

RaI-SLAM: Radar-Inertial SLAM for Autonomous Vehicles

Daniel Casado Herraiez

Matthias Zeller

Dong Wang

Jens Behley

Michael Heidingsfeld

Cyrill Stachniss

Abstract—Simultaneous localization and mapping are essential components for the operation of autonomous vehicles in unknown environments. While localization focuses on estimating the vehicle’s pose, mapping captures the surrounding environment to enhance future localization and decision-making. Localization is commonly achieved using external GNSS systems combined with inertial measurement units, LiDARs, and/or cameras. Automotive radars offer an attractive onboard sensing alternative due to their robustness to adverse weather and low lighting conditions, compactness, affordability, and widespread integration into consumer vehicles. However, they output comparably sparse and noisy point clouds that are challenging for pose estimation, easily leading to noisy trajectory estimates. We propose a modular approach that performs radar-inertial SLAM by fully leveraging the characteristics of automotive consumer-vehicle radar sensors. Our system achieves smooth and accurate onboard simultaneous localization and mapping by combining automotive radars with an IMU and exploiting the additional velocity and radar cross-section information provided by radar sensors, without relying on GNSS data. Specifically, radar scan-matching and IMU measurements are first incorporated into a local pose graph for odometry estimation. We then correct the accumulated drift through a global pose graph backend that optimizes detected loop closures. Contrary to existing radar SLAM methods, our graph-based approach is divided into distinct submodules and all components are designed specifically to exploit the characteristics of automotive radar sensors for scan matching and loop closure detection, leading to enhanced system performance. Our method achieves state-of-the-art accuracy on public autonomous driving data.

Index Terms—Odometry, Mapping, Localization, SLAM, Autonomous Vehicle Navigation

I. INTRODUCTION

SIMULTANEOUS localization and mapping (SLAM) plays a crucial role in the operation of autonomous vehicles within unknown environments. Estimating the pose of an autonomous car with respect to its surroundings is necessary for path planning, and collecting the surrounding information




Fig. 1: Qualitative results of our radar-inertial SLAM framework on the SNAIL-Radar dataset [21]. (left) Radar point cloud map (purple) projected into a Google Maps satellite view of the location of the dataset. (right) Different sections of the created map are shown with corresponding colored boxes in the left image with a visualization of the estimated trajectory (blue).

helps to localize and make more informed decisions when a place is being revisited. Pose estimation is commonly carried out using GNSS sensors. However, GNSS requires satellite visibility, making it unreliable in GNSS-denied areas like parking lots and tunnels. In such situations, autonomous vehicles must rely on other onboard sensing, typically LiDARs and/or cameras, to estimate their pose over time [41] [42]. These sensors can struggle under adverse weather or are dependent on good lighting conditions [4]. Several approaches leverage inertial measurement units (IMUs) [11] as they are economic high-frequency sensors accurate in short periods of time [13], but considered alone, they drift substantially.

Automotive radars are robust to bad weather and low lighting conditions, and their production costs are significantly lower than those of a LiDAR [8]. Additionally, they have been proven as successful alternatives for performing odometry and localization [6] [8]. A combination of automotive radars and IMUs can exploit the advantages of both sensing modalities while being cost-effective and, as we will show, offering the potential for high accuracy in SLAM.

However, as a result of the radar’s working principle, the output point cloud of the sensor is sparse and noisy compared to LiDAR, leading to low performance of methods designed

Manuscript received: Dec 20, 2024; Revised: Feb 6, 2025; Accepted: Mar 21, 2025. This paper was recommended for publication by Editor Javier Civera upon evaluation of the Associate Editor and Reviewers’ comments. (Corresponding author: Daniel Casado Herraiez.)

D. Casado Herraiez and M. Zeller are with CARIAD SE, 38440 Wolfsburg, Germany, and also with the University of Bonn, 53113 Bonn, Germany (e-mail: dcasadoherraiez@gmail.com).

D. Wang is with the Department of Informatics XVII Robotics, Julius Maximilians-University Wuerzburg 97070, Germany.

J. Behley is with the Center for Robotics, University of Bonn, 53113 Bonn, Germany.

M. Heidingsfeld is with the CARIAD SE, 38440 Wolfsburg, Germany.

C. Stachniss is with the Center for Robotics, University of Bonn, 53113 Bonn, Germany, and also with the Lamarr Institute for Machine Learning and Artificial Intelligence, 44227 Dortmund, Germany.

for LiDARs when directly applied to radar. This limitation is evident in some radar SLAM approaches that adopt LiDAR techniques for loop closure detection [21] [49]. Moreover, automotive radars also provide additional information including a measurement of the point's relative radial velocity and a radar-cross section (RCS) related to the target's reflectivity and angle of incidence that can be exploited for scan-matching and place recognition [6] [8] [32].

The main contribution of this work is a novel radar-inertial SLAM system for autonomous vehicles that fully leverages the characteristics of automotive radars for odometry estimation and loop closure, see Fig. 1. We formulate the problem in a modular manner, with three distinct components: a local pose graph, a global pose graph, and a loop closure detector. The local pose graph integrates local information from the IMU and radar scan registration, leveraging the Doppler velocities for scan matching. This serves as an initial pose estimate for the global pose graph. Global pose graph optimization is performed including the loop closures from our loop detector. Our system is open-source¹ and achieves state-of-the-art performance in radar-inertial SLAM using public autonomous driving data.

In sum, we make three key claims. Our approach (i) achieves state-of-the-art results in radar-inertial SLAM for autonomous driving on publicly available data; (ii) introduces a two-layered system structure for radar-inertial SLAM, comprising a local graph and a global graph leading to enhanced accuracy; (iii) exploits the velocity and radar cross-section to improve SLAM accuracy in sparse and noisy radar scans.

II. RELATED WORK

We present an overview of relevant approaches in odometry and SLAM, focusing on point cloud-based methods. Our discussion covers existing LiDAR and radar techniques, as well as methods that combine these sensors with IMUs.

LiDAR-based odometry and SLAM techniques estimate the relative pose of the vehicle using onboard LiDAR sensors while simultaneously creating a map of the environment. Methods that rely only on LiDAR scans usually perform iterative alignment between point clouds to estimate the transformation between them. The most popular approaches in this domain are based on point-to-point matching [16] [41], plane feature extraction and matching [10] [12], or a combination of both [29]. LiDAR-only odometry may degrade in featureless environments [40], which can be addressed by introducing additional pose information from an IMU. The high-frequency accelerometer and gyroscope readings can be integrated over time [13] to provide an additional constraint during odometry. Some approaches exploit this by combining the LiDAR and IMU measurements with an iterated Kalman filter [47] to estimate the pose of the robot. Others employ pose graphs [30] [37] [38] combining IMU and LiDAR odometry information that can be jointly optimized within a single graph. A major disadvantage of single-factor graph approaches like LIO-SAM [38] is that in their implementation, they do not combine the IMU and radar odometry information within

the global pose estimation. Instead, they use the IMU as a source of high-frequency odometry but construct the backend graph containing only the estimated poses from LiDAR scan registration. More recent approaches propose to integrate a local submap and a global map [26] [27] maintaining global consistency. Grisetti et al. [15] also propose a hierarchical optimization method that combines multiple graphs and enables accurate odometry. Nevertheless, although LiDAR scans are dense and high resolution, the sensors are affected by adverse weather and are hard to pack within consumer autonomous vehicles. Moreover, due to the sparse and noisy properties of radar point clouds, applying LiDAR methods to radar scans reduces pose estimation performance. These challenges call for SLAM techniques specifically designed for radar data.

Radar-based odometry and SLAM use radar sensors to estimate the relative pose of the vehicle and construct the environment map. It is important to differentiate between two radar categories, scanning and automotive radars. Scanning radars provide a 2D intensity image of the vehicle's surroundings. While some authors extract features from the intensity images and match them over time [2] [18], others perform signal processing techniques to directly estimate the motion from the radar image scans [46]. Full SLAM approaches with loop closure have also been proposed proving successful in adverse weather scenarios [1] [20] [43]. However, similar to LiDARs, scanning radars are bulky and too expensive to be mounted in consumer vehicles. In contrast, automotive radar sensors are smaller, more affordable, and provide a sparse and noisy 3D point cloud that also contains the Doppler velocity and RCS information of each point. The earliest work in the automotive radar domain directly computes the ego-pose based on the Doppler velocities of the points and the kinematics of autonomous cars [24] [33]. Others have adapted LiDAR methods to work with radar velocities [33] or modeled the uncertainty in radar measurements leading to an enhancement in odometry performance [?]. Another way of handling sparsity and uncertainty in radar point clouds is using the normal distribution transform [18], using scan-to-map matching [8] and/or filtering the scans based on RCS information [7] [32]. More recent approaches employ scan-matching between radar frames [8] [28] [44] or use semantic features from the environment [22] to estimate the change in pose over time. Full automotive radar SLAM approaches have also been proposed optimizing a pose graph with loop closures [17] [19] [31] [36] [49]. The noisyness of radar scans, however, can lead to noisy trajectory estimates, which can be improved using additional input from an IMU. Several methods that combine automotive radar and IMU information use Kalman filters [3] [11] [50] or continuous-time optimization [34] to estimate the pose of the vehicle over time. Recent research has shown how graph-based methods are advantageous in terms of accuracy and simplicity of the system [14] [44]. However, they rely on a single global factor graph that contains all the measurement information and utilize LiDAR techniques for loop closure limiting their performance.

In our method, we propose a full SLAM framework specific for automotive radars taking inspiration from LiDAR approaches that maintain multiple pose graphs [26] [27] and

¹<https://github.com/PRBonn/RaI-SLAM>

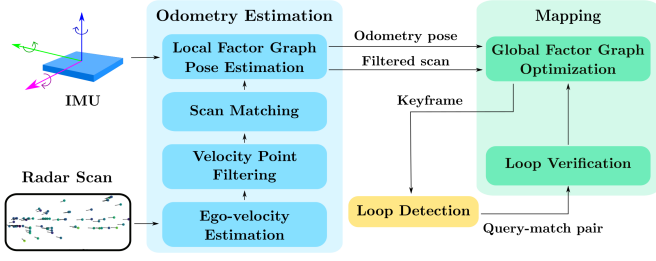


Fig. 2: Structure of our RAI-SLAM system. It consists of three main modules. The odometry estimation module optimizes the local graph for radar-inertial odometry and filtering of the point clouds. The mapping module optimizes the global graph containing odometry information and loop closures. The loop detection module recognizes revisited places and adds them to the global graph.

leverage hierarchical pose graph optimization [15]. Our system is composed of a radar-inertial odometry frontend with a loop detection and loop closure backend, and leverages the additional Doppler velocity and RCS information provided by radars for scan matching and loop closure detection.

III. OUR APPROACH FOR RADAR-INERTIAL SLAM

Our approach estimates the pose of an autonomous vehicle over time and simultaneously constructs a map of the environment. Our system architecture comprises three core modules for odometry pose estimation, loop detection, and global optimization, see Fig. 2. Furthermore, having a double-graph structure enables our system to effectively incorporate short-term and large-scale information, as shown in Fig. 3. In the odometry estimation frontend, the local factor graph collects the most recent information from the IMU and radar scan-matching and marginalizes nodes outside of a fixed window to maintain computational efficiency. In the backend, the global factor graph collects the optimized data from the local pose graph, verifies loop detections, and integrates loop closure information, improving large-scale consistency. This reduces odometry drift and performs global optimization by combining all available data.

We employ incremental smoothing and mapping [23] to optimize the local and global graphs that compute the state of the vehicle. The optimization problem for a single factor graph can be expressed as

$$\mathbf{x}^* = \arg \min_{\mathbf{x}} \sum_k \|f_k(\mathbf{x}_{k_1}, \mathbf{x}_{k_2}, \dots, \mathbf{x}_{k_M}; o_k)\|_{\Sigma_k}^2, \quad (1)$$

where \mathbf{x}^* represents the optimal solution to the optimization problem, Σ_k is the information matrix associated to the k^{th} factor, and $f_k(\mathbf{x}_{k_1}, \mathbf{x}_{k_2}, \dots, \mathbf{x}_{k_M}; o_k)$ represents the residual function for the k^{th} factor involving states $\mathbf{x}_{k_1}, \mathbf{x}_{k_2}, \dots, \mathbf{x}_{k_M}$ and the observed measurement o_k . In the following, we will represent states and observations from the local and global graphs using superscripts l and g , respectively.

A. Local Factor Graph

The goal of the local factor graph is to estimate the pose of the vehicle by combining scan-matching information with

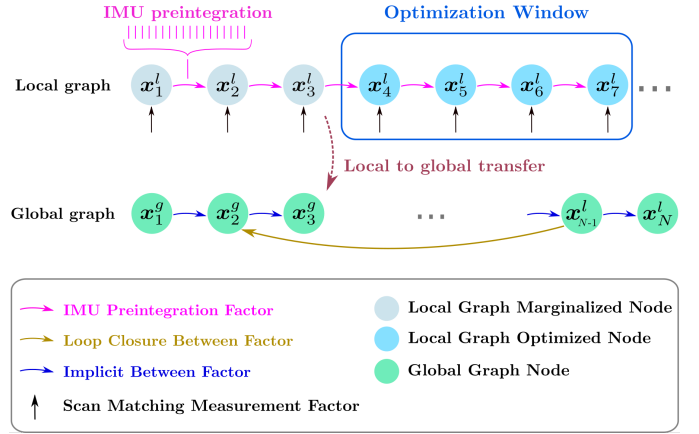


Fig. 3: Diagram of the local and global graph in our system. The scan-matching and IMU information is collected within the optimization window. The global graph incorporates this data and performs loop closure optimization within the estimated poses.

pre-integrated accelerometer and gyroscope measurements. We exploit the robustness of the scan matcher while leveraging the short-term accuracy provided by high-frequency IMU sensors employing two separate factors.

1) *Scan-Matching* : In our framework, the main source of odometry is derived from estimating the relative transform between radar scans using scan-to-map registration. However, an incorrect alignment of one single radar scan can lead to an incorrect absolute trajectory estimate. We mitigate this by utilizing additional sources of information beyond the radar point coordinates for scan alignment. One approach is to employ the Doppler velocities during ICP registration [8]. However, this is sensitive to sensor calibration. We expand on Radar-ICP [8] relying only on the Doppler velocities as an initial prior for scan registration, and as a pre-filtering step to remove dynamic object outliers.

Following Kellner et al. [24], given the azimuth angle of each measurement θ_i and its measured Doppler velocity v_i , we first estimate the velocity of the automotive radar sensor $\mathbf{v}_s = [v_{s,x}, v_{s,y}]^T$ as a least squares problem of the form:

$$\begin{bmatrix} v_1 \\ \vdots \\ v_N \end{bmatrix} = \begin{bmatrix} \cos(\theta_1) & \sin(\theta_1) \\ \vdots & \vdots \\ \cos(\theta_N) & \sin(\theta_N) \end{bmatrix} \begin{bmatrix} v_{s,x} \\ v_{s,y} \end{bmatrix}. \quad (1)$$

Assuming an autonomous vehicle with no vertical and lateral movement, the linear velocity $\mathbf{v}_c \in \mathbb{R}^3$ and angular rate $\omega_c \in \mathbb{R}^3$ can be computed as

$$\mathbf{v}_c = \mathbf{R}_s^c \mathbf{v}_s - \omega_c \times \mathbf{t}_{sc}, \quad (2)$$

where $\mathbf{R}_s^c \in SO(3)$ and $\mathbf{t}_{sc} \in \mathbb{R}^3$ are the rotation and translation from the car to the sensor.

Given the time difference Δt between the current and previous scan, the estimated ego-motion is given as

$$\mathbf{T}_{\text{init}} = \begin{bmatrix} \exp([\omega_c]_{\times} \Delta t) & \mathbf{J} \mathbf{v}_c \Delta t \\ \mathbf{0}^T & 1 \end{bmatrix}, \quad (3)$$

where $[\omega_c]_{\times}$ is the skew-symmetric matrix of ω_c , and \mathbf{J} is the left Jacobian of $SO(3)$ [39].

We use $\mathbf{T}_{\text{init}} \in SE(3)$ as the initial estimate for ICP and perform scan-to-map matching, which helps to handle the sparsity of radar scans. The scan-to-map factor contains its corresponding odometry measurement $o_{\text{odom}_i}^l$ associated to the ICP transformation estimate \mathbf{T}_i according to

$$f_{\text{odom}}^l(\mathbf{x}_i^l; o_{\text{odom}_i}^l) = f_{\text{odom}}^l(\mathbf{x}_i^l; \mathbf{T}_i). \quad (4)$$

2) *IMU Preintegration* : While we consider scan-matching as our main source of information within the global graph, the accelerometer and gyroscope readings from the IMU help to smooth the trajectory and provide high-frequency pose estimation in environments with few geometric features. The raw measurements from the sensor are given as

$$\tilde{\omega}_t = \omega_t + \mathbf{b}_t^g + \boldsymbol{\eta}_t^g, \quad (5)$$

$$\tilde{\mathbf{a}}_t = \mathbf{R}_t^\top (\mathbf{a}_t - \mathbf{g}) + \mathbf{b}_t^a + \boldsymbol{\eta}_t^a, \quad (6)$$

which consist of the true gyroscope and accelerometer measurement $\omega_t, \mathbf{a}_t \in \mathbb{R}^3$ an added white noise $\boldsymbol{\eta}_t^g, \boldsymbol{\eta}_t^a \in \mathbb{R}^3$, an added bias $\mathbf{b}_t^g, \mathbf{b}_t^a \in \mathbb{R}^3$, the gravity vector $\mathbf{g} \in \mathbb{R}^3$ and the rotation of the IMU in the world frame $\mathbf{R}_t \in SO(3)$ [13].

Using the IMU measurements, we estimate the rotation $\mathbf{R}_{t+\Delta t} \in SO(3)$, velocity $\mathbf{v}_{t+\Delta t} \in \mathbb{R}^3$, and position $\mathbf{p}_{t+\Delta t} \in \mathbb{R}^3$ after a time increment Δt following the integration of IMU measurements over time [13]:

$$\mathbf{R}_{t+\Delta t} = \mathbf{R}_t \text{Exp}((\tilde{\omega}_t - \mathbf{b}_t^g - \boldsymbol{\eta}_t^g)\Delta t), \quad (7)$$

$$\mathbf{v}_{t+\Delta t} = \mathbf{v}_t + \mathbf{g}\Delta t + \mathbf{R}_t(\tilde{\mathbf{a}}_t - \mathbf{b}_t^a - \boldsymbol{\eta}_t^a)\Delta t, \quad (8)$$

$$\mathbf{p}_{t+\Delta t} = \mathbf{p}_t + \mathbf{v}_t\Delta t + (\mathbf{g} + \mathbf{R}_t(\tilde{\mathbf{a}}_t - \mathbf{b}_t^a - \boldsymbol{\eta}_t^a))\frac{1}{2}\Delta t^2. \quad (9)$$

The estimate of the relative motion of the sensor $\Delta \mathbf{R}_{i,j}$, $\Delta \mathbf{v}_{i,j}$, and $\Delta \mathbf{p}_{i,j}$ measured by the IMU between two radar measurements at times i and j is given by:

$$\Delta \mathbf{R}_{i,j} = \prod_{k=i}^{j-1} \text{Exp}((\tilde{\omega}_k - \mathbf{b}_k^g + \boldsymbol{\eta}_k^g)\Delta t), \quad (10)$$

$$\Delta \mathbf{v}_{i,j} = \sum_{k=i}^{j-1} \Delta \mathbf{R}_{i,k} (\tilde{\mathbf{a}}_k - \mathbf{b}_k^a - \boldsymbol{\eta}_k^a) \Delta t, \quad (11)$$

$$\Delta \mathbf{p}_{i,j} = \sum_{k=i}^{j-1} \left[\Delta \mathbf{v}_{i,k} \Delta t + \Delta \mathbf{R}_{i,k} (\tilde{\mathbf{a}}_k - \mathbf{b}_k^a - \boldsymbol{\eta}_k^a) \frac{1}{2} \Delta t^2 \right]. \quad (12)$$

The relative motion information estimated from the IMU is added to the local factor graph as an additional constraint for the optimization, in a similar manner to LIO-SAM [38]. Our approach, however, is adapted to incorporate radar data that is optimized together with the IMU factors. Following the formulation in Eq. (1), this results in the following factors with the corresponding observations of the rotation $o_{\mathbf{R}_{i,i-1}}^l$, velocity $o_{\mathbf{v}_{i,i-1}}^l$, and position $o_{\mathbf{p}_{i,i-1}}^l$:

$$f_{\text{IMU}}^l(\mathbf{x}_{i-1}^l, \mathbf{x}_i^l; o_{\mathbf{R}_{i,i-1}}^l) = f_{\text{IMU}}^l(\mathbf{x}_{i-1}^l, \mathbf{x}_i^l; \Delta \mathbf{R}_{i,i-1}) \quad (13)$$

$$f_{\text{IMU}}^l(\mathbf{x}_{i-1}^l, \mathbf{x}_i^l; o_{\mathbf{v}_{i,i-1}}^l) = f_{\text{IMU}}^l(\mathbf{x}_{i-1}^l, \mathbf{x}_i^l; \Delta \mathbf{v}_{i,i-1}) \quad (14)$$

$$f_{\text{IMU}}^l(\mathbf{x}_{i-1}^l, \mathbf{x}_i^l; o_{\mathbf{p}_{i,i-1}}^l) = f_{\text{IMU}}^l(\mathbf{x}_{i-1}^l, \mathbf{x}_i^l; \Delta \mathbf{p}_{i,i-1}). \quad (15)$$

We jointly optimize the local graph that contains scan matching factors from Eq. (4), and the IMU factors from

Eqs. (13), (14), (15) in a windowed manner, following Eq. (1). All the factors outside the marginalization window are excluded from the optimization, keeping a bounded size of the optimizable graph.

B. Global Factor Graph

The global factor graph corrects for the drift accumulated in the local factor graph by performing loop closure optimization. Two key processes are performed in the global graph, as illustrated in Fig. 3. First, we transfer the information from the local to the global factor graph. This results in a global factor graph with implicit factors that contain information about the IMU and scan registration. We, then, include loop closure factors into the graph to correct odometry drift.

1) *Transfer from Local to Global Graph* : The data transfer from the local to the global graph is performed on the latest node of the windowed optimization. It serves as an initialization of the global poses before loop closure [15]. We create an implicit factor containing the relative pose measurement o_i^g between the current and the previous scan coming from both, scan-to-map matching and IMU, following

$$f(\mathbf{x}_{i-1}^g, \mathbf{x}_i^g; o_i^g) = f(\mathbf{x}_{i-1}^g, \mathbf{x}_i^g; \mathbf{T}_{i-1}^{l-1} \mathbf{T}_i^l), \quad (16)$$

where $\mathbf{T}_{i-1}^l, \mathbf{T}_i \in SE(3)$ are the poses of the previous and current frames in the local graph, respectively.

2) *Loop Closure* : Including the implicit factors in the global graph does not solve the problem of a drifting trajectory. We introduce loop detection and closure to correct the accumulated errors. We employ three criteria for loop closing. First, a radar place recognition model [6] finds a matching location and returns a similarity score. Then, an odometry distance measurement estimates the feasibility of the loop closure in space. Finally, we propose an intuitive point correspondence metric that replaces the common ICP distance score to measure the matching quality [38].

The place recognition module identifies locations that have been visited in the past. While other works use ScanContext [44] [45], learning-based approaches [6] have shown superior place recognition performance. We employ SPR [6] as our scan encoder, which captures point-neighbor information from the radar scan and an overall RCS distribution of the point cloud. We match two radar scans if

$$\|\mathbf{X}_{\text{query}} - \mathbf{X}_{\text{match}}\| < \delta_{\text{PR}}, \quad (17)$$

where δ_{PR} is a predefined threshold and $\mathbf{X}_{\text{query}}, \mathbf{X}_{\text{match}} \in \mathbb{R}^{256}$ are the SPR encodings of the query and matching scans.

The odometry distance metric is inspired by TBV-SLAM [1] and verifies that the position of the query frame $\mathbf{t}_{\text{query}} \in \mathbb{R}^3$ and matching frame $\mathbf{t}_{\text{match}} \in \mathbb{R}^3$ are within a reasonable spatial threshold δ_{odom} proportional to the accumulated length of the current trajectory d_{length} :

$$\frac{\|\mathbf{t}_{\text{query}} - \mathbf{t}_{\text{match}}\|}{d_{\text{length}}} < \delta_{\text{odom}}. \quad (18)$$

After a loop candidate has been detected based on similarity and odometry distance, we use ICP to estimate the transformation $\mathbf{T}_m^q \in SE(3)$ between the match and the

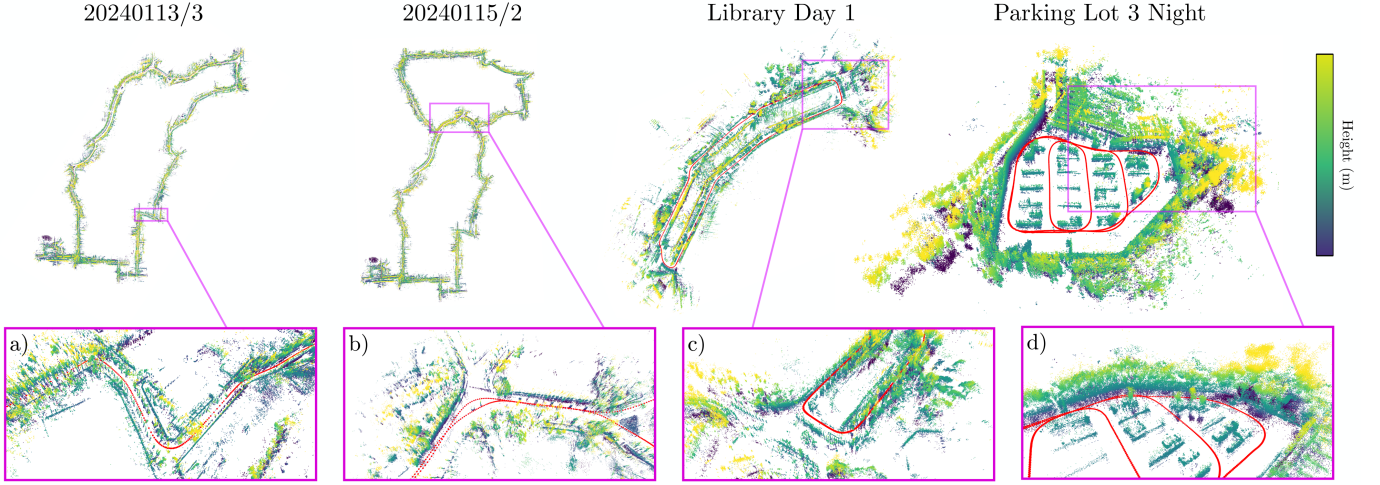


Fig. 4: Pose estimation (shown in **red**) and mapping results for automotive sequences of the SNAIL-Radar [21] and HeRCULES [25] datasets. Snapshots (b) and (d) show segments with loop closures. Movement in the vertical direction has been omitted for the map plots.

	20240113/3			20240113/1			20240115/2			20240123/2			20240123/3			Mean		
	RTE [m]	RRE [°]	ATE [m]	RTE [m]	RRE [°]	ATE [m]	RTE [m]	RRE [°]	ATE [m]	RTE [m]	RRE [°]	ATE [m]	RTE [m]	RRE [°]	ATE [m]	RTE [m]	RRE [°]	ATE [m]
LIO-SAM (lidar)	0.014	0.066	52.6	0.015	0.131	0.3	0.019	0.084	48.5	5.983	0.786	57.6	0.013	0.080	5.9	1.209	0.229	33.0
KISS-ICP (radar)	0.240	0.155	68.4	0.118	0.179	4.4	0.232	0.134	147.1	0.269	0.117	167.8	0.222	0.155	45.9	0.216	0.148	86.7
Radar-ICP	<u>0.238</u>	0.156	<u>18.2</u>	0.120	0.174	<u>3.9</u>	0.229	0.131	<u>31.6</u>	<u>0.252</u>	0.112	37.5	0.221	0.151	7.9	0.212	0.145	<u>19.8</u>
4DRadarSLAM	0.737	1.170	53.2	0.460	1.074	8.9	0.663	1.179	491.2	0.864	0.901	454.5	0.503	0.983	142.1	0.645	1.061	230.0
Graph-RIO	-	-	-	0.169	0.172	9.5	0.195	0.172	763.4	-	-	-	0.266	0.168	497.0	<u>0.210</u>	0.170	423.3
RIV-SLAM	0.213	0.142	30.2	0.113	<u>0.171</u>	4.1	<u>0.219</u>	<u>0.128</u>	33.1	0.224	0.101	<u>35.5</u>	0.201	<u>0.140</u>	<u>6.1</u>	0.194	<u>0.137</u>	21.8
Ours	0.242	<u>0.147</u>	4.7	<u>0.115</u>	0.164	3.4	0.220	0.117	7.9	0.256	<u>0.104</u>	8.4	<u>0.219</u>	0.139	3.5	<u>0.210</u>	0.134	5.6

TABLE I: Comparison to state-of-the-art radar and LiDAR odometry and SLAM approaches on the SNAIL-Radar dataset [21].

	Mountain Day 1			Library Day 1			Sports Complex Day 1			Parking Lot 3 Night			Street Day 1			Mean		
	RTE [m]	RRE [°]	ATE [m]	RTE [m]	RRE [°]	ATE [m]	RTE [m]	RRE [°]	ATE [m]	RTE [m]	RRE [°]	ATE [m]	RTE [m]	RRE [°]	ATE [m]	RTE [m]	RRE [°]	ATE [m]
KISS-ICP (LiDAR)	0.064	0.068	21.9	0.057	0.065	8.4	0.057	0.084	8.3	0.077	0.146	2.1	-	-	-	0.064	0.091	10.2
KISS-ICP (radar)	0.057	<u>0.067</u>	<u>113.3</u>	0.046	<u>0.061</u>	8.7	<u>0.046</u>	<u>0.069</u>	8.2	0.065	0.101	2.5	0.033	0.044	5.5	0.049	0.069	<u>27.7</u>
Radar-ICP	<u>0.055</u>	<u>0.067</u>	118.6	0.049	0.064	10.2	0.049	0.071	<u>7.1</u>	0.058	0.089	3.4	0.022	<u>0.028</u>	11.7	<u>0.029</u>	<u>0.064</u>	30.2
RIV-SLAM	0.077	0.084	206.9	0.014	0.064	<u>4.2</u>	-	-	-	0.020	<u>0.075</u>	<u>2.4</u>	<u>0.010</u>	0.042	<u>10.7</u>	0.030	0.066	56.0
Ours	0.019	0.024	11.2	<u>0.019</u>	0.021	2.1	0.020	0.023	3.9	<u>0.022</u>	0.033	2.2	0.008	0.005	11.5	0.018	0.021	6.2

TABLE II: Comparison of the best-performing odometry and SLAM methods on the HeRCULES dataset [25].

query frames. However, we observe that the distance score used in LiDAR approaches is not a reliable matching quality measure in automotive radars. While in LiDAR a low average distance between point correspondences indicates a good scan alignment, this is not always the case in sparse and noisy radar point clouds. Good alignments may still return high average distance values due to noise outliers, bad alignments may still return reasonable average distance values, and if two radar scans are taken at different times, the amount of noise and reflection outliers may vary.

To address these limitations, we propose a new scoring method for loop-matching quality. Instead of using the mean distance of corresponding points, we measure that the number of points after ICP registration that have a correspondence within a radius r is greater than a threshold δ_d . This gives an intuitive and reliable estimate of the matching quality. The comparison of the score with δ_d is given as:

$$\left[\frac{1}{|\mathcal{C}|} \sum_{(\mathbf{q}, \mathbf{m}) \in \mathcal{C}} \mathbb{I}\{\|\mathbf{q} - \mathbf{m}\| < r\} \right] > \delta_d, \quad (19)$$

where $(\mathbf{q}, \mathbf{m}) \in \mathcal{C}$ is the set of correspondences between query \mathbf{q} and matching scan \mathbf{m} , and $\mathbb{I}\{c\}$ is the indicator function, returning 1 if condition c is true, and 0 otherwise.

Once the three criteria have been verified and the transformation between query and match \mathbf{T}_m^q has been computed, we express the final loop factor with the corresponding observation $o_{\text{loop}_i}^g$ following

$$f(\mathbf{x}_q^g, \mathbf{x}_m^g; o_{\text{loop}_i}^g) = f(\mathbf{x}_q^g, \mathbf{x}_m^g; \mathbf{T}_m^{q^{-1}}). \quad (20)$$

The global graph containing the implicit factors from Eq. (16), and the loop factors in Eq. (20), is optimized independent of the local graph, following the procedure in Eq. (1). We optimize the global graph including all transferred and loop factors as a backend in a separate thread.

IV. EXPERIMENTAL EVALUATION

The main focus of our research is to develop a radar-inertial SLAM system for autonomous vehicles that exploits the information provided by radar sensors for large-scale pose estimation and mapping. We present our experiments to

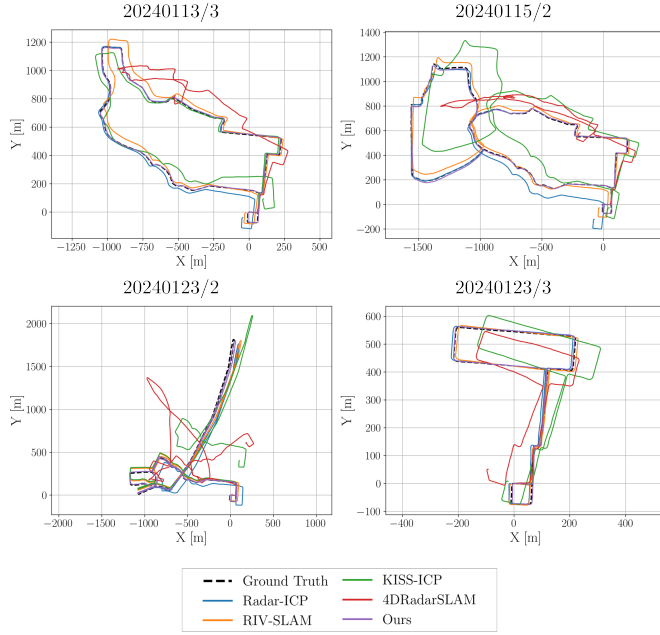


Fig. 5: Qualitative comparison of our approach to the state of the art on the SNAIL-Radar dataset [21].

show the capabilities of our method, called RaI-SLAM. The results explicitly support our key claims that our approach: (i) achieves state-of-the-art results in radar-inertial SLAM for autonomous driving on publicly available data; (ii) introduces a two-layered system structure for radar-inertial SLAM, comprising a local graph and a global graph leading to enhanced accuracy; (iii) exploits the velocity and radar cross-section to improve SLAM accuracy in sparse and noisy radar scans.

A. Implementation Details and Experimental Setup

We implement RaI-SLAM using ROS2 for communication between modules and a GTSAM [9] factor graph optimization framework, employing a fixed-lag smoother for the local graph frontend and assuming constant covariance values for odometry and loop closure. We evaluate our approach on the SNAIL-Radar [21] and HeRCULES [25] datasets, as other existing datasets have either short sequences without loop closures [5] [35], inaccuracies in the ground truth [48], or no IMU sensors [31]. Within the SNAIL-Radar dataset, we select sequences 20240113/3, 20240113/1, 20240115/2, 20240123/2, and 20240123/3, which cover the entire area of the dataset and are recorded using an SUV vehicle. From the HeRCULES [25] dataset, we select five diverse sequences containing loop closures, “Mountain Day 1”, “Library Day 1”, “Sports Complex Day 1”, “Parking Lot 3 Night”, and “Street Day 1”. For the loop closure module with SPR [6], we train the model on the 4DRadarDataset [31], which also uses a different radar sensor as the SNAIL-Radar dataset, indicating the generalizability of our system. We use the ARS548 radar for the evaluation of all methods except in the comparison with 4DRadarSLAM [49] whose parameters have been optimized by the original authors for the Oculii Eagle. We evaluate all trajectories on the plane, discarding all vertical movement.

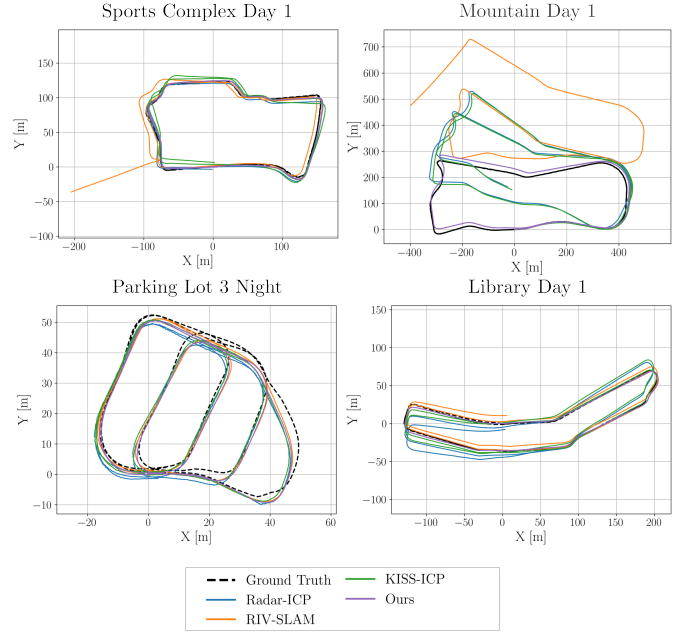


Fig. 6: Qualitative comparison of our approach to the state of the art on the HeRCULES dataset [25].

	20240115/2			Library Day 1		
	RTE [m]	RRE [°]	ATE [m]	RTE [m]	RRE [°]	ATE [m]
Scan matching	0.247	0.258	112.4	0.087	0.361	18.47
Scan matching with vel. prior	0.244	0.252	40.04	<u>0.085</u>	<u>0.346</u>	17.03
Local graph with IMU	0.222	0.145	39.90	0.024	0.210	16.99
Local and global without IMU	0.252	0.253	5.642	0.087	0.362	<u>3.652</u>
Local and global with IMU	<u>0.229</u>	<u>0.146</u>	<u>5.726</u>	0.024	0.210	3.389

TABLE III: Ablation studies on sequence 20240115/2 from the SNAIL-Radar dataset [21] and “Library Day 1” from the HeRCULES [25] dataset.

The metrics used for evaluation consist of the relative translation (RTE) and rotation (RRE) error between pose estimations measuring short-term accuracy, and the absolute trajectory error (ATE) to estimate absolute performance. The best and second best results on automotive radar sensors are **bolded** and underlined, respectively.

B. Comparison with the State of the Art

The first experiment evaluates the performance of our method and demonstrates that it achieves state-of-the-art results in publicly available automotive radar-inertial SLAM data. We compare our method against other approaches in the SNAIL-Radar dataset [21], and then compare the best-performing methods in the HeRCULES dataset [25]. Quantitative results are shown in Tabs. I and II, and qualitative results are presented in Figs. 4, 5 and 6.

The baseline methods employed in the comparison are RIV-SLAM [44], which leverages a single graph integrating all sensor data; Graph-RIO [14], which uses a factor graph but excludes scan-matching and loop closure; 4DRadarSLAM [49], which relies solely on the Oculii radar data without IMU information; Radar ICP [8], a state-of-the-art radar odometry method; KISSL-ICP [41], a LiDAR odometry approach applied directly on radar and LiDAR point clouds; and LIO-SAM [38],

a LiDAR-inertial odometry system used in this work to benchmark sensor performance.

As it can be observed in Fig. 5 and Fig. 6, LiDAR approaches like KISS-ICP applied to radar point clouds have an increased drift, sometimes leading to failures in scenarios with few geometric features like in the long road from sequence 2024123/2. Odometry-only Radar-ICP has less drift but lacks loop closure to correct it, visible in sequences like “Mountain Day 1”. RIV-SLAM achieves high overall relative and absolute accuracy but accumulates drift over longer sequences. It relies on intensity ScanContext [45] and the regular ICP distance score, failing to detect certain loops and correct accumulated errors. Our method, which relies on a radar-oriented loop detection procedure, identifies the loops and effectively integrates local and global information, leading to a lower mean ATE while maintaining competitive accuracy for relative pose estimation. RIV-SLAM also fails in sequence “Sports Complex Day 1”, where an incorrect point cloud alignment leads to an inconsistent trajectory. This is prevented with our velocity-based ICP initialization. Additionally, our radar-inertial SLAM technique is on par with LiDAR approaches. Specifically, it presents an advantage in the highly dynamic scenario from “Street Day 1”, where the LiDAR odometry approach fails to complete the sequence due to the high amount of point outliers corresponding to moving objects.

C. Ablation Studies

The second experiment evaluates how our two-layered graph structure and how exploiting the characteristics of radar point clouds for odometry and loop closure contribute to the final accuracy. We perform the evaluation on the sequence 20240115/2 [21] and “Library Day 1” [25] as good examples of long trajectories that include loops. The main components are the velocity prior for scan-matching, the local graph with and without IMU, and the global loop closure for drift correction. The results are shown in Tab. III.

Employing the estimated ego velocity as the prior for ICP results in a notable improvement of the absolute errors. Without the velocity prior, a wrong match between the current scan and the local map would lead to an incorrect trajectory for the following measurements, affecting the absolute trajectory error. Furthermore, adding the IMU measurements reduces relative errors between scans but preserves a similar absolute error due to the way the local factor graph is constructed. Adding the loop closure additionally improves the global results, resulting in a system with reduced local and global trajectory errors.

Based on the results of pose graph optimization with constant covariances in sequence “20240115/2”, however, we often see that well-aligned local pose estimates can get misaligned in the global optimization. Therefore, the estimation of realistic pose and loop closure covariances is a promising avenue of future research that could potentially increase the consistency of the resulting trajectory and map.

V. CONCLUSION

In this paper, we presented a system for radar-inertial SLAM that exploits the advantages of automotive radars and deals

with their sparse and noisy output point clouds *without* relying on external GNSS information. We proposed a novel system composed of a local and global pose graph that incorporates IMU and radar information. Our odometry frontend keeps a low local translation and rotation error and our loop closure procedure reduces drift in the absolute trajectory. We implemented and evaluated our approach on real-world scenarios supporting all claims made in this paper. The experiments suggest that our method achieves high performance for estimating the global pose of the vehicle by solely relying on onboard radar-inertial sensing, with each component of the system contributing to the final pose estimation accuracy.

REFERENCES

- [1] D. Adolfsson, M. Karlsson, V. Kubelka, M. Magnusson, and H. Andreasson. Tbv radar slam: Trust but verify loop candidates. *IEEE Robotics and Automation Letters (RA-L)*, 8(6):3613–3620, 2023.
- [2] R. Aldera, D. De Martini, M. Gadd, and P. Newman. What could go wrong? introspective radar odometry in challenging environments. In *Proc. of the IEEE Intl. Conf. on Intelligent Transportation Systems (ITSC)*, 2019.
- [3] Y. Almalioglu, M. Turan, C.X. Lu, N. Trigoni, and A. Markham. Millirio: Ego-motion estimation with low-cost millimetre-wave radar. *IEEE Sensors Journal*, 21(3):3314–3323, 2020.
- [4] 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.
- [5] 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.
- [6] D. Casado Herraiez, L. Chang, M. Zeller, L. Wiesmann, J. Behley, M. Heidingsfeld, and C. Stachniss. Spr: Single-scan radar place recognition. *IEEE Robotics and Automation Letters (RA-L)*, 9(10):9079 – 9086, 2024.
- [7] D. Casado Herraiez, F. Kaschner, M. Zeller, D. Muhle, J. Behley, M. Heidingsfeld, D. Cremers, and C. Stachniss. Ground-aware automotive radar odometry. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2025.
- [8] D. Casado Herraiez, M. Zeller, L. Chang, I. Vizzo, M. Heidingsfeld, and C. Stachniss. Radar-only odometry and mapping for autonomous vehicles. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2024.
- [9] F. Dellaert and M. Kaess. *Factor Graphs for Robot Perception*. Foundations and Trends in Robotics, 2017.
- [10] P. Dellenbach, J.E. 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.
- [11] C. Doer and G.F. Trommer. An ekf based approach to radar inertial odometry. In *Proc. of the Intl. Conf. on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, 2020.
- [12] S. Ferrari, L.D. Giammarino, L. Brizi, and G. Grisetti. Mad-icp: It is all about matching data – robust and informed lidar odometry. *IEEE Robotics and Automation Letters (RA-L)*, 9(11):9175–9182, 2024.
- [13] C. Forster, L. Carlone, F. Dellaert, and D. Scaramuzza. On-manifold preintegration for real-time visual-inertial odometry. *IEEE Trans. on Robotics (TRO)*, 33(1):1–21, 2016.
- [14] R. Girod, M. Hauswirth, P. Pfreundschuh, M. Biasio, and R. Siegwart. A robust baro-radar-inertial odometry m-estimator for multicopter navigation in cities and forests. In *Proc. of Intl. Conf. on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, 2024.
- [15] G. Grisetti, R. Kümmerle, C. Stachniss, U. Frese, and C. Hertzberg. Hierarchical Optimization on Manifolds for Online 2D and 3D Mapping. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2010.
- [16] T. Guadagnino, B. Mersch, I. Vizzo, S. Gupta, M.V. Malladi, L. Lofefaro, G. Doisy, and C. Stachniss. Kinematic-icp: Enhancing lidar odometry with kinematic constraints for wheeled mobile robots moving on planar surfaces. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2025.
- [17] M. Hilger, V. Kubelka, D. Adolfsson, H. Andreasson, and A.J. Lilienthal. Towards introspective loop closure in 4d radar slam. *arXiv preprint arXiv:2404.03940*, 2024.

- [18] M. Hilger, N. Mandischer, and B. Corves. Randt slam: Radar slam based on intensity-augmented normal distributions transform. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2024.
- [19] M. Holder, S. Hellwig, and H. Winner. Real-time pose graph slam based on radar. In *Proc. of the IEEE Vehicles Symposium (IV)*, 2019.
- [20] Z. Hong, Y. Petillot, A. Wallace, and S. Wang. Radarslam: A robust simultaneous localization and mapping system for all weather conditions. *Intl. Journal of Robotics Research (IJRR)*, 41(5):519–542, 2022.
- [21] J. Huai, B. Wang, Y. Zhuang, Y. Chen, Q. Li, Y. Han, and C. Toth. Snail-radar: A large-scale diverse dataset for the evaluation of 4d-radar-based slam systems. *arXiv preprint arXiv:2407.11705*, 2024.
- [22] 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.
- [23] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. Leonard, and F. Dellaert. isam2: Incremental smoothing and mapping with fluid relinearization and incremental variable reordering. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2011.
- [24] 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.
- [25] H. Kim, M. Jung, C. Noh, S. Jung, H. Song, W. Yang, H. Jang, and A. Kim. Hercules: Heterogeneous radar dataset in complex urban environment for multi-session radar slam. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2025.
- [26] K. Koide, M. Yokozuka, S. Oishi, and A. Banno. Globally consistent and tightly coupled 3d lidar inertial mapping. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2022.
- [27] K. Koide, M. Yokozuka, S. Oishi, and A. Banno. Glim: 3d range-inertial localization and mapping with gpu-accelerated scan matching factors. In *Proc. of Robotics: Science and Systems (RSS)*, 2024.
- [28] V. Kubelka, E. Fritz, and M. Magnusson. Do we need scan-matching in radar odometry? In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2024.
- [29] D. Lee, H. Lim, and S. Han. Genz-icp: Generalizable and degeneracy-robust lidar odometry using an adaptive weighting. *IEEE Robotics and Automation Letters (RA-L)*, 10(1):152–159, 2024.
- [30] K. Li, M. Li, and U.D. Hanebeck. Towards high-performance solid-state-lidar-inertial odometry and mapping. *IEEE Robotics and Automation Letters (RA-L)*, 6(3):5167–5174, 2021.
- [31] X. Li, H. Zhang, and W. Chen. 4d radar-based pose graph slam with ego-velocity pre-integration factor. *IEEE Robotics and Automation Letters (RA-L)*, 8(8):5124–5131, 2023.
- [32] Y. Li, Y. Liu, Y. Wang, Y. Lin, and W. Shen. The millimeter-wave radar slam assisted by the rcs feature of the target and imu. *Sensors*, 20(18):5421, 2020.
- [33] 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 Intl. Conf. on Microwaves for Intelligent Mobility (ICMIM)*, 2017.
- [34] 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.
- [35] A. Palffy, E. Pool, S. Baratam, J.F. Kooij, and D.M. Gavrilu. 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.
- [36] Y.S. Park, Y.S. Shin, J. Kim, and A. Kim. 3d ego-motion estimation using low-cost mmwave radars via radar velocity factor. *IEEE Robotics and Automation Letters (RA-L)*, 6(4):7691–7698, 2021.
- [37] C. Qian, Z. Xiang, Z. Wu, and H. Sun. Rf-lío: Removal-first tightly-coupled lidar inertial odometry in high dynamic environments. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2021.
- [38] T. Shan, B. Englot, D. Meyers, W. Wang, C. Ratti, and D. Rus. LIO-SAM: Tightly-coupled Lidar Inertial Odometry via Smoothing and Mapping. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2020.
- [39] J. Sola, J. Deray, and D. Atchuthan. A micro lie theory for state estimation in robotics. *arXiv preprint arXiv:1812.01537*, 2018.
- [40] T. Tuna, J. Nubert, Y. Nava, S. Khattak, and M. Hutter. X-icp: Localizability-aware lidar registration for robust localization in extreme environments. *IEEE Trans. on Robotics (TRO)*, 40:452–471, 2024.
- [41] 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):1029–1036, 2023.
- [42] O. Vysotska and C. Stachniss. Effective Visual Place Recognition Using Multi-Sequence Maps. *IEEE Robotics and Automation Letters (RA-L)*, 4(2):1730–1736, 2019.
- [43] D. Wang, Y. Duan, X. Fan, C. Meng, J. Ji, and Y. Zhang. Maroam: Map-based radar slam through two-step feature selection. *arXiv preprint arXiv:2210.13797*, 2022.
- [44] D. Wang, S. May, and A. Nuechter. Riv-slam: Radar-inertial-velocity optimization based graph slam. In *Proc. of the Intl. Conf. on Automation Science and Engineering (CASE)*, 2024.
- [45] H. Wang, C. Wang, and L. Xie. Intensity scan context: Coding intensity and geometry relations for loop closure detection. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2020.
- [46] R. Weston, M. Gadd, D. De Martini, P. Newman, and I. Posner. Fast-mbyim: 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.
- [47] W. Xu, Y. Cai, D. He, J. Lin, and F. Zhang. FAST-LIO2: Fast Direct LiDAR-Inertial Odometry. *IEEE Trans. on Robotics (TRO)*, 38(4):2053–2073, 2022.
- [48] J. Zhang, H. Zhuge, Y. Liu, G. Peng, Z. Wu, H. Zhang, Q. Lyu, H. Li, C. Zhao, D. Kircali, et al. Ntu4dradlm: 4d radar-centric multi-modal dataset for localization and mapping. In *Proc. of the IEEE Intl. Conf. on Intelligent Transportation Systems (ITSC)*, 2023.
- [49] J. Zhang, H. Zhuge, Z. Wu, G. Peng, M. Wen, Y. Liu, and D. Wang. 4dradarslam: A 4d imaging radar slam system for large-scale environments based on pose graph optimization. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2023.
- [50] 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.