes: retro-reflectors and noise. Retro-reflectors were installed as roadside landmarks to supplement radar sensors for road-path perception. We leveraged the geometrical pattern and RCS of retro-reflector pathways by using density-based clustering. The DBSCAN model was experimentally validated and achieved a classification accuracy of 84.13 % and $F1$ score of 83.71 %. This work contributes to the development of radar processing techniques, which could be used in ADAS lane keeping systems or AV localization systems. This work also studied the effectiveness of unsupervised learning techniques for segmenting radar data.

In the future, we would like to explore other popular clustering techniques, such as Gaussian mixture models and k-means. Additionally, we plan to collect a labeled radar data set and train supervised deep learning models that make use of temporal data for path prediction. Specifically, we hope to implement Reccurent Neural Networks (RNN) and Long Short Term Memory (LSTM) networks for retro-reflector path prediction. RNN and LSTM are widely used deep learning sequential models that have recently been applied to temporal sensor data [21]. The feature vector will grow to include the relative lateral and longitudinal velocities of inferred objects. We would like to implement these systems in real time and evaluate their performance using a closed loop path-following control system. We plan to compare our method with LiDAR-based retro-reflector detection. Finally, we would like to extend this work to curved roads in outdoor environments and conduct tests in poor weather conditions.

VII. ACKNOWLEDGMENTS

The authors would like to thank their colleagues at the Ingenuity Labs Research Institute, specifically Kai Wang and Franc Marrato, for providing support with experimentation, software debugging, and hardware set-up.

REFERENCES

[1] M. Shah, Z. Huang, A. Laddha, M. Langford, B. Barber, S. Zhang, C. Vallespi-Gonzalez, and R. Urtasun, "Liranet: End-to-end trajectory prediction using spatio-temporal radar fusion," in *CoRL*, pp. 31–48, Nov. 16-18 2020.

[2] R. Nabati and H. Qi, "CenterFusion: Center-based Radar and Camera Fusion for 3D Object Detection," *2021 IEEE Winter Conference on Applications of Computer Vision*, pp. 1526–1535, Jan. 5-9 2021.

[3] D. Adolfsson, M. Magnusson, A. Alhashimi, A. J. Lilienthal, and H. Andreasson, "Cfear radarodometry - conservative filtering for efficient and accurate radar odometry," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5462–5469, Sep. 27 - Oct. 1 2021.

[4] F. Jin, A. Sengupta, S. Cao, and Y.-J. Wu, "MmWave Radar Point Cloud Segmentation using GMM in Multimodal Traffic Monitoring," in *2020 IEEE International Radar Conference*, pp. 732–737, Apr. 2020.

[5] A. Danzer, T. Griebel, M. Bach, and K. Dietmayer, "2D Car Detection in Radar Data with PointNets," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pp. 61–66, Oct. 27-30 2019.

[6] A. Ouaknine, A. Newson, P. Pérez, F. Tupin, and J. Rebut, "Multi-View Radar Semantic Segmentation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 15671–15680, Oct. 11-17 2021.

[7] D. Barnes, M. Gadd, P. Murcutt, P. Newman, and I. Posner, "The oxford radar robotcar dataset: A radar extension to the oxford robotcar dataset," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6433–6438, May 31-Aug. 31 2020.

[8] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuscenes: A multimodal dataset for autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 14-19 2020.

[9] K. Burnett, A. P. Schoellig, and T. D. Barfoot, "Do we need to compensate for motion distortion and doppler effects in spinning radar navigation?," *IEEE Robotics and Automation Letters*, vol. 6, pp. 771–778, Apr. 2021.

[10] A. Greisman, K. Hashtrudi-Zaad, and J. A. Marshall, "Detection of conductive lane markers using mm wave fmcw automotive radar," in *2021 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, pp. 1–6, Sep. 23-25 2021.

[11] D. Brodeski, I. Bilik, and R. Giryes, "Deep radar detector," in *2019 IEEE Radar Conference (RadarConf)*, pp. 1–6, Apr. 22-26 2019.

[12] D. Barnes and I. Posner, "Under the radar: Learning to predict robust keypoints for odometry estimation and metric localisation in radar," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 9484–9490, May 31 - Aug. 31 2020.

[13] D. Kellner, J. Klappstein, and K. Dietmayer, "Grid-based dbscan for clustering extended objects in radar data," in *2012 IEEE Intelligent Vehicles Symposium*, pp. 365–370, June 3-7 2012.

[14] N. Scheiner, N. Appenrodt, J. Dickmann, and B. Sick, "A multi-stage clustering framework for automotive radar data," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pp. 2060–2067, Oct. 27-30 2019.

[15] M. Kalaagi and D. Seetharamdoo, "Radar cross section enhancement using metasurfaces for road safety applications," in *2022 16th European Conference on Antennas and Propagation (EuCAP)*, pp. 1–5, Mar. 27 - Apr. 1 2022.

[16] "Standard radar interface technical documentation," tech. rep., Continental Engineering Services GmbH, D-60489 Frankfurt Breitlacher Straße 94, 2019.

[17] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[18] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise," in *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, KDD'96, p. 226–231, AAAI Press, 1996.

[19] M. Ester, H.-P. Kriegel, and X. Xu, "A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise," in *Proceedings of the Thirteenth National Conference on Artificial Intelligence and Eighth Innovative Applications of Artificial Intelligence Conference (AAAI)*, Aug. 4-8 1996.

[20] A. Voigt, "Ars 408-21 long range radar sensor technical data," tech. rep., Continental Engineering Services GmbH, D-60489 Frankfurt Breitlacher Straße 94, 2020.

[21] K. Min, D. Kim, J. Park, and K. Huh, "Rnn-based path prediction of obstacle vehicles with deep ensemble," *IEEE Transactions on Vehicular Technology*, vol. 68, pp. 10252–10256, Oct 2019.