2024 IEEE International Conference on Robotics and Automation (ICRA 2024)
May 13-17, 2024. Yokohama, Japan

Radar-Only Odometry and Mapping for Autonomous Vehicles

Daniel Casado Herraez Matthias Zeller Le Chang Ignacio Vizzo Michael Heidingsfeld Cyrill Stachniss

***Abstract*— Odometry and mapping play a pivotal role in the navigation of autonomous vehicles. In this paper, we address the problem of pose estimation and map creation using only radar sensors. We focus on two odometry estimation approaches followed by a mapping step. The first one is a new point-to-point ICP approach that leverages the velocity information provided by 3D radar sensors. The second one is advantageous for 2D radars with a low number of samples, and particularly useful for scenarios where the sensor is being blocked by large dynamic obstacles. It exploits a constant velocity filter and the measured Doppler velocities to estimate the vehicle's ego-motion. We enrich this with a filtering step to improve the accuracy of the points in the resulting map. We put our work to the test using the View of Delft and NuScenes datasets, which involve 3D and 2D radar sensors. Our findings illustrate state-of-the-art performance of our odometry techniques in terms of accuracy when compared to existing alternatives. Moreover, we demonstrate that our map filtering methodology achieves higher similarity rates than the raw unfiltered map when benchmarked against a corresponding LiDAR map.**

I. Introduction

Sensor odometry and mapping are fundamental tasks for autonomous vehicles navigating unfamiliar terrain. This process is indispensable for two primary reasons. Firstly, it empowers these vehicles to estimate their own motion by leveraging data from their surroundings, also proving an orthogonal relative pose cue to GNSS systems. Secondly, it facilitates the creation of environment maps, a valuable feature for localization, path planning, and spatial scene understanding. Pose estimation and map creation are commonly achieved using cameras, which are sensitive to low lighting, or LiDARs, which are costly, affected by bad weather conditions, and often difficult to pack into an end-user vehicle.

In this paper, we investigate the problem of radar odometry and mapping relying solely on automotive radars, without the need for wheel odometry or additional sensor information. These sensors are already present in several consumer vehicles. They are compact, low cost, and robust to adverse weather. They also provide additional information like the velocity and radar cross-section (RCS) of the measurements. However, the noise and sparsity of their output scan poses a new set of questions, making radar odometry and mapping a challenging task.

D. Casado Herraez and M. Zeller are with CARIAD SE and with the University of Bonn, Germany. L. Chang is with CARIAD SE and with the University of Stuttgart, Germany. I. Vizzo is with Dexory, UK; during the work on the paper, he was with the University of Bonn. M. Heidingsfeld is with CARIAD SE, Germany. C. Stachniss is with the Center for Robotics, University of Bonn and with the Lamarr Institute for Machine Learning and Artificial Intelligence, Germany.

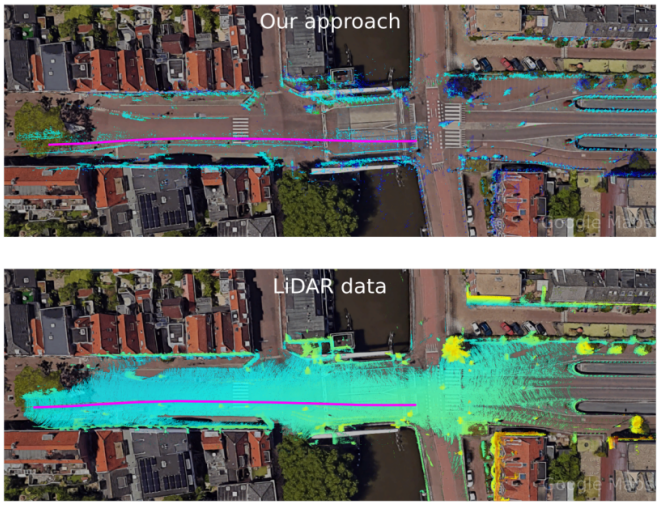

Fig. 1: (top) Odometry estimation and map generated using our radar point-to-point ICP approach for seq. 10 of the View of Delft dataset. (bottom) LiDAR data shown for comparison.

While some state-of-the-art methods use the Doppler velocity information to estimate the vehicle's ego-motion, others use positional information of the measurements [3] [21]. Prior research has also explored various avenues, leveraging both, signal processing and machine learning techniques. Some approaches have transformed radar scans to look similar to LiDAR point clouds [32] [48], while others have harnessed geometric models to identify surface reflections [26] [29] [39] [43]. Furthermore, robust methods that primarily rely on vehicle velocity [19] [27] have demonstrated success in rejecting dynamic object and noise outliers from single radar scans. However, these methods tend to face limitations when dealing with high outlier densities that block large parts of the radar's view. They also place little focus on the mapping task, and very few authors evaluate their methods on publicly available datasets.

The main contribution of this paper is a comprehensive exploration of odometry techniques plus mapping utilizing 3D and 2D radar sensors for autonomous vehicles. Our focus encompasses two ego-motion estimation approaches and a novel mapping step, collectively enhancing the capabilities of radar-driven navigation. The first method presents a novel point-to-point iterative closest point (ICP) technique specifically designed to harness the velocity information provided by radar sensors. The second method is tailored for low-speed scenarios where the radar is being covered by large moving obstacles such as trucks or buses, and is a

979-8-3503-8457-4/24/$31.00 ©2024 IEEE

configuration particularly advantageous for 2D radar sensors with a low number of samples. Our approach leverages a constant velocity filter in conjunction with Doppler velocity measurements to robustly estimate the ego-motion of the vehicle. We also give weight to the mapping stage and introduce a simple, yet effective, filtering step to improve the precision of the resulting map. Finally, we carry out an extensive evaluation on public real-world datasets.

In sum, we make three key claims: Our work (i) can achieve state-of-the-art results in 3D radar odometry using a point-to-point radar ICP approach, (ii) can accurately estimate odometry in 2D radars, particularly in scenarios where most of the radar is blocked by large dynamic objects, and (iii) provides a simple yet effective filtering step for radar map creation leading to a high similarity compared to its corresponding LiDAR map. These claims are backed up by the paper and our experimental evaluation.

II. Related Work

This section presents a comprehensive overview of the state-of-the-art approaches in the domain of sensor odometry and mapping. We divide it into three main categories: radar sensors and their use in vehicles, LiDAR scan matching, and radar odometry and mapping.

Radar sensors have emerged as essential components of autonomous driving, operating effectively under low lighting and adverse weather conditions. Instead of emitting and receiving light pulses like LiDAR, they use frequency modulated continuous wave (FMCW) signals to measure the range and velocity of the target. Overall, we can categorize radars into two groups based on their output, scanning radars and automotive radars. The former, output a 2D image where the value of each pixel represents the intensity of the measurement. Meanwhile, automotive radars output a point cloud with positional and velocity information. They are compact thus easy to pack inside a vehicle. They are the ones used in our work.

A big challenge for both is that the output is characterized as being sparse with a large number of noisy measurements, which can affect pose estimation, mapping, and other perception tasks. Some authors approach this by predicting the radar scan's occupancy within a grid [32] [48]. Others use machine learning [30] or geometric models [26] [29] [39] [43] to identify surface reflections. However, for odometry, RANSAC outlier rejection based on the vehicle's velocity has proven successful in rejecting dynamic objects and noise outliers from a single scan [27]. A shortcoming of this method is that it does not account for possible changes in velocity and can lead to failure when most of the radar measurements come from a moving object.

LiDAR scan matching solves the problem of estimating the transformation between two point clouds. Some popular approaches include LOAM [52], which minimizes the distance between corners and edges, and surfel-based matching [7]. However, the iterative closest point algorithm and its variations [15] [45] [51] [53] remain a highly effective solution to the problem. More specifically, point-to-point approaches like KISS-ICP [47] can lead to high accuracy results without added complexity. The problem in applying these methods directly to radar point clouds is the different nature of the points. While LiDAR point clouds are commonly dense (millions of points per scan) and accurate, radar point clouds are sparse (hundreds of points per scan) and comparably noisy. This can be problematic for point matching, resulting in point-to-point methods that consider only positional information being easily outperformed [21] [31].

Combined with the Doppler velocity information, we take advantage of the radar functionalities to improve the accuracy of our scan matching approach. The closest works in this area are related to FMCW LiDARs [22] [50], but the high density and low noise of the output point cloud means their point-to-plane technique does not directly work in the radar domain, which is sparser and with a higher amount of outliers. We extend their method to radar sensing, but rely on point-to-point, scan-to-map alignment, as well using velocity measurements for outlier rejection and better initialization.

Radar odometry and mapping define those methods that estimate the transformation between two point clouds and store the points into a map. These may not necessarily require point correspondence matching between the scans thanks to the radar's Doppler velocity information [6] [27] [28] [38] [49]. Many of the existing approaches are based on the scanning radar and are usually either feature-based [1] [5] [8] [9] [10] [12] [13] [25] [33], where keypoints are extracted and matched across the output intensity images, or signal processing based [6] [38] [49], where the car's ego-pose is estimated using correlation and the Fourier Mellin transform between frames. In the context of automotive radars, some take advantage of combining them with other sensors including LiDAR for operation or training [2] [16] [17], cameras [55], or IMUs to improve accuracy [4] [18] [36] [54]. Robust kernels [14] and probabilistic approaches like the normal distribution transform [40] [41] or the Gaussian mixture model [21] account for the uncertainty of a measurement as a way of dealing with the high amount of outliers. Others approach the problem based only on the Doppler velocity information [27] [28], and/or on the vehicle's motion kinematics [34] [42] [44]. We observe that focusing on the velocity can be useful especially when the measurements are planar and have a low number of points. Additionally, using previous pose information can help us reject outliers originating from large moving obstacles. Aldera et al. [3] also use positional information for estimating point matches between scans. In our approach, we demonstrate how combining both, velocity and positional information can lead to almost LiDAR level accuracy, especially regarding 3D radar point clouds. Furthermore, despite the existence of SLAM systems [46] operating on scanning radars [23] [9], are assisted by IMUs [54], or that make use of semantically labelled point clouds [24], little emphasis has been put on the structure of the resulting map. We present a filtering strategy that uses velocity and point matches across scans to eliminate of outliers, and compare the results against its relative LiDAR map.

III. Our Approach

This work aims to estimate the pose of the vehicle and store the sensor readings in a map of the environment. For this, we introduce two new odometry methods for 3D and 2D automotive radars. For the 3D case, we demonstrate how our scan-to-map point-to-point ICP strategy can achieve state-of-the-art accuracy using the positional and velocity information of the measurements. For the 2D case, we improve previous approaches by addressing scenarios where most of the radar's field of view is being covered by a moving obstacle. Additionally, we propose a filtering step to improve the consistency of the points in the resulting map.

A. Point-to-Point Radar ICP

Automotive radars produce sparse point clouds. Therefore, scan matching methods that involve finding point correspondences have been commonly outperformed by NDT [31] and other probabilistic approaches [21]. However, we demonstrate how Doppler velocity information provided by the radar can be used to augment the point-to-point ICP algorithm leading to high registration accuracy. We demonstrate how the method benefits from the higher amount of points and the additional axis from 3D automotive radars, still providing accurate results for 2D radar sensors.

The goal of point-to-point ICP is to obtain the transformation $\mathsf{T} \in SE(3)$ that minimizes the distance between a source point set $\mathcal{P} = \{p_1, p_2, ..., p_M\}$, and a target point set $\mathcal{Q} = \{q_1, q_2, ..., q_N\}$. In our method, we address the sparsity of radar point clouds by building $\mathcal{Q}$ as a submap aggregating the ten previous scans. Each iteration involves two steps.

We obtain the set of point correspondences $\mathcal{M} = \{m_1, m_2, ..., m_K\}$ based on the Euclidean distance between matches, represented as tuples $\boldsymbol{m}_k = (p_i, q_j)$.

This approach then minimizes the point position error E_t function between the matches, where subscript t denotes position. Error terms can be weighted adaptively using the Geman McClure kernel ρ_t as suggested by Vizzo et al. [47]:

$$E_t(\mathsf{T}) = \sum_{k=1}^{N} \rho_t(||\boldsymbol{q}_k - \mathsf{T}\boldsymbol{p}_k||). \tag{1}$$

We can derive the Jacobian of the point position error for a point k using its Lie algebra formulation where $^\wedge$ denotes the skew-symmetric operator:

$$\mathcal{J}_{t_k} = \begin{bmatrix} \mathsf{I}_{3\times3} & , & -(\mathsf{T}\boldsymbol{p}_k)^\wedge \end{bmatrix}. \tag{2}$$

Originally, Eq. (1) used on laser sensors only includes positional information. We extend it to radars taking inspiration from Hexsel et al. [22] by assuming rigid body geometry, between the vehicle, the sensor, and the point measurement, see Fig. 2. The expected velocity of the point ${}^V\boldsymbol{v}_{\mathrm{est}_k}$ is given by the projection of the estimated sensor velocity ${}^V\boldsymbol{v}_S$ in the vehicle frame, denoted by V, into the radial measurement direction ${}^V\boldsymbol{d}_k$ for a point k in the vehicle frame:

$${}^V\boldsymbol{v}_{\mathrm{est}_k} = -\,{}^V\boldsymbol{d}_k \cdot {}^V\boldsymbol{v}_S. \tag{3}$$

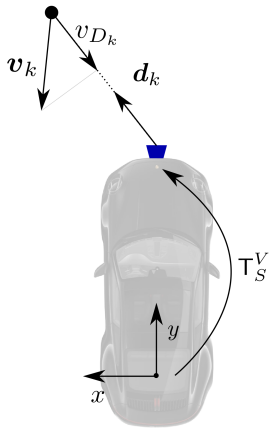

Fig. 2: Geometry of the measurements in the sensor frame S, where v_k is the velocity of the static target, $\boldsymbol{v}_{D_k}$ is the projection of $\boldsymbol{v}_k$ in the radial direction of the sensor $\boldsymbol{d}_k$, and T_S^V is the calibration matrix from vehicle to sensor frame.

The point velocity error E_v is the difference between the estimated velocity of the point from the currently computed update, and the measured Doppler velocity, where subscript v denotes velocity:

$$E_v(\mathsf{T}) = \sum_{k=1}^{N} \rho_v(||\boldsymbol{v}_{D_k} - \boldsymbol{v}_{\mathrm{est}_k}(\mathsf{T})||) \tag{4}$$

The Geman McClure kernel ρ_v is tuned empirically based on the distribution of the velocity residuals in a radar scan.

The Jacobian of the velocity error depends on the calibration matrix from the vehicle to the sensor $\mathsf{T}_S^V = (\mathsf{R}_S^V, \boldsymbol{t}_S^V)$, represented by a rotation matrix $\mathsf{R}_S^V \in SO(3)$ and a translation vector $\boldsymbol{t}_S^V \in \mathbb{R}^3$. It is given as:

$$\mathcal{J}_{v_k} = \begin{bmatrix} \frac{-\,{}^V\boldsymbol{d}_k}{\Delta t} & , & \frac{-\,{}^V\boldsymbol{d}_k \times \boldsymbol{t}_S^V}{\Delta t} \end{bmatrix}. \tag{5}$$

The optimization problem outputs the optimal transformation between the previous pose and the current one $\mathsf{T} \in SE(3)$. The final error function depends on the point position and velocity errors, Eq. (6). Parameter $\gamma \in [0, 1]$ allows us to weigh each error type individually. A value of $\gamma = 0$ would result in simple point-to-point ICP without considering Doppler velocities.

$$\mathsf{T}^* = \underset{\mathsf{T}}{\operatorname{argmin}} (1-\gamma) E_t(\mathsf{T}) + \gamma E_v(\mathsf{T}). \tag{6}$$

Note that as the first estimate for the optimization, we could assume a constant velocity from the previous pose as done by DICP [22] and KISS-ICP [47]. One of the advantages of radar sensors is, however, that they provide the velocity measurement for each observed target. With this information, we can make a first estimate of the ego-velocity directly from the current scan by following a least squares approach introduced by Kellner et al. [27]. This also enables us to remove with RANSAC dynamic point outliers prior to ICP optimization if they don't match the velocity estimate.

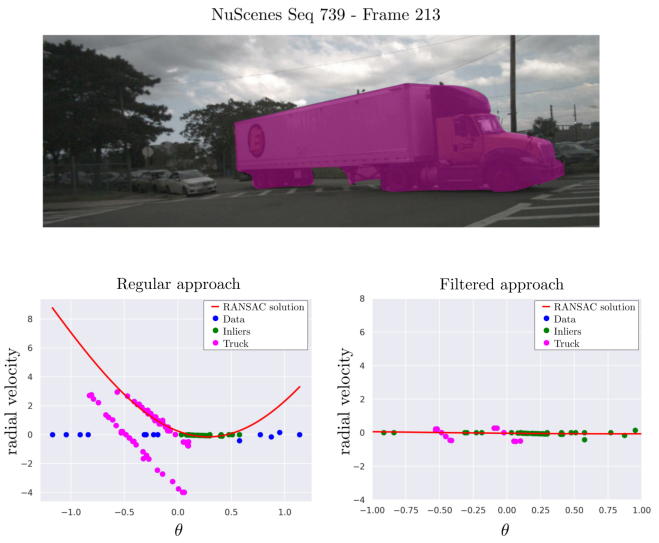

Fig. 3: Scenario where a truck is moving in front of the sensor. Kellner et al. [27](left) fail by considering the truck's rear wheels as inliers. Our velocity filter (right) Sec. III-B eliminates the points that don't match the vehicle's motion leading to an accurate estimation.

The final Jacobian for a point k is obtained by combining the point-to-point Jacobian with the velocity Jacobian:

$$\mathcal{J}_k = \begin{bmatrix} \mathsf{I}_{3\times 3} & -(\mathsf{T}\boldsymbol{p}_k)^{\wedge} \\ \frac{-{}^{V}\boldsymbol{d}_k}{\Delta t} & \frac{-{}^{V}\boldsymbol{d}_k \times \boldsymbol{t}_S^V}{\Delta t} \end{bmatrix}. \tag{7}$$

B. Single-Scan Odometry and Velocity Filter

One interesting feature to exploit when using radar sensors is that they provide a Doppler measurement of the point's velocity. Under the assumptions of a mostly static environment and an Ackerman vehicle with no lateral displacement, Kellner et al. [27] propose a least squares solution to estimate the vehicle's odometry using RANSAC for dynamic outlier rejection. It requires a minimum of two velocity measurements from stationary targets to output a solution, making it appropriate for radars with a low number of points like the 2D ARS-408 used in NuScenes (125 points/scan).

However, solely relying on the current scan makes the system vulnerable to changes in the environment, such as cases where most of the points in the sensor scan are outliers coming from large moving objects (trucks, buses, etc.). In situations like Fig. 3, the RANSAC solution may fail to estimate the vehicle's velocity. This can occur when the vehicle approaches an intersection and a truck comes by, or when the vehicle is parking with other moving cars.

Our method addresses this problem by introducing an initial filtering step previous to RANSAC that adopts a constant velocity model, inspired by ego-pose prediction in KISS-ICP [47]. Our filter only keeps feasible samples according to the car's previous motion. For this, we approximate the translation and angular velocities in the current time step by using pose estimates from the two previous scans, T_{i-1} and T_{i-2}. Then, we can derive the corresponding car velocities in the vehicle frame as:

$${}^{V}\boldsymbol{v}_V = [\boldsymbol{v}_t, \boldsymbol{v}_\theta]^T = \frac{\mathrm{Log}(\mathsf{T}_{i-2}^{\top}\mathsf{T}_{i-1})}{\Delta t}, \tag{8}$$

where Δt is the difference in timestamps between the previous and the current scan, and Log: $SE(3) \to \mathbb{R}^6$ extracts the linear translation for each axis, as well as the axis-angle representation of the rotation.

Similar to Eq. (3), the predicted velocity in the radial direction for each point k in the current scan can be estimated, computing the projection of the sensor velocity in the vehicle frame to the direction of the measurement:

$${}^{V}v_{\mathrm{pred}_k} = -{}^{V}\boldsymbol{d}_k \cdot (\boldsymbol{v}_t + \boldsymbol{v}_\theta \times \boldsymbol{t}_S^V). \tag{9}$$

Then, it is possible to filter each point with a measured Doppler velocity v_{D_k} that doesn't match the predicted point velocity within a margin ε, adjusted empirically with the radar velocity resolution:

$$v_{D_k} - {}^{V}v_{\mathrm{pred}_k} < \varepsilon. \tag{10}$$

This simple thresholding acts as a first-step filtering executed before RANSAC, helping to reduce the risk of failure when rejecting the RANSAC outliers. As shown in Tab. I, it can also be applied to 3D radar point clouds. The final velocity of the vehicle can be estimated using the least squares approach from Kellner et al. [27].

C. Radar Mapping

In the previous sections, we presented our methods for accurately estimating the pose of the vehicle. Now we shift the focus toward mapping. It involves storing the points to obtain a useful representation of the environment that can be used for localization and place recognition. As radar point clouds are noisy, it is important that only valuable points are stored. We propose a fairly simple yet effective pre-filtering strategy to obtain a map with consistent measurements.

Dynamic points that don't match the vehicle's speed are removed with RANSAC after estimating odometry. We treat these points as moving objects or noise and discard them from our global map.

Static point outliers pose a challenge due to the sparsity of radar point clouds. We approach this using the principle of scan-to-submap matching [35] [54] and constructing two different maps, a local map, and a global map.

We create the local map by aggregating multiple previous scans. Then, we transform the current radar scan into local map coordinates using our last odometry estimate. For each point, we measure the distance to its closest correspondence, similar to the matching step in Sec. III-A. If the distance is greater than a threshold δ, we classify the point as an outlier. Otherwise, we consider it an inlier, see Fig. 4. After determining the point inliers of the current scan using the local map, we add them to the global map. The global map stores all the inliers from all scans, and it is the map we will evaluate.

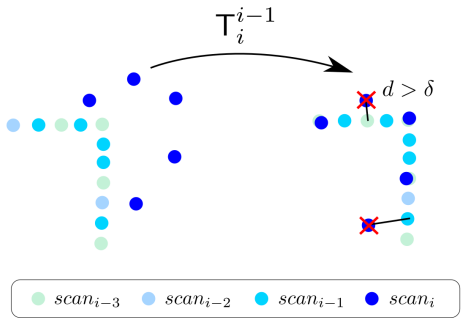

Fig. 4: The local map is constructed using the previous K=3 scans. Comparing the distance d of each point in the current scan to its closest neighbor in the local submap we determine if a point is an inlier or an outlier.

TABLE I: Evaluation on the View of Delft dataset using 3D radars.

Method	Mean RPE/RMSE/KITTI [m/m/%]	Mean RPE/RMSE/KITTI [°]	σ [m]	σ [°]
KISS-ICP *on LiDAR data*	**0.048 / 0.060 / 4.02**	**0.051 / 0.080 / 0.056**	**0.036**	**0.061**
Point-to-Point ICP	0.067 / 0.088 / 10.9	0.064 / 0.100 / 0.143	0.055	0.073
4DRVO-Net	0.080 / - / -	0.070 / - / -	-	-
DICP	0.079 / 0.101 / 5.35	0.073 / 0.168 / 0.131	0.061	0.145
Kellner et al.	0.078 / 0.116 / **4.36**	0.097 / 0.171 / 0.173	0.059	0.120
KISS-ICP on radar data	0.061 / 0.073 / 4.66	0.074 / 0.114 / **0.098**	0.040	0.080
Velocity filtered (ours)	0.078 / 0.115 / **4.36**	0.097 / 0.170 / 0.155	0.059	0.116
Radar ICP (ours)	**0.049 / 0.062** / 4.41	**0.059 / 0.089** / 0.100	**0.036**	**0.066**

IV. Experimental Evaluation

This work provides two radar-only odometry methods, accompanied by a comprehensive mapping strategy tailored for radar sensors. While our odometry approaches are characterized by effective strategies for high outlier scenarios and sparse point clouds, our mapping manages to record consistent and valuable points in the environment.

We present our experiments to show the capabilities of our work. The results of our evaluation support our key claims that our work (i) can achieve state-of-the-art results on 3D radar odometry using a point-to-point radar ICP approach, (ii) can accurately estimate odometry in 2D radars, particularly in scenarios where most of the radar is blocked by large dynamic objects, and (iii) provides a simple yet effective filtering step for radar map creation leading to a high similarity compared to its corresponding LiDAR map.

A. Experimental Setup

On the evaluation of our method, we run experiments on 3D and 2D radar datasets in real-world driving scenarios, View of Delft [37] (VoD) and NuScenes [11] respectively. We first evaluate our work on odometry comparing it to the GPS ground truth and to other radar and LiDAR methods. We denote our ICP approach from Sec. III-A as "Radar ICP" and our constant velocity filter method from Sec. III-B as "Velocity Filter". We show the results for generic scenarios as well as for some edge cases that particularly benefit from our work. Then, we evaluate the filtered radar map with respect to the LiDAR map qualitatively and quantitatively using the symmetric and asymmetric chamfer distance similarity metrics.

TABLE II: Evaluation on the NuScenes dataset using 2D radars

Method	Mean RPE/RMSE [m]	Mean RPE/RMSE [°]	σ [m]	σ [°]
KISS-ICP *on LiDAR data*	0.013 / 0.016	**0.035 / 0.051**	0.012	**0.051**
Point-to-Point ICP	0.354 / 0.641	0.468 / 0.875	2.246	4.871
KISS-ICP on radar data	0.133 / 0.162	0.149 / 0.196	0.104	0.140
CREME	- / 0.012	- / **0.100**	-	-
Kellner et al.	0.010 / **0.011**	0.108 / 0.174	0.009	0.206
Velocity filtered (ours)	**0.009 / 0.011**	0.088 / 0.117	**0.008**	0.110
Radar ICP (ours)	0.012 / 0.015	**0.085** / 0.126	0.009	**0.089**

B. Performance on View of Delft 3D Radar Dataset

The first experiment evaluates the accuracy of our 3D radar odometry estimation pipeline in comparison to other methods. This evaluation supports our claim that leveraging Doppler velocity information with point-to-point ICP can yield high accuracy, obtaining state-of-the-art results in 3D radar odometry. We evaluate both of our methods in the 3D radar dataset [37] and compare them with other approaches in Tab. I including raw ICP, KISS-ICP [47] on radar and LiDAR data, Kellner et al.'s ego-motion estimation [27], point-to-plane DICP [22], and the deep learning-based 4DRVO-Net [55] that combines camera and radar data. For comparison with 4DRVO-Net, we use the results from their paper including only the evaluation sequences (3, 4, 9, 17, 19, 22, 24). However, it gives us a good reference of its performance. In our evaluation, we focus on planar movement. We use the evo library to compute the RPE root-mean-squared error (RMSE) and the relative position error (RPE) for translation and rotation, its standard deviation, and the KITTI RPE (KITTI) [20] over 1, 10, 20, and 30 m. Moreover, $\gamma = 0.1$ is set for all sequences. Our radar ICP approach gives the best scan-to-scan RPE and RMSE results compared to other radar methods, close to LiDAR accuracy. Note that KISS-ICP is a state-of-the-art method for LiDAR scan registration, based on classical point-to-point matching. The good results of KISS-ICP on radar data show the importance of point-to-point positional information, improved in our technique that exploits the Doppler velocity measurements.

C. Performance in Presence of Large Dynamic Objects

The second experiment studies the performance of our method in 2D radar sensors when big moving objects are within the radar's field of view. While View of Delft [37] is good for evaluating odometry and mapping, the scenarios are limited to a specific area in a city, and most outliers belong to small objects easily removable by RANSAC. However, NuScenes [11] has a wider variety of cases, some of which include trucks and buses almost blocking the radar's view, see Fig. 3. During evaluation, we set the value of ε to 0.5 for all scenarios. Our results from Tab. III show scenes where a big truck or bus is directly turning in front of the estimating vehicle. We also evaluate our method on the full dataset in Tab. II comparing it against ICP, KISS-ICP [47] on radar and LiDAR, Kellner et al.'s ego-motion estimation [27], and CREME [21]. Our 2D odometry method improves its baseline by Kellner et al. [27] in all scenarios.

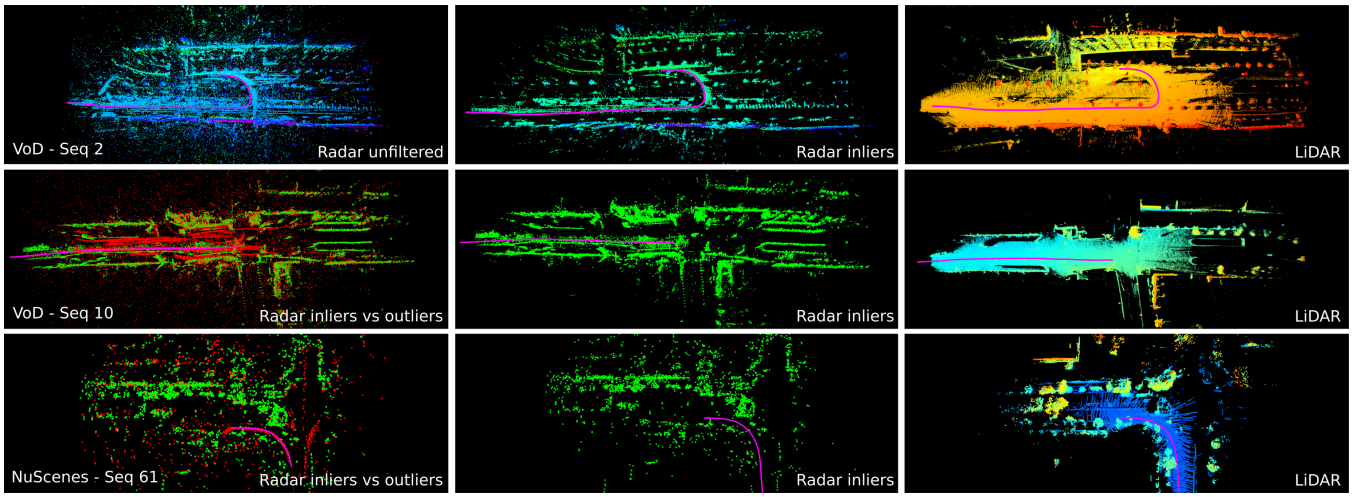

Fig. 5: (1^{st} row) Comparison between the map created using unfiltered radar scans, the map created using only the inliers from our map filtering strategy, and the LiDAR map. (2^{nd}, 3^{rd} rows) Display of the map inliers (green) and outliers (red), the resulting map constructed only from inliers, and the LiDAR map.

TABLE III: Evaluation of our 2D odometry estimation on NuScenes sequences with the presence of big moving objects.

Method	**0224**		**0242**		**0587**		**0739**		**0749**		**0875**		**0998**			
	RPE [m]	RPE [°]	RPE [m]	RPE [°]	RPE [m]	RPE [°]	RPE [m]	RPE [°]	RPE [m]	RPE [°]	RPE [m]	RPE [°]	RPE [m]	RPE [°]	Mean RPE [m]	Mean RPE [°]
Kellner et al.	0.0075	0.3088	0.0088	0.4698	0.0031	0.1087	0.0066	0.1832	0.0060	0.4157	0.0071	0.2390	0.0120	0.3103	0.0071	0.2859
CREME	0.0044	0.0812	0.0051	0.0867	0.0027	0.0675	0.0056	0.0714	0.0030	0.0630	0.0032	0.0528	0.0066	0.0800	0.0044	0.0718
Velocity filtered (ours)	**0.0036**	**0.0435**	**0.0031**	**0.0672**	**0.0020**	**0.0294**	**0.0052**	**0.0514**	**0.0022**	**0.0359**	**0.0024**	**0.0434**	**0.0060**	**0.0653**	**0.0034**	**0.0472**
Radar ICP (ours)	0.0164	0.3288	0.0224	0.3461	0.0104	0.1494	0.0106	0.1161	0.0178	0.0617	0.0123	0.204	0.0328	0.2621	0.0175	0.2097

D. Mapping Results

The last experiment evaluates our mapping stage. We show quantitatively and qualitatively that our filtered map obtains a high degree of similarity to the LiDAR map.

For the evaluation procedure, we construct the reference map by removing dynamic objects from the LiDAR point clouds based on ground truth annotations. Then, we transform the LiDAR scans to the global poses estimated by our odometry. As we are only using the front-facing radars in both datasets, we only consider LiDAR points inside the camera FOV. Moreover, in NuScenes [11], the radar range, 250 m, is much larger than the LiDAR's, 100 m, thus we only consider radar points within the LiDAR range for evaluation.

We use the symmetric chamfer distance SCD($\mathcal{R} \leftrightarrows \mathcal{L}$) as a measure of similarity between the radar map $\mathcal{R}$ and the LiDAR map $\mathcal{L}$. A smaller distance indicates better results, meaning that two point clouds are similar to each other, while a larger one indicates a big difference between them. Additionally, in order to reduce the effect of 2D radar noise points being matched with the LiDAR ground points in NuScenes, we also show the asymmetric chamfer distance ACD($\mathcal{R} \rightarrow \mathcal{L}$) measuring only the similarity from the radar to the LiDAR map. Results from Tab. IV(left) and Fig. 5 show that our approach leads to an improvement in SCD and ACD, and while it removes a high amount of noise and dynamic object outliers, it preserves the structures in the environment. This means the filtered map has a remarkably higher similarity than the unfiltered map and the map filtered using random point sampling. We also show the results on View of Delft for various thresholds δ in Tab. IV(right).

TABLE IV: (left) Similarity metrics between reference LiDAR and radar map. (right) Experiments for different distance thresholds δ. Given in m^2.

Method	VoD		NuScenes	
	SCD	ACD	SCD	ACD
No filter	10.98	1.20	44.22	2.44
Random	10.91	1.20	45.36	2.44
Ours	**3.61**	**0.66**	**40.47**	**1.92**

δ	SCD	ASC
0.5	2.52	0.52
1.0	2.81	0.61
2.0	3.13	0.69

V. Conclusion

In this paper, we presented a comprehensive exploration of odometry methods plus mapping for automotive 3D and 2D radar sensors. Our work encompasses two novel ego-motion estimation approaches and a mapping step, collectively enhancing radar-driven navigation capabilities. Our 3D radar method uses classical point-to-point ICP leveraging the Doppler velocity measurements provided by the radar. Our 2D radar approach exploits a constant velocity filter for scenarios where most of the sensor's field of view is covered by large dynamic objects. Moreover, our filtering steps also lead to a highly outlier-free map with high similarity to LiDAR. This allows us to successfully estimate the vehicle's ego-motion and construct an accurate map of the environment. We implemented and evaluated our approach on different datasets and provided comparisons to other existing techniques, supporting all claims made in this paper. The experiments suggest that our approach achieves state-of-the-art odometry estimates and demonstrates how simple filtering can lead to improvements in mapping accuracy.

REFERENCES

[1] D. Adolfsson, M. Magnusson, A. Alhashimi, A.J. Lilienthal, and H. Andreasson. Lidar-level localization with radar? the cfear approach to accurate, fast, and robust large-scale radar odometry in diverse environments. *IEEE Trans. on Robotics (TRO)*, 39(2):1476–1495, 2022.

[2] R. Aldera, D. De Martini, M. Gadd, and P. Newman. Fast radar motion estimation with a learnt focus of attention using weak supervision. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2019.

[3] R. Aldera, M. Gadd, D. De Martini, and P. Newman. What goes around: Leveraging a constant-curvature motion constraint in radar odometry. *IEEE Robotics and Automation Letters (RA-L)*, 7(3):7865–7872, 2022.

[4] Y. Almalioglu, M. Turan, C.X. Lu, N. Trigoni, and A. Markham. Milli-rio: Ego-motion estimation with low-cost millimetre-wave radar. *IEEE Sensors*, 21(3):3314–3323, 2020.

[5] D. Barnes and I. Posner. Under the radar: Learning to predict robust keypoints for odometry estimation and metric localisation in radar. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2020.

[6] D. Barnes, R. Weston, and I. Posner. Masking by moving: Learning distraction-free radar odometry from pose information. *arXiv preprint arXiv:1909.03752*, 2019.

[7] J. Behley and C. Stachniss. Efficient Surfel-Based SLAM using 3D Laser Range Data in Urban Environments. In *Proc. of Robotics: Science and Systems (RSS)*, 2018.

[8] 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 (RA-L)*, 6(2):771–778, 2021.

[9] K. Burnett, Y. Wu, D.J. Yoon, A.P. Schoellig, and T.D. Barfoot. Are we ready for radar to replace lidar in all-weather mapping and localization? *IEEE Robotics and Automation Letters (RA-L)*, 7(4):10328–10335, 2022.

[10] K. Burnett, D.J. Yoon, A.P. Schoellig, and T.D. Barfoot. Radar odometry combining probabilistic estimation and unsupervised feature learning. *arXiv preprint arXiv:2105.14152*, 2021.

[11] H. Caesar, V. Bankiti, A. Lang, S. Vora, V. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom. nuScenes: A Multimodal Dataset for Autonomous Driving. In *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020.

[12] S. Cen and P. Newman. Precise ego-motion estimation with millimeter-wave radar under diverse and challenging conditions. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2018.

[13] S. Cen and P. Newman. Radar-only ego-motion estimation in difficult settings via graph matching. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2019.

[14] N. Chebrolu, T. Läbe, O. Vysotska, J. Behley, and C. Stachniss. Adaptive Robust Kernels for Non-Linear Least Squares Problems. *IEEE Robotics and Automation Letters (RA-L)*, 6:2240–2247, 2021.

[15] P. Dellenbach, J. Deschaud, B. Jacquet, and F. Goulette. CT-ICP Real-Time Elastic LiDAR Odometry with Loop Closure. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2022.

[16] F. Ding, A. Palffy, D.M. Gavrila, and C.X. Lu. Hidden gems: 4d radar scene flow learning using cross-modal supervision. In *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2023.

[17] F. Ding, Z. Pan, Y. Deng, J. Deng, and C.X. Lu. Self-supervised scene flow estimation with 4-d automotive radar. *IEEE Robotics and Automation Letters (RA-L)*, 7(3):8233–8240, 2022.

[18] C. Doer and G.F. Trommer. An ekf based approach to radar inertial odometry. In *Proc. of the IEEE Intl. Conf. on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, 2020.

[19] P. Gao, S. Zhang, W. Wang, and C.X. Lu. Dc-loc: Accurate automotive radar based metric localization with explicit doppler compensation. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2022.

[20] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun. Vision meets Robotics: The KITTI Dataset. *Intl. Journal of Robotics Research (IJRR)*, 32(11), 2013.

[21] K. Haggag, S. Lange, T. Pfeifer, and P. Protzel. A credible and robust approach to ego-motion estimation using an automotive radar. *IEEE Robotics and Automation Letters (RA-L)*, 7(3):6020–6027, 2022.

[22] B. Hexsel, H. Vhavle, and Y. Chen. Dicp: Doppler iterative closest point algorithm. *arXiv preprint arXiv:2201.11944*, 2022.

[23] Z. Hong, Y. Petillot, A. Wallace, and S. Wang. Radar slam: A robust slam system for all weather conditions. *arXiv preprint arXiv:2104.05347*, 2021.

[24] S.T. Isele, F. Haas-Fickinger, and J.M. Zöllner. Seraloc: Slam on semantically annotated radar point-clouds. In *Proc. of the IEEE Intl. Conf. on Intelligent Transportation Systems (ITSC)*, 2021.

[25] E. Jose and M.D. Adams. Relative radar cross section based feature identification with millimeter wave radar for outdoor slam. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2004.

[26] A. Kamann, P. Held, F. Perras, P. Zaumseil, T. Brandmeier, and U.T. Schwarz. Automotive radar multipath propagation in uncertain environments. In *Proc. of the IEEE Intl. Conf. on Intelligent Transportation Systems (ITSC)*, 2018.

[27] D. Kellner, M. Barjenbruch, J. Klappstein, J. Dickmann, and K. Dietmayer. Instantaneous ego-motion estimation using doppler radar. In *Proc. of the IEEE Intl. Conf. on Intelligent Transportation Systems (ITSC)*, 2013.

[28] D. Kellner, M. Barjenbruch, J. Klappstein, J. Dickmann, and K. Dietmayer. Instantaneous ego-motion estimation using multiple doppler radars. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2014.

[29] J. Kopp, D. Kellner, A. Piroli, and K. Dietmayer. Fast rule-based clutter detection in automotive radar data. In *Proc. of the IEEE Intl. Conf. on Intelligent Transportation Systems (ITSC)*, 2021.

[30] F. Kraus, N. Scheiner, W. Ritter, and K. Dietmayer. Using machine learning to detect ghost images in automotive radar. In *Proc. of the IEEE Intl. Conf. on Intelligent Transportation Systems (ITSC)*, 2020.

[31] P.C. Kung, C.C. Wang, and W.C. Lin. A normal distribution transform-based radar odometry designed for scanning and automotive radars. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2021.

[32] P.C. Kung, C.C. Wang, and W.C. Lin. Radar occupancy prediction with lidar supervision while preserving long-range sensing and penetrating capabilities. *IEEE Robotics and Automation Letters (RA-L)*, 7(2):2637–2643, 2022.

[33] H. Lim, K. Han, G. Shin, G. Kim, S. Hong, and H. Myung. Orora: Outlier-robust radar odometry. *arXiv preprint arXiv:2303.01876*, 2023.

[34] S. Lupfer, M. Rapp, K. Dietmayer, P. Brosseit, J. Lombacher, M. Hahn, and J. Dickmann. Increasing fastslam accuracy for radar data by integrating the doppler information. In *Proc. of the IEEE MTT-S Intl. Conf. on Microwaves for Intelligent Mobility (ICMIM)*, 2017.

[35] R.A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A.J. Davison, P. Kohli, J. Shotton, S. Hodges, and A. Fitzgibbon. KinectFusion: Real-Time Dense Surface Mapping and Tracking. In *Proc. of the Intl. Symposium on Mixed and Augmented Reality (ISMAR)*, 2011.

[36] Y.Z. Ng, B. Choi, R. Tan, and L. Heng. Continuous-time radar-inertial odometry for automotive radars. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2021.

[37] A. Palffy, E. Pool, S. Baratam, J.F. Kooij, and D.M. Gavrila. Multi-class road user detection with 3+ 1d radar in the view-of-delft dataset. *IEEE Robotics and Automation Letters (RA-L)*, 7(2):4961–4968, 2022.

[38] Y.S. Park, Y.S. Shin, and A. Kim. Pharao: Direct radar odometry using phase correlation. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2020.

[39] R. Prophet, J. Martinez, J.C.F. Michel, R. Ebelt, I. Weber, and M. Vossiek. Instantaneous ghost detection identification in automotive scenarios. In *Proc. of the IEEE Radar Conf. (RadarConf)*, 2019.

[40] M. Rapp, M. Barjenbruch, K. Dietmayer, M. Hahn, and J. Dickmann. A fast probabilistic ego-motion estimation framework for radar. In *Proc. of the Europ. Conf. on Mobile Robotics (ECMR)*, 2015.

[41] M. Rapp, M. Barjenbruch, M. Hahn, J. Dickmann, and K. Dietmayer. Probabilistic ego-motion estimation using multiple automotive radar sensors. *Journal on Robotics and Autonomous Systems (RAS)*, 89:136–146, 2017.

[42] K. Retan, F. Loshaj, and M. Heizmann. Radar odometry on $se(3)$ with constant velocity motion prior. *IEEE Robotics and Automation Letters (RA-L)*, 6(4):6386–6393, 2021.

[43] F. Roos, M. Sadeghi, J. Bechter, N. Appenrodt, J. Dickmann, and C. Waldschmidt. Ghost target identification by analysis of the doppler distribution in automotive scenarios. In *Proc. of the International Radar Symposium (IRS)*, 2017.

[44] J. Schlichenmaier, L. Yan, M. Stolz, and C. Waldschmidt. Instantaneous actual motion estimation with a single high-resolution radar

sensor. In *Proc. of the IEEE MTT-S Intl. Conf. on Microwaves for Intelligent Mobility (ICMIM)*, 2018.

[45] A. Segal, D. Haehnel, and S. Thrun. Generalized-ICP. In *Proc. of Robotics: Science and Systems (RSS)*, 2009.

[46] C. Stachniss, J. Leonard, and S. Thrun. *Springer Handbook of Robotics, 2nd edition*, chapter Chapt. 46: Simultaneous Localization and Mapping. Springer Verlag, 2016.

[47] I. Vizzo, T. Guadagnino, B. Mersch, L. Wiesmann, J. Behley, and C. Stachniss. KISS-ICP: In Defense of Point-to-Point ICP – Simple, Accurate, and Robust Registration If Done the Right Way. *IEEE Robotics and Automation Letters (RA-L)*, 8(2):1–8, 2023.

[48] R. Weston, S. Cen, P. Newman, and I. Posner. Probably unknown: Deep inverse sensor modelling radar. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2019.

[49] R. Weston, M. Gadd, D. De Martini, P. Newman, and I. Posner. Fastmbym: Leveraging translational invariance of the fourier transform for efficient and accurate radar odometry. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2022.

[50] Y. Wu, D.J. Yoon, K. Burnett, S. Kammel, Y. Chen, H. Vhavle, and T.D. Barfoot. Picking up speed: Continuous-time lidar-only odometry using doppler velocity measurements. *IEEE Robotics and Automation Letters (RA-L)*, 8(1):264–271, 2022.

[51] J. Yang, H. Li, and Y. Jia. Go-icp: Solving 3d registration efficiently and globally optimally. In *Proc. of the IEEE/CVF Intl. Conf. on Computer Vision (ICCV)*, 2013.

[52] J. Zhang and S. Singh. LOAM: Lidar Odometry and Mapping in Real-time. In *Proc. of Robotics: Science and Systems (RSS)*, 2014.

[53] J. Zhang, Y. Yao, and B. Deng. Fast and robust iterative closest point. *IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI)*, 44(7):3450–3466, 2021.

[54] Y. Zhuang, B. Wang, J. Huai, and M. Li. 4d iriom: 4d imaging radar inertial odometry and mapping. *IEEE Robotics and Automation Letters (RA-L)*, 8(6):3246–3253, 2023.

[55] G. Zhuo, S. Lu, H. Zhou, L. Zheng, and L. Xiong. 4drvo-net: Deep 4d radar-visual odometry using multi-modal and multi-scale adaptive fusion. *arXiv preprint arXiv:2308.06573*, 2023.