This article has been accepted for publication in IEEE Transactions on Intelligent Vehicles. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TIV.2023.3324941

R^3O: Robust Radon Radar Odometry

Daniel Louback S. Lubanco, Ahmed Hashem, Markus Pichler-Scheder, Andreas Stelzer, Reinhard Feger, Thomas Schlechter

***Abstract*—Being able to measure and track positions of mobile systems is an important capability in many applications, autonomous driving being one major example. With the Robust Radon Radar Odometry algorithm, this paper presents an approach for estimating the odometry of vehicles based on radar data only. The algorithm embodies a robust method for estimating the change in orientation as a key feature. The odometry algorithm is under the realm of direct methods, and it exploits properties of the Fourier transform for decoupling the changes in orientation from the changes in translation. In the first step, the Radon transform along with phase-correlation, outlier removal, robust measure of central tendency, keyframe selection and graph optimization are used in order to achieve a robust method for estimating the change in orientation, next the translation is estimated with the support of phase-correlation. The algorithm's performance was evaluated with real world data. Significant improvements in position and orientation error in terms of relative pose error and the KITTI odometry error metric are shown as compared to other direct methods for radar based odometry.**

***Index Terms*—Radar, Odometry, Radon Transform, Autonomous Vehicles, Phase-Correlation.**

I. Introduction

THE progress of autonomy in vehicles comes with higher demand for the sensors embedded on the vehicle as well as the algorithms that use the sensors' data. This is a consequence of the fact that highly autonomous vehicles will be required to be able to perform a myriad of tasks without the active intervention of the driver.

Self-localization, i.e. self-determination of the position and orientation with regards to a reference frame over time [1], is among the many capabilities that are required for vehicles to achieve highly automated driving (HAD), i.e. level 4 and 5 of autonomy. The usual choice of sensors in self-localizing vehicles is through combining information from Global Navigation Satellite Systems (GNSS) with Inertial Measurement Units (IMU) [2]. However, GNSS-based estimation is impaired in GNSS denied regions and under multi-path conditions. Also, IMU-based odometry estimation has its pitfalls. For example, drift error, which affects any odometry algorithm, is aggravated with IMU due to the double integration required to calculate the position from acceleration data [3].

Daniel Louback S. Lubanco and Thomas Schlechter are with the University of Applied Sciences Upper Austria, 4600 Wels, Austria (e-mail: lubanco@ieee.org; thomas.schlechter@ieee.org). Markus Pichler-Scheder is with the Linz Center of Mechatronics, 4040 Linz, Austria (e-mail: markus.pichler-scheder@lcm.at). Ahmed Hashem, Andreas Stelzer and Reinhard Feger are with the Institute for Communications Engineering and RF Systems, Johannes Kepler University Linz, 4040 Linz, Austria (e-mail: name.lastname@jku.at). This paper was produced by the IEEE Publication Technology Group. They are in Piscataway, NJ.

Manuscript received Sep 29, 2022

The aforementioned requirements with regards to a vehicle's autonomy as well as possible drawbacks with the common methods for self-localization generate the motivation to analyze other possibilities for self-localization of vehicles. One alternative is to consider radar sensors. Radar is a pivotal sensor in the development of HAD due to, e.g., its high robustness under many conditions and its long range capabilities [4].

Radar ego-motion estimation is currently a topic of intensive research [5]–[21]. The approaches for calculating odometry using radar data can be roughly divided into direct and indirect odometry methods [6]. The first category relies on estimating the change in orientation and position through using area-based methods, while the latter falls under the realm of feature-based methods. Direct and feature-based odometry methods have their own advantages and disadvantages. For example, one advantage of direct methods is that they do not need to perform the feature detection step; conversely, one advantage of feature-based methods is that they usually have lower computational complexity.

In general, indirect radar odometry methods operate by detecting features in at least a pair of radar images, followed by associating the features common to the set of images, and finally estimating the ego-motion. Examples of indirect radar odometry methods include [5]–[17].

In contrast, direct radar odometry methods usually calculate the ego-motion by exploiting the decoupling between rotation and translation of the magnitude spectrum of an image.

In this paper, we present an algorithm named Robust Radon Radar Odometry (R^3O) which consists of a direct method for calculating the odometry of an agent, e.g. a vehicle, by exclusively using scanning radar data. R^3O combines the advantages of the direct radar odometry method with a procedure for calculating the rotation between radar image pairs which is more robust to interpolation errors, rotationally dependent aliasing and noise. The improvement in robustness is achieved by exploiting a property of the Radon transform along with filtering steps for removing outliers and less reliable estimations. The first step in R^3O's rotation estimation consists of finding multiple candidates for the rotation via using the Radon domain images. Next, an outlier removal method is used for removing anomalous candidates. Then, a robust location estimator and a robust scale estimator are used for calculating the rotation and its uncertainty. In addition, a keyframe method is used which dismiss pairs with high uncertainty in their rotation estimations. Finally, a local optimization is employed in order to refine the orientation estimation. The benefits of having a robust method for estimating the rotation are at least twofold. First, clearly by improving the rotation estimation

0000–0000/00$00.00 © 2022 IEEE

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Intelligent Vehicles. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TIV.2023.3324941

between pairs, the overall orientation estimation will tend to be more accurate. Second, due to the fact that, in general, in direct odometry methods one of the images is rotated using the estimated rotation before the translation estimation, a more reliable rotation estimation will likely lead to better performance in the translation estimation which induces a more reliable position estimation.

In particular, we judge that the innovative odometry pipeline (see Fig. 1) in R^3O, the use of multiple candidates for rotation estimation by exploiting the Radar transform and the subsequent steps based on the statistics of this candidates including filtering the candidates as well as keyframe selection and graph optimization based on the uncertainty estimated from the candidates, represents an important development for direct based odometry algorithms. For example, rotation estimation and its uncertainty estimation exploring other statistical methods are some of the possibilities that our framework enables

This paper is organized as follows: in section II the related work is discussed, followed by the theoretical foundations in section III that are necessary for implementing the proposed algorithm. After R^3O is explained in section IV, its performance is evaluated in section V. In section VI an ablation study along with computational complexity evaluation is presented and a conclusion is provided in section VII.

II. Related Work

A. Direct Odometry

The idea behind direct methods can be traced at least back to [22]. In that canonical paper the goal was to calculate the translation, rotation and scaling between a pair of images. For that, the authors used the rotation and translation properties of the Fourier transform (FT) which leads to the conclusion that when comparing two images, where one is a reference image I^r and the second is a translated, rotated and scaled replica of I^r called I^m, the magnitude of the FT of I^m is invariant to translation and therefore, it is a rotated replica of the FT of I^r. Following this, the two magnitude spectra are converted to their polar representation (when scaling estimation is not relevant) or to log-polar representation (when both scaling and rotation are desired to be estimated, which is in general not the case in odometry estimation); then, the polar or log-polar pair of images is phase-correlated and the resulting peak represents the estimated rotation, or the estimated rotation and scaling, respectively. Finally, one of the two Cartesian images embodies the rotation (and scaling) compensated using the estimations, and the translation between the pair of images is estimated by phase-correlating them, where the peak from the correlation represents the translation between the pair.

In [23] the authors used the aforementioned rationale for performing scan matching of sonar data. In that work, since scaling registration was not relevant, the author used the variant of the algorithm which converts the magnitude spectra to polar representation. Moreover, that research group expanded their algorithm for 3D mapping in [24]. Furthermore, in [25] it is argued that it might be worth to use directly a polar image, i.e. range-azimuth image, for estimating the rotation between the two images. According to this work, when the translation is relatively small, the error caused by the coupling between translation and rotation in the spatial domain might be small enough, so that the motion estimation can be done as follows: estimate the rotation between the pair by phase-correlating the image pair in the polar domain, next, convert the polar images to Cartesian representation. Finally, compensate the rotation on the Cartesian domain using the estimated rotation followed by phase-correlating the Cartesian images in order to estimate the translation between the pair.

B. Scanning Radar Odometry

In radar odometry we can classify the odometry approaches between feature-based and direct [26], [27]. Further, we can also group the radar odometry approaches based on the type of radar that is used, i.e. single chip or scanning radar. Because the scope of this paper is on scanning radar, on this section we will discuss the related work with this kind of radar.

Feature-based odometry approaches utilizing a scanning radar can be traced back to Cen et al. [5], [6]. Cen et al. [5] algorithm consisted on finding features in each azimuth bin by finding the peaks in the 1D signal combining low- and high-frequency components of that signal and a scan-matching approach tailored to radar data was also included. Kung et al. [15] developed an algorithm based on the Normal Distribution Transform along with a filter for radar images which pixels below a threshold were not considered feature candidates. Moreover, Hong et al. in [13] used SURF features [28] for radar odometry and SLAM, and in [14] introduced an odometry and SLAM approach based on radar data which a blob feature detector was developed which is based on a Hessian matrix for finding keypoint candidates and these are further filtered by observing their homogeneous spatial distribution. Moreover, in [14] a graph-based outlier rejection for keypoints was proposed.

Additionally, Burnett et al. developed the MC-RANSAC algorithm by extending the work of Cen et al. through exploiting Random Sample Consensus in [29] for data association as well as implementing a motion-distortion correction. In addition, Burnett et al. proposed HERO in [30] using a deep neural-network, unsupervised learning and Exactly Sparse Gaussian Variation Inference for finding keypoints and estimating the relative pose. Another approach using neural networks and supervised learning was proposed by Barnes et al. [7]. Furthermore, Aldera et al. proposed in [8] a mechanism to detect failures patterns in the feature matching step and in [9] a smooth-curvature constraint based on Ackermann drive model was exploited. Recently, Hyungtae et al. presented the ORORA algorithm which uses anisotropic component-wise translation estimation and graduated non-convexity based rotation estimation to estimate the rotation and translation separately. Finally, Adolfsson et al. [10]–[12] developed the CFEAR algorithm which is the current state of the art radar odometry algorithm. In [12] the latest versions of CFEAR, namely CFEAR-3 and CFEAR-3-s50, use the k-strongest filtering for pre-filtering the radar image as well as motion compensation and surface point distance minimization in the registration step.

The pivotal work in direct methods for odometry estimation using scanning radar data can be traced back to Checchin

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Intelligent Vehicles. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TIV.2023.3324941

et al. [19], where the authors followed the procedure previously explained in [22] in order to estimate the position and orientation of the agent; the relative motion estimation was further used for performing Simultaneous Localization and Mapping (SLAM) using radar data. Moreover, Park et al. [31] proposed a coarse to fine strategy for estimating the motion. Also, a keyframe selection approach was employed as well as pose graph optimization, resulting in a trajectory with pose correction as well as a map of the traversed region.

Additionally, still within the scope of direct methods, Barnes et al. [21] deployed a supervised learning approach which employed a brute-force approach to find the pose between frames using correlative scan matching along with a filter used to reduce the effect of, e.g. noise and occlusion. Following the work from [21], Weston et al. [32] proposed a faster algorithm based on [21] by exploiting the Fourier-based registration. For further discussion in the recent progress in radar odometry, we refer the reader to [27], [33].

It can be noticed that most of the literature for radar odometry has concentrated its efforts in developing feature-based methods which can be justified because of, e.g. reduced computation requirements with a sparse number of points/features. Nonetheless, a sparse set of features does not use all the information available in the radar image (similar to what happens in the mapping domain when one has to choose between landmark and volumetric mapping). Moreover, its computational performance is proportional to the amount of features, leading to a trade-off between computational performance and algorithm robustness. On the other hand, in our paper we show that exploiting the analytical characteristics of the direct approach, e.g. the decoupling between rotation and translation, along with the Radon transform and a robust framework for odometry estimation lead to results comparable with the state of the art radar odometry algorithm and better than all the previously documented direct radar odometry approaches.

III. Theory

In this section the fundamental theoretical aspects used in order to develop R^{3}O are explained.

A. Phase-Correlation

1) Introduction: The Generalized Cross-Correlation is frequently used for estimating shifts between two 1D or 2D signals [34], [35]. Among the numerous filters that can be used in a Generalized Cross-Correlation framework, the Phase-Transform filter is one of them. When such a filter is employed, the correlation method is named Phase-Correlation (PC).

Consider two signals $x_1(m)$ and $x_2(m)$ such that $x_1, x_2 \in \mathbb{R}$ and m is the independent variable in spatial domain. Let x_2 be a shifted copy of x_1, shifted by Δ_m. By using the shift property of the FT, we have that the spectra of the two signals will be related by $X_2(u) = e^{-j2\pi u \Delta_m} X_1(u)$, where u is the independent variable in the frequency domain, and $X_1(u)$ and $X_2(u)$ are the spectra of $x_1(m)$ and $x_2(m)$, respectively. The cross-spectrum is defined as

$$G(u) = X_2(u) X_1^*(u) = e^{-j2\pi u \Delta_m} |X_1(u)|^2 \tag{1}$$

where $(\cdot)^*$ is the conjugate of $(\cdot)$. Finally, their phase-normalized cross-spectrum $C(u)$ can be found by dividing the cross-spectrum by its magnitude. Therefore, $C(u)$ can be written as

$$C(u) = \frac{X_2(u) X_1^*(u)}{|X_2(u) X_1^*(u)|} = e^{-j2\pi u \Delta_m} \tag{2}$$

The next step for finding the shift consists of using the inverse Fourier transform (IFT) of $C(u)$ which results in the phase-correlation vector denoted by Υ, or, by extending to two dimensions, the phase-correlation matrix denoted by Ξ. To conclude, the shift Δ_m is found by locating the highest peak in Υ or Ξ.

Choosing the Phase-Transform filter for the Generalized Cross-Correlation, i.e., the inverse of the magnitude of the cross-spectrum as shown in (2), can be justified by several reasons. For example, the peak resulting from the IFT of (2) is theoretically a unit impulse function in comparison to the usually broader peak resulting from IFT of the non-filtered Generalized Cross-Correlation. Moreover, the PC gives more emphasis to the location, e.g. spatial structures, than energy [36]. Besides, in [37], the PC is evaluated as a robust correlation method.

2) Enhancements to PC: Even though the PC may offer advantages when compared to other Generalized Cross-Correlation options, it is important to note that there are also disadvantages within this approach. For example, while PC has advantageous performance in the case of frequency-dependent noise, it has degraded performance with frequency-independent white noise [38]. Nonetheless, by applying some signal processing techniques, the drawbacks can be mitigated. In [39] it is suggested that limiting the bandwidth of $C(u)$ leads to accomplishing better results. The justification is that most of the artifacts that reduce the performance of the PC are present in the higher frequency components and thus employing a lowpass filter in the $C(u)$ improves PC results. In this paper, we will make use of an ideal lowpass filter designed in the frequency domain with a cutoff $\kappa = 0.3$ and $\kappa = 0.6$ for the PC used for rotation and translation estimation respectively, which implies that 70% and 40% of the higher-frequency components are rejected by the filter respectively.

3) Subpixel Estimation: In many applications, the estimation restricted to integer pixel positions using the PC demonstrated in section III-A1 is not enough. As a consequence, a multitude of subpixel methods for the PC have been developed in literature. These methods diverge in many aspects. For example, the subpixel method might estimate the subpixel shift by using $C(u)$ in the frequency domain or the Ξ in spatial domain. The reader may find an up to date review of subpixel methods in [38]. In the present study, the authors have employed the subpixel method explained in [40] because of its relatively high performance reported in [38], [39].

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Intelligent Vehicles. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TIV.2023.3324941

B. Fourier-Based Registration

1) Overview: Fourier-based registration is an efficient registration approach which is under the family of area-based methods [38]. The original algorithm once proposed in [22] consisted of registering images in which there was scaling, rotation and translation between the pair. However, for the purpose of this paper, the scaling estimation is not relevant given that the radar images generated by a scanning radar can be seen as a top view 2D image around the vehicle where the height of the radar is fixed. Consequently, the theoretical foundations for Fourier-based registration will be explained only with regards to rotation and translation registration.

2) Polar Method: The polar method is the straightforward approach for calculating the rotation and translation between two images. Let the reference image represented in Cartesian coordinates be called $I^r(x,y)$. The moved image $I^m(x,y)$ evolves from $I^r(x,y)$ by applying a clockwise rotation by ϕ and a translation by (p,q). Hence, $I^m(x,y)$ can be written as

$$\begin{aligned} I^m(x,y) = I^r(&x\cdot\cos(\phi)+y\cdot\sin(\phi)-p,\\ &-x\cdot\sin(\phi)+y\cdot\cos(\phi)-q) \end{aligned} \tag{3}$$

Let then the FT of $I^r(x,y)$ be called $F^r(u,v)$ and the FT of $I^m(x,y)$ be called $F^m(u,v)$. By using the translation and rotation property of the FT, the two spectra can be associated by

$$\begin{aligned} F^m(u,v) = F^r(&u\cdot\cos(\phi)+v\cdot\sin(\phi),\\ &-u\cdot\sin(\phi)+v\cdot\cos(\phi))\cdot e^{-j2\pi(u\cdot p+v\cdot q)} \end{aligned} \tag{4}$$

Therefore, in order to find the rotation between the reference and moving image, one can represent both magnitude spectra in polar coordinates, and finally the estimation of the rotation $\hat{\phi}$ can be achieved by using the PC, where the peak in the angle direction of the Ξ will represent the rotation between the pair. Finally, one compensates the rotation by rotating one of the images using $\hat{\phi}$ and the translation component can be found by using the PC between the images in Cartesian coordinates, where the peaks of this Ξ will represent the estimation of the translation: $(\hat{p},\hat{q})$.

3) Radon Method: Even though Fourier-based registration using the polar (or log-polar) representation for estimation of rotation is a straightforward option, it might not be the one which will lead to the best estimation of ϕ. The problems mostly arise from the discrete nature of images and their respectively discrete magnitude spectra which have shown to be particularly detrimental to the performance of the straightforward option. For example, because values have to be interpolated onto different grids, larger interpolation errors are typically present especially in the Cartesian to log-polar conversion which has led to the design of filters in order to mitigate these effects [22], [38]. Motivated by that, research has been made in order to look for alternatives to the polar representation option, and one that has shown positive results is to use the Radon transform for obtaining the rotation estimation [41]. The intuition behind the improvement in performance when compared to the straightforward option, is that the Radon transform alleviates the interpolation errors because each "pixel" in the Radon domain comes from a summation of pixels which represent an exact line in the magnitude spectra image while in the (log)-polar representation the conversion is done pixel to pixel [41] on interpolated grid positions. Effectively, this implies that elements affected by local interpolation errors have smaller influence because these local errors are amalgamated under the summation.

The Radon transform $R(\rho,\theta)$ of a continuous image $f(x,y)$ is defined as

$$R(\rho,\theta) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(x,y)\delta(x\cos(\theta)+y\sin(\theta)-\rho)\,\mathrm{d}x\mathrm{d}y \tag{5}$$

with $-\infty<\rho<\infty$ and $0\leq\theta<\pi$ and where δ is the impulse function. It should be noted that in the discrete case finite sums substitute the integrals, and ρ becomes finitely bounded according to the image size. By inspecting the definition of the Radon transform, the aforementioned property is clear: Every point defined by an angle θ and a displacement ρ in the Radon domain is constituted by a line integral (or summing along a line in the discrete case) over the original image. Moreover, a key property of the Radon transform is that if an image is rotated by ϕ, its Radon transform will be shifted by ϕ, leading to $R(\rho,\theta+\phi)$ [42]. By using this property, one can then convert the magnitude spectra of the pair of moving and reference image, earlier defined as $I^m(x,y)$ and $I^r(x,y)$, to the Radon domain, and then the rotation can be estimated by phase-correlating the two Radon domain images in which the peak in the θ dimension represents the estimated rotation $\hat{\phi}$. After that, analogous to the polar method, one can then compensate the rotation on the Cartesian images in the frequency or spatial domain and use the PC to estimate the translation.

C. Outlier Removal

Outliers are observations which differ from the normal behaviour of the rest of the observations. The Median Absolute Deviation (MAD) is a viable and often attractive option for identifying outliers. That might be explained by the simplicity and effectiveness of the algorithm. For a data set Y with elements Y_m, $m\in[1,\ldots,M]$ the MAD of the data set, denoted by $\mathrm{MAD}(Y)$, is defined as

$$\mathrm{MAD}(Y) = M\big(|Y-M(Y)|\big) \tag{6}$$

where M is the sample median, and according to [43], [44] the standard deviation of normally distributed data using the MAD can be estimated as

$$\hat{\sigma}_{\mathrm{MAD}}(Y) = 1.4826\cdot\mathrm{MAD}(Y) \tag{7}$$

Then, the following evaluation can be done to find outliers:

$$\left|\frac{Y_m-M(Y)}{\hat{\sigma}_{\mathrm{MAD}}(Y)}\right| > \zeta \tag{8}$$

where ζ is a threshold value, and Y_m is an outlier when (8) is evaluated as true, i.e. the left hand side is greater than ζ.

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Intelligent Vehicles. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TIV.2023.3324941

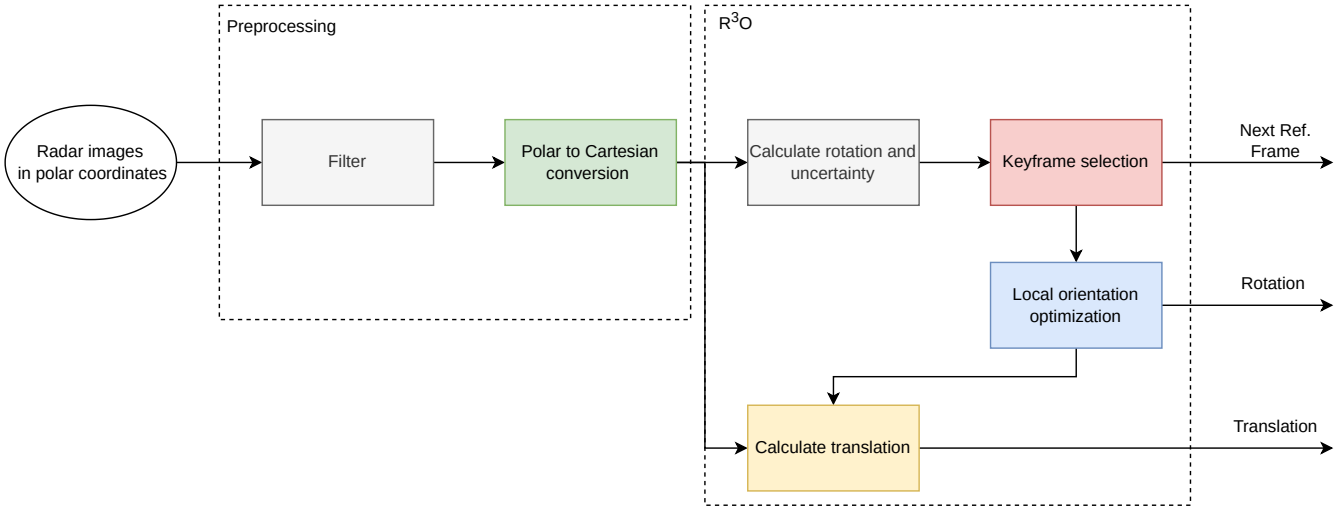

Fig. 1. Diagram of the R^{3}O algorithm. By looking at the diagram is possible to observe the importance of a good estimation of rotation is not only on the estimation of that parameter itself, but also on the translation estimation.

D. Tukey-biweight

The final component of this section consists of the Tukey biweight for location and scale estimation. This estimator is under a family of estimators called M-estimators. These estimators aim to attain efficiency, robustness and resistance [45]. The tukey biweight location estimator, denoted by μ_{tukey}, of a set Z with elements Z_i, $i \in [1, \ldots, N]$ is defined as

$$\mu_{\text{tukey}}(Z) = M(Z) + \frac{\sum_{|u_i|<1} (Z_i - M(Z))(1-u_i^2)^2}{\sum_{|u_i|<1} (1-u_i^2)^2} \quad (9)$$

where M is the sample median and u_i is given by:

$$u_i = \frac{Z_i - M(Z)}{c \cdot \text{MAD}(Z)} \quad (10)$$

where c is a constant that must be chosen. In all our location estimations, we have set c equal to 6. Furthermore, the Tukey-biweight scale estimator, denoted by σ_{tukey}, is defined as

$$\sigma_{\text{tukey}}(Z) = \sqrt{n}\, \frac{\sqrt{\sum_{|u_i|<1} (Z_i - M)^2 (1-u_i^2)^4}}{|(\sum_{|u_i|<1} (1-u_i^2)(1-5u_i^2))|} \quad (11)$$

where u is defined in (10). We used c equal to 9 for all scale estimations.

IV. Proposed Method

A. Introduction

R^{3}O employs a careful method for estimating the rotation. First, this is inspirited by having in mind that usually radar sensors have lower angular resolution when compared to other sensors, e.g. lidar, which demands exceptional performance of the registration algorithm. Second, this is motivated by noting that the rotation estimation is a cardinal step for registering the translation in the Fourier-based registration. This can be considered true, because a poor rotation estimation does not only imply a bad estimation of that variable on its own, but also affects the estimation of the translation.

The proposed method takes advantage of the previously documented improvement in performance of rotation estimation by using the Radon domain instead of the (log)-polar representation [41] along with a sequence of steps that are performed for finding $\hat{\phi}$ (see Fig. 1 for a diagram of how the relative motion is estimated by using R^{3}O). Note that we will assume that the radar is placed looking forward and the radar Cartesian image has its center collocated with the radar's center.

The robustness of the algorithm can be primarily explained by the use of multiple candidates for the rotation estimation using Radon domain slices. The use of different displacements can be related to the work of [46] in the sense that the multiple displacements can be understood to be resemblant to different high pass filters, e.g., the higher the displacement the less low frequency components are going to be used, applied to the magnitude spectra, thus reducing the effects of rotationally dependent aliasing. This kind of aliasing is caused by artifacts that are introduced due to the discretization of the FT, and its effect can deteriorate the performance of the rotation estimation using the Fourier-based image registration. For example, [46] has shown that this aliasing component reduces the height of the peak corresponding to the correct rotation. We noted in our evaluation that a selection of ρ values which included displacements very close or equal to zero along with displacements with relatively large magnitude (up to 66% of maximum absolute displacement) had the most positive impact in the algorithm performance. Nonetheless, note that the larger number of displacement selected, the higher is the computational requirements.

Moreover, the use of outlier removal can be thought as a processing step which remove divergent estimations, e.g. caused by interpolation errors and noise. Furthermore, the Tukey-biweight location estimator, which is a robust measure of central tendency, is used for the estimation of rotation between two frames. This process is repeated over a window of frames which keyframe selection is performed followed by

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Intelligent Vehicles. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TIV.2023.3324941

graph-based local orientation optimization which refines the rotation estimations. Finally, the optimized rotation between frames is used for de-rotating images in the Cartesian spatial coordinates and the relative translation is found by maximizing the phase-correlation between each pair.

Note that is possible to restrict the range of displacement values to choose in half due to the symmetry of the magnitude spectra of real images and one of property of the Radon transform, concretely that $R(\rho, \theta) = R(-\rho, \theta \pm 180°)$. Also, we note that it is possible to achieve similar results by removing the MAD outlier removal along with adjusting the constant values in the Tukey-biweight location and scale estimators. However, we decided to keep the outlier removal and location/scale estimation as different modules, so that other alternatives for location/scale estimation as well as other outlier removal approaches can be straightforwardly tested with the current algorithm's pipeline.

B. Preprocessing

Scanning radar images are contaminated by a series of unwanted artifacts, e.g. multipath, ghost targets and speckle noise, which may reduce the performance of odometry algorithms [5]. Therefore, before performing the odometry estimation, it is sensible to try to reduce the influence of unwanted components. In our algorithm, we have included as a preprocessing step an adaptation of the k-strongest filter [12] tailored for direct based odometry. In this filter, for every azimuth bin on the radar polar image, the filter sets to zero all range bins that have a power return smallest than the k-th greatest power return in that azimuth bin, where k is a parameter. Moreover a threshold, denoted by z_{min}, is applied to filter out, i.e. in our implementation this means set their pixel intensity to zero, all the pixels in the range-polar image that have power smaller than a threshold. In [12] it was shown that this filtering method achieves better performance with scanning radar than the constant false alarm (CFAR) filter. Nonetheless, we differ from [12] because we employ afterwards a direct method instead of a feature-based method. Consequently, we may choose a larger number for the parameter k as well as a smaller threshold for the smallest power return, i.e. we keep a larger number of pixels/targets. The advantage of that is that while in [12] it was necessary to select the parameters of the filtering algorithm leveraging a trade-off between robustness of the motion estimation against computational performance (it was shown in [12] that a less conservative filtering leads to high robustness while a more conservative filtering leads to higher computation performance), in our approach we can select the parameters solely based on the odometry performance. Finally, after filtering, the polar image is interpolated to Cartesian coordinates using bilinear interpolation.

C. R^3O

1) Rotation estimation: For every pair of radar images which the relative motion is estimated, their relative-rotation and the uncertainty of the estimation are calculated. In order to

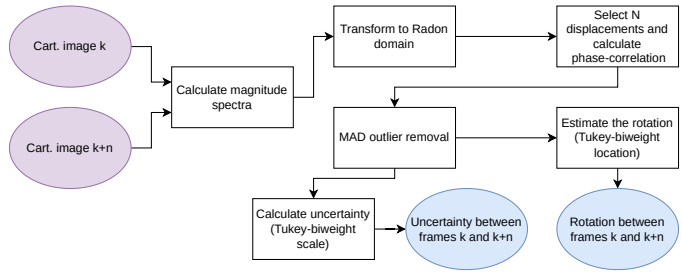

Fig. 2. Rotation estimation and its uncertainty estimation.

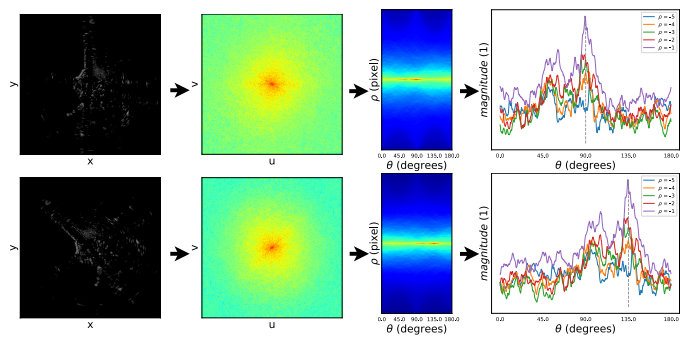

Fig. 3. Example of how a relative rotation between two Cartesian images becomes a shift in Radon domain. In this example we manually rotated one Cartesian radar image by $45°$. From left to right we have: (1) the images in Cartesian coordinates; (2) the magnitude spectra of the images in dB (we plot in dB to facilitate the visualization, but we use the linear scale for calculations); (3) the Radon domain images of the magnitude spectra (note that the $\rho = 0$ coincides with the middle/center row, and the rows above and below the center are from negative and positive displacement respectively); (4) magnitude vs. angle plots for some selected displacements.

explain the procedure for the rotation estimation, it is sensible to define the phase-correlation operator, denoted by $\mathcal{P}$, as

$$\mathcal{P}(s_1(\alpha), s_2(\alpha)) = \underset{\alpha}{\operatorname{argmax}}\, \mathcal{F}^{-1}\left(\frac{\mathcal{F}(s_1)\mathcal{F}^*(s_2)}{|\mathcal{F}(s_1)\mathcal{F}^*(s_2)|}\right) \quad (12)$$

where $\mathcal{F}$ and $\mathcal{F}^{-1}$ are the FT and IFT, respectively, s_1 and s_2 are 1D or 2D signals for which $\mathcal{P}$ is being calculated, and α represents the arguments of the signals which can be either a tuple of 2 arguments or a single argument corresponding to 2D or 1D signals, respectively. Consequently, $\mathcal{P}$'s output will be either 2D or 1D depending on the dimension of the evaluated signals.

See Fig. 2 for an overview of how the rotation estimation estimation and its uncertainty is calculated (the methods described in Section III-B would consist of applying a 2D phase-correlation after the transformation to Radon domain step or to use the polar transform replacing the Radon transform and them calculating the 2D phase-correlation). Having a moving and reference image in Cartesian coordinates, denoted by I^m and I^r respectively, the first step consists of calculating the FT of each image and the output of this step is the magnitude of each spectrum. Next, the magnitude spectra are converted to the Radon domain, let $R^r(\rho, \theta)$ and $R^m(\rho, \theta)$ be the Radon domain image of the reference image's magnitude spectrum and moving image's magnitude spectrum, respectively. Afterwards, instead of performing the 2D $\mathcal{P}(R^r(\rho, \theta), R^m(\rho, \theta))$ for finding the estimated rotation, R^{3}O calculates N 1D estimations of the

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Intelligent Vehicles. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TIV.2023.3324941

rotation by selecting N slices of the Radon domain images. This is accomplished by noting that at a fixed ρ value the 2D PC becomes 1D. For example, when ρ is fixed to a selected displacement ρ_k, the 1D PC is calculated by

$$\mathcal{P}\big(R^r(\rho_k,\theta), R^m(\rho_k,\theta)\big) = \mathcal{P}\big(R^r_{\rho_k}(\theta), R^m_{\rho_k}(\theta)\big) \tag{13}$$

where we use the subscript ρ_k on each Radon domain image to indicate the slice with all the θ values of a selected ρ_k displacement. Therefore, by defining a set W as the set with elements W_k, $k \in [1,\ldots,N]$ as the initial set of candidates for rotation estimation, we have that every candidate is calculated by finding the phase-correlation between the reference and moving Radon domain images:

$$W_k = \mathcal{P}\big(R^r_{\rho_k}(\theta), R^m_{\rho_k}(\theta)\big) \tag{14}$$

where ρ_k is the selected ρ value coming from a list of selected displacements which has length N. Fig. 3 shows an illustration of how a rotation in the Cartesian domain becomes a shift in the Radon domain.

Succeeding that, an outlier removal step comes into place in which we have selected the outlier removal based on the MAD as explained in section III-C for removing anomalous candidates. The surviving candidates are defined as the set X with elements X_k, $k \in [1,\ldots,J]$, $J \leq N$.

Finally, the rotation between the frames is estimated by

$$\bar{\phi} = \mu_{\text{tukey}}(X) \tag{15}$$

Furthermore, we estimate the uncertainty of the rotation estimation by calculating the Tukey-biweight scale estimation on the final set of candidates X:

$$\bar{\sigma} = \sigma_{\text{tukey}}(X) \tag{16}$$

2) Keyframe Selection: The keyframe selection step aims to dismiss radar images which generated high uncertainty in rotation estimation with regards to the current reference frame as well as to find the reference frame for the next iteration. Given the current reference frame, denoted by F_0, the algorithm estimates the relative rotation and its uncertainty with regards to the frames F_n where $n \in [1,2,\ldots,L]$ with L being the length of the window. Each estimated rotation is denoted by $\bar{\phi}_{0n}$ and its uncertainty is denoted by $\bar{\sigma}_{0n}$ where $n \in [1,2,\ldots,L]$. Next, in order to filter the estimations with high uncertainty, a fixed and an adaptive threshold are used. The fixed threshold, denoted by γ_{fixed}, removes all the frames F_n where the corresponding $\bar{\sigma}_{0n}$ was greater than the threshold value. In addition to it, an adaptive threshold is used. This is calculated by

$$\gamma_{\text{apt}} = \min(\bar{\sigma}_{0n}) \cdot \beta_{\text{apt}} \tag{17}$$

where β_{apt} is a constant with $\beta_{\text{apt}} \geq 1$ and $\min(\bar{\sigma}_{0n})$ is the smallest uncertainty in the current window. Analogous to the fixed threshold, the adaptive threshold removes all the frames F_n where the $\bar{\sigma}_{0n}$ was greater than γ_{apt}. See Fig. 4 for a graphical representation of the keyframe selection step as well as for an example of how the relative rotations' candidates between F_0 and F_n can be distributed after the MAD outlier removal. In this example, frames 4 and 6 were removed because their uncertainties were greater than the γ_{apt}.

The effect of the adaptive threshold is to avoid frames in the current window which had a considerable greater uncertainty than the smallest possible uncertainty in this window. On the other hand, the fixed threshold assures that no frame with uncertainty greater than what is determined to be acceptable can be taken (this could come into place in the case that the smallest estimated uncertainty is still high, then frames with unacceptably high uncertainty could be kept). In the case where all the frames' uncertainties are greater than the fixed threshold, the frame with the smallest uncertainty is the only one kept and selected.

Finally, the surviving frame which is the furthest in time from the current reference is selected as the reference frame for the next iteration (see example in Fig. 4).

3) Local Orientation Optimization: We have used a graph-based optimization with a robust loss function in order to refine our local orientation estimation. For that we consider all the pair-wise combinations of the remaining frames after the keyframe selection step. The graph is formed by nodes which represents the orientation at frames F_n relative to the reference frame F_0. We denote the nodes as o_i where i has the same subscript of the surviving Frame F_n; and the vector with all the nodes is defined as $\boldsymbol{o} = (o_i)^T \ \forall i$. Between each node, there is a rotation factor with its uncertainty. These are equal to the relative rotation estimation and its uncertainty estimated following the steps shown in Section IV-C1. The relative rotation constraint and its uncertainty are denoted by $\bar{\phi}_{ij}$ and $\bar{\sigma}_{ij}$ respectively, where i is the subscript of the node where the edge starts and j is the subscript of the node where the edge ends. An example of the orientation graph is shown in 5.

The objective function of the local orientation graph-optimization is written as [47]–[49]:

$$\boldsymbol{F}(\boldsymbol{o}) = \sum_{<i,j>\in V} \rho_{\text{cauchy}}\left(\sqrt{e(o_i,o_j,\bar{\phi}_{ij})^T \Sigma_{ij}^{-1} e(o_i,o_j,\bar{\phi}_{ij})}\right) \tag{18}$$

where V is the set of pair of indices for which a constraint in $\boldsymbol{o}$ exists, e is the error function which evaluates how much the parameters satisfy the constraints, Σ_{ij} is the uncertainty of the measurement which is equal to $\bar{\sigma}_{ij}$, and ρ_{cauchy} is the Cauchy loss function defined by

$$\rho_{\text{cauchy}}(x) = \frac{c^2_{\text{cauchy}}}{2} \cdot ln\left(1 + \left(\frac{x}{c_{\text{cauchy}}}\right)^2\right) \tag{19}$$

where c_{cauchy} is a constant. Finally, it is evident that in order to get the relative rotation between each successive node requires only the calculation of the difference between their local orientation. For example, assuming an orientation vector $\boldsymbol{o} = (o_0, o_1, o_2)^T$, then the estimated rotation $\hat{\phi}$ between nodes o_1 and o_2 is $\hat{\phi} = o_2 - o_1$. In order to solve the graph, we use GTSAM[1] [49]

4) Global Orientation Estimation: Over the course, the orientation Φ_n (with a given initial orientation Φ_0) of an agent

[1]https://github.com/borglab/gtsam

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Intelligent Vehicles. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TIV.2023.3324941

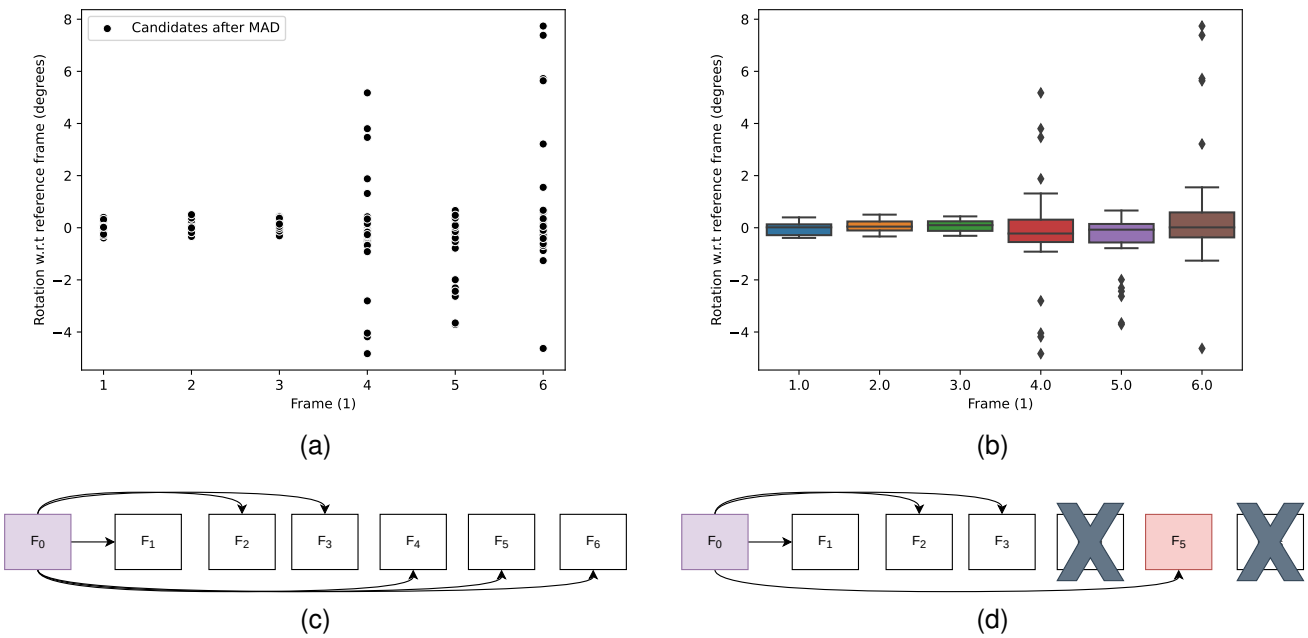

Fig. 4. Keyframe selection starts by estimating the rotation and the uncertainty between the reference frame and the frames in the current window. In this example, the uncertainties between F_0 and $F_{4,6}$ were greater than the defined threshold parameters. The reasoning for dismissing $F_{4,6}$ is elucidated by observing the boxplot (with the whiskers limits based on 1.5 times the interquartile range) where the larger dispersion on their estimations is clear. Therefore $F_{4,6}$ were not used for this local orientation estimation. Note that F_5 will be the reference frame in the next iteration. (a) Scatter-plot showing the distribution of the surviving candidates after the MAD outlier removal. (b) Boxplot of the surviving candidates after MAD outlier removal. (c) A graphical representation of the keyframe selection. (d) A graphical representation of the keyframe selection after the thresholds were applied.

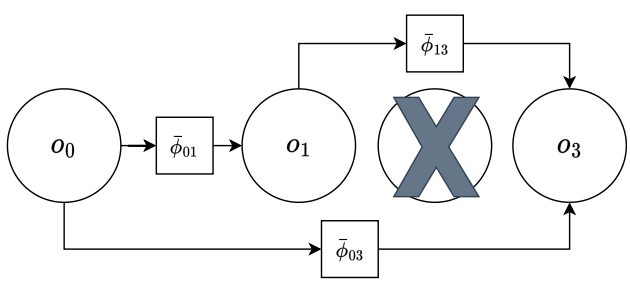

Fig. 5. Illustration of the orientation graph used in R^{3}O. In this example, the keyframe window has length 3 and the frame F_2 was pruned in the keyframe selection step. The nodes $[o_0, o_1, o_3]$ represent the local orientation in the iteration. Also, $[\bar{\phi}_{01}, \bar{\phi}_{03}, \bar{\phi}_{13}]$ are the rotation factors between each node.

with regards to an inertial frame at instance k, for $k > 0$, is then calculated by summing all the $\hat{\phi}$ till n.

$$\Phi_N = \Phi_0 + \sum_{n=1}^{N} \hat{\phi}_n \tag{20}$$

D. Position Estimation

After the local orientation is estimated, we use a sequence of the closest frames from the current reference frame to the next reference frame in order to calculate the translation in this window. For example, in the illustration shown in Fig 4, the translation would be computed between the following pairs: F_0-F_1,F_1-F_2,F_2-F_3,F_3-F_5. By using the appropriate $\hat{\phi}$ for each pair, the rotation is then compensated in the Cartesian representation and the translation is estimated by using the operator $\mathcal{P}$ on the Cartesian images. However, in order to get the position, it is necessary to convert the displacement from the local coordinates into the displacement with regards to the inertial frame. We assume that $\hat{p}$ and $\hat{q}$ at instance k represent the displacement between the instance k and $k-1$ in the vehicle's longitudinal and lateral directions, respectively. Therefore, the position (x_k, y_k) for $k > 0$ and for a given x_0 and y_0 can be found by

$$\begin{bmatrix} x_k \\ y_k \end{bmatrix} = \begin{bmatrix} \cos(\Phi_k) & -\sin(\Phi_k) \\ \sin(\Phi_k) & \cos(\Phi_k) \end{bmatrix} \begin{bmatrix} \hat{p_k} \\ \hat{q_k} \end{bmatrix} + \begin{bmatrix} x_{k-1} \\ y_{k-1} \end{bmatrix} \tag{21}$$

V. Evaluation

A. Introduction

The goals of this section are to evaluate and compare the performance of the proposed algorithm with: (1) the current state of the art radar-odometry algorithm [12] which reported its results with the Oxford [50] and Mulran [20] datasets; (2) the direct methods with best reported performance, namely [21] in the Oxford dataset and [31] in the Mulran dataset. Consequently, we have evaluated R^{3}O using Mulran and Oxford dataset.

In order to show that the algorithm did not have its parameters overly calibrated for one dataset or another, we kept all the parameters of the core part of the algorithm fixed. All in all, the only changes that happened between the datasets occurred in the prepocessing stage which is explained by noting that the two datasets use different scanning radars with, e.g. different maximum range. For all the PC operations, the previously mentioned subpixel method explained in section

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Intelligent Vehicles. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TIV.2023.3324941

TABLE I
PARAMETERS USED FOR ALGORITHM'S EVALUATION

Preproccessing	
Parameter	Value
k-strongest: k	80/200 (Oxf./Mulr.)
$z_{\min}$	25.5 (scale 0-255)
Cart. Resolution	0.319/0.39 (Oxf./Mulr.) pixel/m
R^3O	
Parameter	Value
κ (rotation)	0.3
κ (translation)	0.6
ζ	1.5
ρ	$[-54, -53, \ldots, 0]$
γ_{fixed}	$1.8°$
β_{apt}	3.5
Window length	6 frames
c_{cauchy}	0.01

III-A3 was used along with the enhancement described in section III-A2. All the parameters used in the evaluation can be seen in Table I. Over the datasets evaluated and with the parameters described, our algorithm was implemented in Python and tested in a computer with 2-Intel E5-2695 and 2-Nvidia Quadro RTX-4000 (using a second thread to compute all the steps after the keyframe selection run on a second thread) obtaining an output rate of approx. 4.57Hz which makes it suitable for online utilization given that both radars have an output rate of 4Hz. Finally, because the radar was attached to a car, we have used only the translation component in the longitudinal direction for the position update.

B. Oxford

The Oxford Radar dataset uses a vehicle equipped with an assortment of sensors including a Navtech CTS350-X scanning radar with range resolution of 4.38cm, 0.9 degree azimuth resolution, maximum range of 163m, and output rate of 4Hz. This dataset has a total of 32 traversal over approximately the same route in Oxford, UK under different weather, traffic and lighting conditions. Moreover, ground-truth is provided based on camera, GNSS, and inertial navigation system information [50]. For evaluating the performance of the algorithm, we have used the KITTI odometry evaluation [51] which consists of computing the errors in translation and rotation from length $[100m, 200m, \ldots, 800m]$, followed by averaging the errors over the errors obtained with different lengths. We used this metric because this was the metric used by other radar-odometry algorithms which assessed their performance using this dataset. Finally, among the 32 possible traversals, we selected the same 8 sequences used by [12], [14], [30].

A comparison between the trajectories estimated by our algorithm and the ground truth can be seen in Fig. 6. Moreover, in Table II a comparison between R^3O's performance with other radar-based odometry algorithms is shown. Among the direct based methods, R^3O obtained the best performance in terms of translation and rotation error. In addition, comparing with all algorithm (direct and feature based), R^3O had the second best performance in terms of rotation error.

C. Mulran

This dataset includes a variety of sequences which include, e.g. dynamic scenarios over urban areas. The radar used on the dataset was a Navtech CIR204-H which is a 360 degrees scanning radar with 0.9 degrees angular resolution and 0.06m range resolution. Also, this radar has a throughput frequency of 4 Hz and 200 meters maximum range. Moreover, the dataset provides a ground truth trajectory which is calculated by performing SLAM using a fiber optic gyroscope (FOG), Virtual Reference Station Global Positioning System, and iterative closest point (ICP) algorithm. More details of the dataset can be found in [20]. Among the sequences provided by the dataset, we have used the DCC, Riverside and KAIST sequences (we selected the same 9 sequences evaluated in [12], [14]) for evaluating our proposed algorithm. Riverside is a sequence with straight runs along a river and two bridges, DCC is a sequence which is structurally distinctive including a mountain and crossroads, and KAIST is in a campus environment.

For evaluating the performance of the algorithms in this dataset, in addition to the KITTI odometry metric, we have used the relative pose error (RPE) explained in [52], [53]. This was motivated by noting that both [12] and [31] used this metric in their evaluation using the Mulran dataset. The RPE consists of measuring the difference between the estimated and the true motion in a fixed time, frame or spatial interval following by averaging the errors. Because [12] and [31] use different intervals for their RPE evaluation, it was necessary to compute the RPE using different intervals: (1) for the comparison with [12], we computed the RPE in translation with an interval equal to 1 frame; (2) for the comparison with [31], we computed the RPE in rotation angle and the RPE in translation with 5 different spatial interval parameters which corresponded to 10%, 20%, 30%, 40%, 50% of the total trajectory length.

A comparison between the trajectories estimated by our algorithm in contrast with the ground truth can be seen in Fig. 7. In addition, in Table III a comparison between R^3O's performance with other radar-based odometry algorithms using the KITTI evaluation is shown. The results in this table show that the proposed algorithm achieved best performance in terms of rotation error and the second best performance in terms of translation error. Moreover, the results using the RPE with step equal to 1 frame are displayed in Table IV which shows that our algorithm overcame the state of the art [12] in terms of translation RPE.

Furthermore, we calculated the RPE with the same steps as Park et al. for the Riverside03 and DCC02 trajectories (which were the 2 trajectories from [20] used by [31]). The boxplots for the RPE with different segments size of the estimators are shown in Fig. 8. Finally, the comparison between the performance of R^3O and PhaRaO [31] is shown in Table V. It is evident by observing Table V that our algorithm achieved better performance than [31]. For example, in the DCC02 sequence, the proposed algorithm had 38% less translation error than [31].

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Intelligent Vehicles. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TIV.2023.3324941

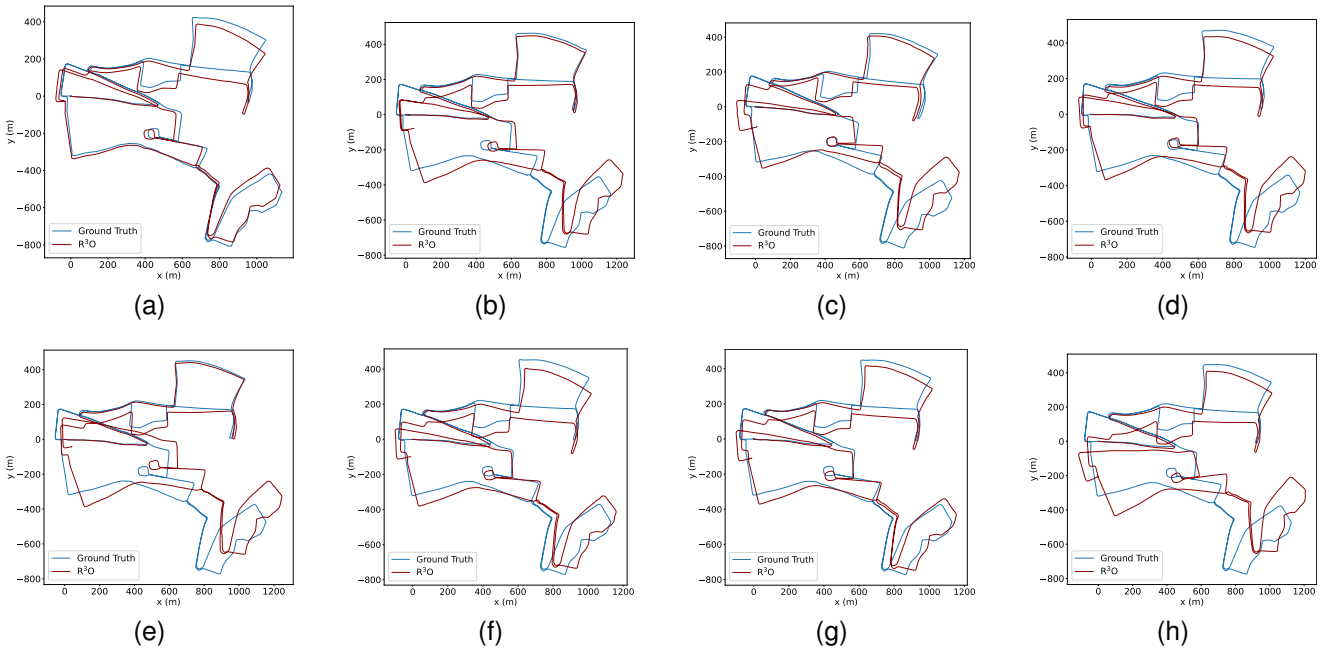

Fig. 6. Comparison between R^{3}O's estimated trajectory and true trajectory in the Oxford dataset. (a) 10-11-46, (b) 10-12-32, (c) 16-11-53, (d) 16-13-09, (e) 17-13-26 , (e) 18-14-14, (f) 18-14-46, (g) 18-15-20.

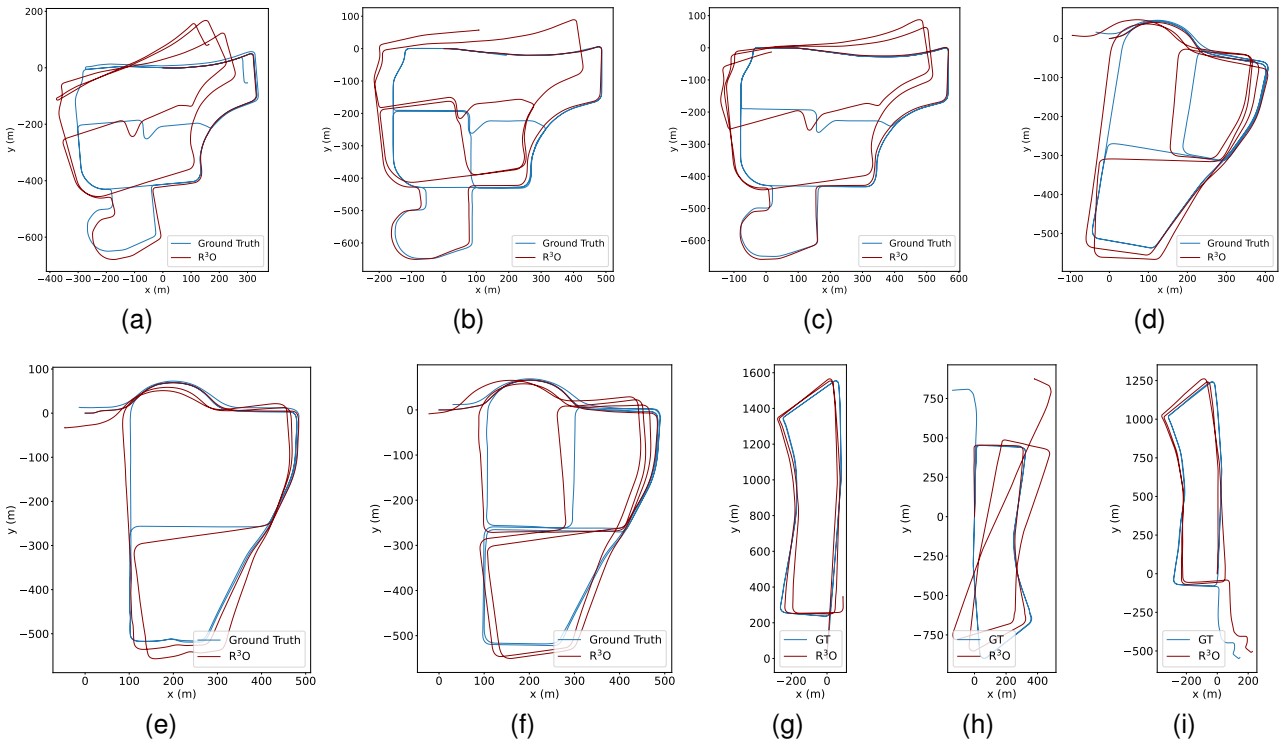

Fig. 7. Comparison between R^{3}O's estimated trajectory and the true trajectory in the Mulran dataset. (a) KAIST01, (b) KAIST02, (c) KAIST03, (d) DCC01, (e) DCC02, (f) DCC03, (g) Riverside01, (h) Riverside01, (i) Riverside03.

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Intelligent Vehicles. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TIV.2023.3324941

TABLE II
RESULTS WITH THE OXFORD DATASET

Algorithm	10-11-46	10-12-32	16-11-53	16-13-09	17-13-26	18-14-14	18-14-46	18-15-20	Mean
R^3O	2.02/0.57	1.89/0.56	2.17/0.60	1.95/0.60	2.09/0.58	2.06/0.61	1.96/0.57	2.18/0.65	**2.04/0.59**
MbyM-Dual Cart SCV [21]	-	-	-	-	-	-	-	-	2.784/0.85
F-MbyM [32]	-	-	-	-	-	-	-	-	2.06/0.63*
RadarSLAM-Odometry [14]	2.16/0.6	2.37/0.7	2.49/0.7	2.62/0.7	2.27/0.6	2.29/0.7	2.12/0.6	2.25/0.7	2.32/0.7
CFEAR-3 [12]	1.26/0.38	1.23/0.36	1.42/0.44	1.25/0.39	1.25/0.4	1.34/0.41	1.42/0.44	1.26/0.41	1.31/0.40
CFEAR-3-s50 [12]	**1.05/0.36**	**1.05/0.34**	**1.18/0.36**	**1.08/0.34**	**1.07/0.36**	**1.11/0.37**	**1.11/0.37**	**1.03/0.37**	**1.09/0.36**
HERO [30]	2.14/0.71	1.77/0.62	2.01/0.61	1.75/0.59	2.04/0.73	1.83/0.61	1.97/0.65	2.2/0.77	3.31/1.09
CC-Mean [9]	2.35/0.8	2.57/0.9	-	-	-	-	-	-	2.53/0.82
MC-RANSAC [29]	-	-	-	-	-	-	-	-	3.31/1.09
Robust Keypoint [7]	-	-	-	-	-	-	-	-	2.05/0.67*
Cen et al. [5]	-	-	-	-	-	-	-	-	3.72/0.95
Kung et al. [15]	-	-	-	-	-	-	-	-	1.96/0.6

We compare the performance of different radar-based odometry algorithm in the Oxford dataset. The first three methods are dense/direct methods while the other are indirect/feature-based algorithms. Results are displayed in (% translation error / deg/100m). Results marked with "*" are from supervised-learning algorithms which have only disclosed results without spatial-cross-validation. Note that Cen et al. and MC-RANSAC results were evaluated in [30].

TABLE III
RESULTS WITH THE MULRAN DATASET: KITTI ODOMETRY EVALUATION

Algorithm	Riv.01	Riv.02	Riv.03	KAIST01	KAIST02	KAIST03	DCC01	DCC02	DCC03	Mean
R^3O	**1.34/0.39**	1.98/0.53	1.81/0.57	1.89/**0.63**	1.55/**0.53**	**1.53/0.5**	2.39/**0.43**	1.40/**0.34**	1.48/0.41	1.70/**0.48**
RadarSLAM-Odometry [14]	2.04/0.5	1.51/**0.5**	1.71/0.5	2.13/0.7	2.07/0.6	1.99/**0.5**	2.7/**0.5**	1.9/0.4	1.64/**0.4**	1.97/0.5
CFEAR-3 [12]	1.59/0.63	1.39/0.51	1.41/0.4	1.59/0.66	1.62/0.66	1.73/0.78	2.28/0.54	1.49/0.46	1.47/0.48	1.62/0.57
CFEAR-3-s50 [12]	1.62/0.62	**1.35**/0.52	**1.19/0.37**	**1.48**/0.65	**1.51**/0.63	1.59/0.75	**2.09**/0.55	**1.38**/0.47	**1.26**/0.47	**1.5**/0.56
MC-RANSAC [29]	3.83/1.07	3.61/1.03	3.94/0.96	-	3.46/1	3.31/0.99	4.17/0.95	3.14/0.66	2.7/0.67	3.66/0.8
MC-RANSAC + DPLR [29]	4.12/1.13	3.93/1.11	3.81/0.92	-	3.59/1.08	3.60/1.05	4.01/0.89	2.87/0.55	2.7/0.66	3.44/0.72
ORORA [16]	3.53/0.84	2.67/0.64	2.11/0.49	-	3.28/0.82	3.04/0.7	3.12/0.67	2.60/0.51	2.37/0.57	3.0/0.67

We compare the performance of different radar-based odometry algorithm in the Mulran dataset. Results are displayed in (% translation error / deg/100m). The proposed algorithm is the only direct method in this table. Note that MC-RANSAC and MC-RANSAC + DPLR and ORORA are using features based on [6]. Moreover, MC-RANSAC and MC-RANSAC + DPLR results were evaluated in [16].

TABLE IV
RESULTS WITH THE MULRAN DATASET - RPE WITH STEP EQUAL TO 1 FRAME

Algorithm	Riv.01	Riv.02	Riv.03	KAIST01	KAIST02	KAIST03	DCC01	DCC02	DCC03	Mean
R^3O	6.54	**4.04**	4.67	**4.76**	**4.33**	**4.7**	**6.44**	**3.55**	**4.55**	**4.84**
CFEAR-3	**5.9**	5.38	**4.5**	6.37	6.01	6.21	7.83	4.54	5.16	5.77
CFEAR-3-s50	6.36	5.65	4.63	6.34	6.04	6.19	7.56	4.45	5.04	5.81

We compare our algorithm with [12] using the RPE metric. Results are displayed in centimeters.

TABLE V
RESULTS WITH THE MULRAN DATASET - RPE WITH STEPS USED IN [31]

Algorithm	Riv.03	DCC02
R^3O	**1.56/4.59**	**1.18/2.29**
PhaRaO [31]	2.36/4.66	1.91/2.34

Results are displayed in (% translation error / degrees).

VI. ANALYSIS

A. Introduction

In this section we evaluate the effect of each step of the algorithm in its overall performance in terms of rotation estimation. Along with that, we also compare our algorithm with Fourier-based registration using polar representation of magnitude spectra as well as by computing a 2D PC over the Radon 2D image of the spectra which we will refer as the Rad2D algorithm (the angle bins had a distance of approx. 0.35 degree for all algorithms). This is motivated on the ground that we can not only assess the improvements of R^3O compared to Rad2D and polar, but also the impact of each cardinal step in R^3O algorithm. Furthermore, we include the influence of the amount of ρ candidates in the algorithm's performance along with the rationale for appropriate selection of candidates in R^3O.

In this analysis, akin to [12], we have used the 32 sequences of the Oxford dataset which amounts to 240088 radar images. Moreover, the parameters of R^3O and preprocessing concur with Table I and Section V. Finally, exactly as we had in Section V, we have calculated the Radon Transform of magnitude spectra of size 1023x1023 using distance between displacements equal to one pixel. This implies a total number of 1023 possible ρ candidates, i.e.

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Intelligent Vehicles. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TIV.2023.3324941

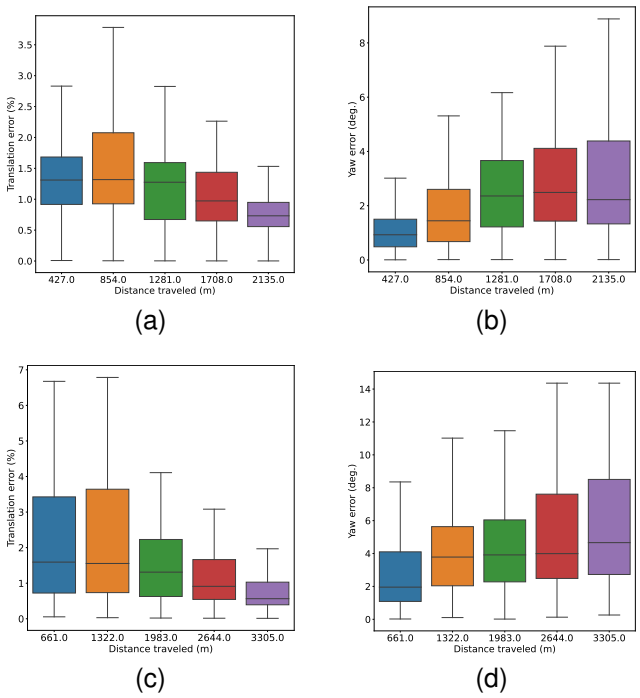

Fig. 8. Boxplots (with the whiskers limits based on 1.5 times the interquartile range) for the RPE using the same step size as [31]. (a) Translation-Riverside03. (b) Rotation-Riverside03. (c) Translation-DCC02. (d) Rotation-DCC02.

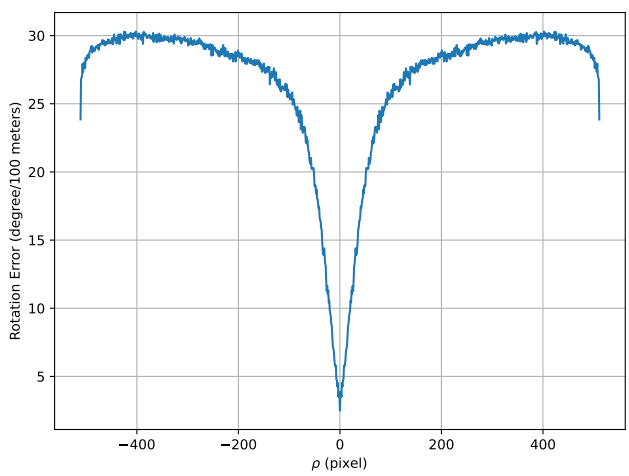

Fig. 9. The rotation error using a single candidate using the KITTI evaluation. It is possible to observe that candidates near the center have the smallest errors.

$$\rho = [-511, -510, \ldots, 0, \ldots, 510, 511] \ .$$

B. Analysis of Individual Candidates

In this analysis, we used only one single ρ for estimating the rotation, i.e. using the rotation estimate of each possible candidate before the outlier removal stage (see Fig. 2). In Fig. 9 the rotation error using the KITTI odometry metric averaged over all sequences' errors versus the displacements is shown. The symmetry around zero is a known property already alluded in Section IV-A; this implies that we can only use effectively half of the displacements. Furthermore, it is also visible in Fig. 9 that the candidates values near zero are those

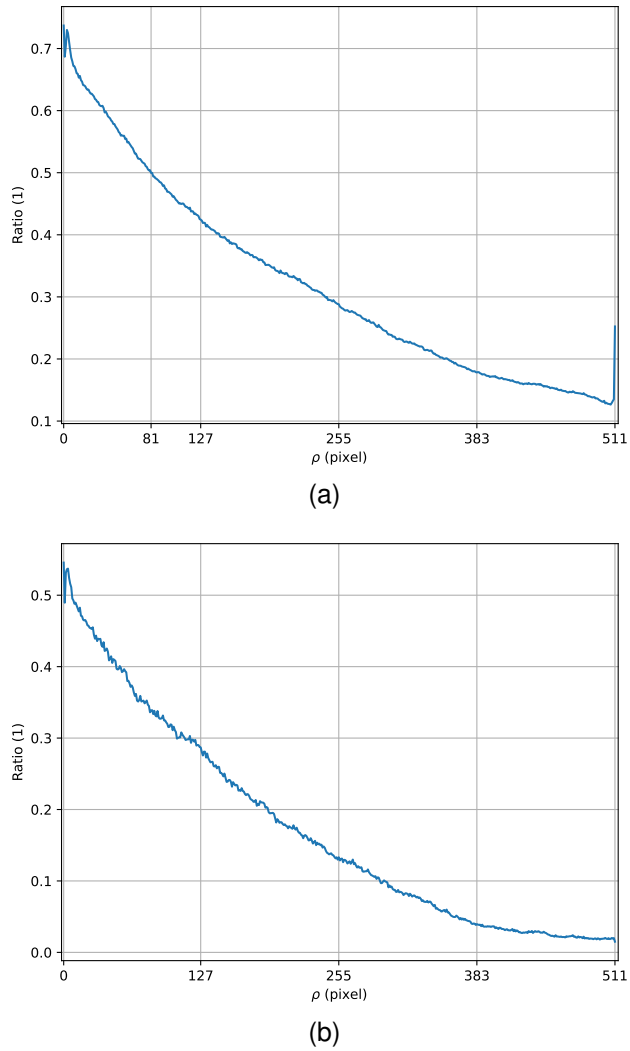

Fig. 10. The ratio which the single candidate rotation estimation was within defined boundary. (a) The ratio over all available samples. (b) The ratio when considering samples which at least 1 degree rotation occurred between the pair of radar images.

which yield the lowest rotation error. Since the performance ρ candidates is symmetrical, for simplicity, from now on we will show the results only using one of the halves of candidates.

Moreover, we performed a complementary evaluation in order to narrow down the candidates which are meaningful for attaining the correct rotation estimation. For that, using the RPE with step equal to 1 frame, we estimated the ratio/frequency which each candidate estimates a rotation with an error within ± 1 angle bin. Therefore, for each RPE estimation, we evaluate whether the magnitude of the error was smaller than 0.35 degree. Finally, we divide the sum of occurrences of all the errors within the acceptable boundary by the total amount of error evaluations. In Fig. 10a the results of this experiment can be seen; it is noticeable that already at displacement equal to 81, the ratio is below 0.5, also as the displacement increases the ratio decreases. Moreover, by observing Figs. 9 and 10a, we can observe that at the latest displacement the error goes down and the ratio increases respectively. This can be explained by noticing that at far-

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Intelligent Vehicles. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TIV.2023.3324941

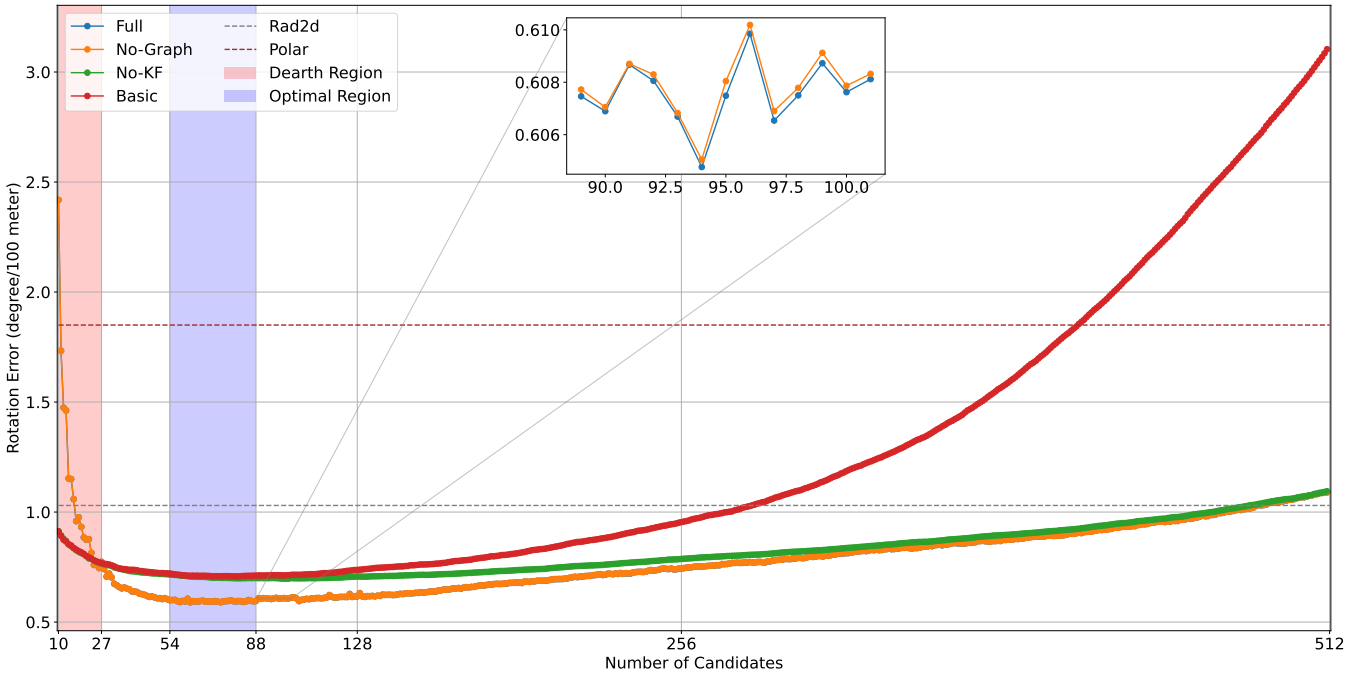

Fig. 11. Ablation study. It is noticeable the effects of each stage of the algorithm in its overall performance as well as that the effects of each stage are directly linked with the number of candidates which are evaluated.

TABLE VI
RESULTS OF THE ABLATION STUDY

Algorithm	Dearth Region	Optimal Region	Overall
Full	1.1009	**0.5952**	**0.7875**
No-Graph	1.1010	0.5953	0.7877
No-KF	**0.8234**	0.7045	0.8229
Basic	0.8246	0.7108	1.2706

displacements the amount of pixels used in the sum for the Radon transform are few and correspond only to the highest frequency components of the magnitude spectra, thus the sum approaches zero, consequently very often the rotation estimation of pairs at the far-displacements are equal to zero. Therefore, when no rotation actually occurred, these far-displacements get the correct result. Fig 10b shows the ratio over the single displacements considering only the samples where there was at least 1 degree rotation between them. By juxtaposing the two figures of Fig. 10, it becomes evident the aforementioned rotation estimation characteristics at far-displacements.

C. Ablation Study

After acquiring the information of the individual performance in the preceding subsection, now we observe the accumulative effect of the candidates as well as the impact of each step in R^3O rotation estimation. For that, we evaluate the algorithm at its different stages with a different number of candidates. The candidates are selected by always starting from displacement zero till the number of candidates minus 1, i.e. when the number of candidates is N, we use the ρ values of $[0, 1, \ldots, \mathrm{N}-1]$. In this study we evaluated the algorithm with the number of candidates were $[10, 11, 12, \ldots, 512]$. Moreover, we evaluated the algorithm in 4 different stages: (1) Full: R^3O using all its stages; (2) No-Graph: which we do not use the graph optimization; (3) No-KF: which we not use the graph optimization as well as the keyframe selection; (4) Basic: No-KF without outlier removal. In this study, we used the rotation error estimation using KITTI odometry metric and averaged over all sequences' errors.

Fig. 11 shows the rotation error over the number of candidates including also the rotation error of polar and Rad2D. In Fig. 11 it is possible to note that:

- When few candidates are used, less than 27, the keyframe selection and consequently the graph optimization fails to improve the results. This can be understood by considering that the keyframe stage effectiveness requires that the uncertainty estimation to be reliable, which at low sampling numbers do not necessarily occurs. We denote this region as the Dearth Region.
- As the number of candidates increase, there is likewise an improvement in the rotation estimation until it reaches a plateau. We denote this plateau as the Optimal Region. Using the number of candidates within this region lead to the best performance.
- After the plateau, we notice that increasing the number of candidates do not lead to an improvement in performance. This can be understood with the aid of Fig. 9 and Fig. 10. We believe that at this point, there is a high number of candidates with estimations which are afar from the correct rotation. Consequently, there is no consistent improvement by using them. Moreover, as the number of outliers become more prevalent, it increases the chance that some will survive after the outlier removal which will

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Intelligent Vehicles. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TIV.2023.3324941

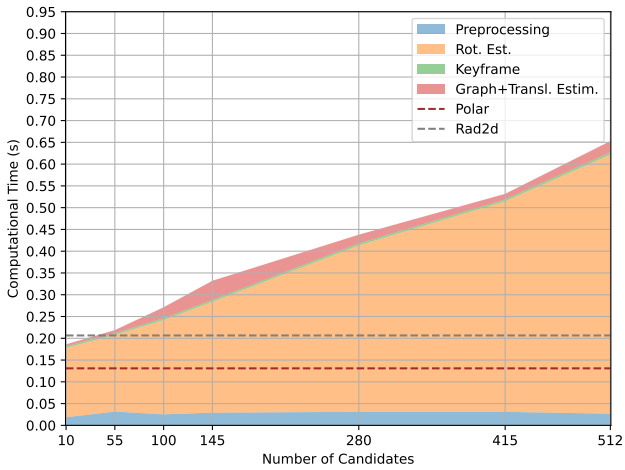

Fig. 12. R^3O's computational time. It is possible to observe that the rotation estimation (before the keyframe stage) have the highest impact in algorithm's computational time.

also deteriorate the performance of the keyframe selection and graph optimization.

- There is a relatively small but consistent improvement by using the graph optimization.
- After the Optimal Region, as the number of candidates increase, the importance of the outlier removal grows. It is noticeable that at very large number of candidates, the performance gap between the Basic and other versions increases.
- Rad2d performance is better than polar. Nonetheless, any version of R^3O in the Optimal Region outperforms Rad2d.

D. Computational Time

Finally, we compare the computational time of R^3O using different number of candidates with Rad2d and polar. For all the evaluations, we use the same machine described in Section V. In Fig. 12 the computational time is show. It can be seen that R^3O's computational time grows as the number of candidates increases. The preprocessing steps is almost constant through time while the rotation estimation (before the keyframe step) grows with the number of candidates. Moreover, the keyframe stage has constant and almost negligible effect on the algorithm computational time and the last stages in the algorithm (note that these stages run in a second thread which explains the non-linearity in time complexity, i.e. they can take shorter or longer time based on the resources which the second and the main thread demand concurrently) also have relatively small effect in the total computational time. Furthermore, we see that the polar-based algorithm has a lower computational time compared to Rad2d due to the additional time complexity of the Radon transform. Moreover, we notice that at the number of candidates used in Section V, R^3O has almost the same time complexity as Rad2d.

VII. Conclusion

A method for odometry estimation using radar sensors has been proposed in this paper. The theoretical foundations of the method and the details for its implementation were explained. Moreover, R^3O was compared to other direct and indirect odometry methods. It was shown that R^3O surpasses all other direct radar odometry methods, and it challenges the state of the art radar odometry algorithm in the Mulran dataset. Furthermore, by evaluating the algorithm across different datasets allowed us to assess that the proposed algorithm holds consistent performance when utilized with different scanning radars under various environments, lightning, traffic and weather conditions.

To conclude, we believe that the pipeline demonstrated in this paper, especially the generation of multiple hypotheses for the rotation estimation by exploiting the Radon transform, can represent a new cornerstone for direct odometry methods. The framework introduced in R^3O can be incorporated in previous approaches, e.g. [31], [32], as well as in a hybrid approach combining dense and feature-based techniques. Moreover, R^3O can be also used with other sensor modalities, e.g. 2D lidar, by adjusting the preprocessing step.

Acknowledgments

This work has been supported by the COMET-K2 Center of the Linz Center of Mechatronics (LCM) funded by the Austrian federal government and the federal state of Upper Austria.

Furthermore, we would like to thank the authors of [31] for their valuable feedback with regards to their paper.

References

[1] S. A. S. Mohamed, M.-H. Haghbayan, T. Westerlund, J. Heikkonen, H. Tenhunen, and J. Plosila, "A survey on odometry for autonomous navigation systems," *IEEE Access*, vol. 7, pp. 97 466–97 486, 2019.

[2] A. Chalvatzaras, I. Pratikakis, and A. A. Amanatiadis, "A survey on map-based localization techniques for autonomous vehicles," *IEEE Transactions on Intelligent Vehicles*, pp. 1–23, 2022.

[3] O. J. Woodman, "An introduction to inertial navigation," 2007.

[4] C. Waldschmidt, J. Hasch, and W. Menzel, "Automotive radar — from first efforts to future systems," *IEEE Journal of Microwaves*, vol. 1, no. 1, pp. 135–148, 2021.

[5] S. H. Cen and P. Newman, "Precise ego-motion estimation with millimeter-wave radar under diverse and challenging conditions," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 6045–6052.

[6] ——, "Radar-only ego-motion estimation in difficult settings via graph matching," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 298–304.

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

[8] R. Aldera, D. D. Martini, M. Gadd, and P. Newman, "What could go wrong? introspective radar odometry in challenging environments," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, 2019, pp. 2835–2842.

[9] 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*, vol. 7, no. 3, pp. 7865–7872, 2022.

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

[11] ——, "Oriented surface points for efficient and accurate radar odometry," 2021.

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Intelligent Vehicles. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TIV.2023.3324941

[12] ——, "Lidar-level localization with radar? the cfear approach to accurate, fast, and robust large-scale radar odometry in diverse environments," *IEEE Transactions on Robotics*, vol. 39, no. 2, pp. 1476–1495, 2023.

[13] Z. Hong, Y. Petillot, and S. Wang, "Radarslam: Radar based large-scale slam in all weathers," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 5164–5170.

[14] Z. Hong, Y. Petillot, A. Wallace, and S. Wang, "Radarslam: A robust simultaneous localization and mapping system for all weather conditions," *The International Journal of Robotics Research*, vol. 41, no. 5, pp. 519–542, 2022. [Online]. Available: https://doi.org/10.1177/02783649221080483

[15] P.-C. Kung, C.-C. Wang, and W.-C. Lin, "A normal distribution transform-based radar odometry designed for scanning and automotive radars," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 14 417–14 423.

[16] H. Lim, K. Han, G. Shin, G. Kim, S. Hong, and H. Myung, "Orora: Outlier-robust radar odometry," 2023.

[17] Y. Almalioglu, M. Turan, C. X. Lu, N. Trigoni, and A. Markham, "Milliro: Ego-motion estimation with low-cost millimetre-wave radar," *IEEE Sensors Journal*, vol. 21, no. 3, pp. 3314–3323, 2021.

[18] C. D. Monaco and S. N. Brennan, "Radarodo: Ego-motion estimation from doppler and spatial data in radar images," *IEEE Transactions on Intelligent Vehicles*, vol. 5, no. 3, pp. 475–484, 2020.

[19] P. Checchin, F. Gérossier, C. Blanc, R. Chapuis, and L. Trassoudaine, "Radar scan matching slam using the fourier-mellin transform," in *Field and Service Robotics*, A. Howard, K. Iagnemma, and A. Kelly, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 151–161.

[20] G. Kim, Y. S. Park, Y. Cho, J. Jeong, and A. Kim, "Mulran: Multimodal range dataset for urban place recognition," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 6246–6253.

[21] D. Barnes, R. Weston, and I. Posner, "Masking by moving: Learning distraction-free radar odometry from pose information," *CoRR*, vol. abs/1909.03752, 2019. [Online]. Available: http://arxiv.org/abs/1909.03752

[22] B. Reddy and B. Chatterji, "An fft-based technique for translation, rotation, and scale-invariant image registration," *IEEE Transactions on Image Processing*, vol. 5, no. 8, pp. 1266–1271, 1996.

[23] H. Bülow, M. Pfingsthorn, and A. Birk, "Using robust spectral registration for scan matching of sonar range data," *IFAC Proceedings Volumes*, vol. 43, no. 16, pp. 611–616, 2010, 7th IFAC Symposium on Intelligent Autonomous Vehicles. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1474667016351254

[24] H. Buelow and A. Birk, "Spectral registration of noisy sonar data for underwater 3d mapping," *Auton. Robots*, vol. 30, pp. 307–331, 04 2011.

[25] N. Hurtós, D. Ribas, X. Cufí, Y. Petillot, and J. Salvi, "Fourier-based registration for robust forward-looking sonar mosaicing in low-visibility underwater environments," *Journal of Field Robotics*, vol. 32, no. 1, pp. 123–151, 2015. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/rob.21516

[26] D. Louback da Silva Lubanco, T. Schlechter, M. Pichler-Scheder, and C. Kastl, "Survey on radar odometry," in *Computer Aided Systems Theory – EUROCAST 2022*, R. Moreno-Díaz, F. Pichler, and A. Quesada-Arencibia, Eds. Cham: Springer Nature Switzerland, 2022, pp. 619–625.

[27] K. Harlow, H. Jang, T. D. Barfoot, A. Kim, and C. Heckman, "A new wave in robotics: Survey on recent mmwave radar applications in robotics," 2023.

[28] H. Bay, A. Ess, T. Tuytelaars, and L. Van Gool, "Speeded-up robust features (surf)," *Computer Vision and Image Understanding*, vol. 110, no. 3, pp. 346–359, 2008, similarity Matching in Computer Vision and Multimedia. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1077314207001555

[29] K. Burnett, A. P. Schoellig, and T. D. Barfoot, "Do we need to compensate for motion distortion and doppler effects in spinning radar navigation?" *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 771–778, 2021.

[30] K. Burnett, D. J. Yoon, A. P. Schoellig, and T. D. Barfoot, "Radar odometry combining probabilistic estimation and unsupervised feature learning," *CoRR*, vol. abs/2105.14152, 2021. [Online]. Available: https://arxiv.org/abs/2105.14152

[31] Y. S. Park, Y.-S. Shin, and A. Kim, "Pharao: Direct radar odometry using phase correlation," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 2617–2623.

[32] R. Weston, M. Gadd, D. De Martini, P. Newman, and I. Posner, "Fast-mbym: Leveraging translational invariance of the fourier transform for efficient and accurate radar odometry," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 2186–2192.

[33] A. Venon, Y. Dupuis, P. Vasseur, and P. Merriaux, "Millimeter wave fmcw radars for perception, recognition and localization in automotive applications: A survey," *IEEE Transactions on Intelligent Vehicles*, vol. 7, no. 3, pp. 533–555, 2022.

[34] C. Knapp and G. Carter, "The generalized correlation method for estimation of time delay," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 24, no. 4, pp. 320–327, 1976.

[35] G. Carter, "Coherence and time delay estimation," *Proceedings of the IEEE*, vol. 75, no. 2, pp. 236–255, 1987.

[36] Q.-S. Chen, M. Defrise, and F. Deconinck, "Symmetric phase-only matched filtering of fourier-mellin transforms for image registration and recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 16, no. 12, pp. 1156–1168, 1994.

[37] J. L. Horner and P. D. Gianino, "Phase-only matched filtering," *Appl. Opt.*, vol. 23, no. 6, pp. 812–816, Mar 1984. [Online]. Available: http://opg.optica.org/ao/abstract.cfm?URI=ao-23-6-812

[38] X. Tong, Z. Ye, Y. Xu, S. Gao, H. Xie, Q. Du, S. Liu, X. Xu, S. Liu, K. Luan, and U. Stilla, "Image registration with fourier-based image correlation: A comprehensive review of developments and applications," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 12, no. 10, pp. 4062–4081, 2019.

[39] A. Alba, J. F. Vigueras-Gomez, E. R. Arce-Santana, and R. M. Aguilar-Ponce, "Phase correlation with sub-pixel accuracy: A comparative study in 1d and 2d," *Computer Vision and Image Understanding*, vol. 137, pp. 76–87, 2015. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1077314215000685

[40] M. Guizar-Sicairos, S. T. Thurman, and J. R. Fienup, "Efficient subpixel image registration algorithms," *Opt. Lett.*, vol. 33, no. 2, pp. 156–158, Jan 2008. [Online]. Available: https://opg.optica.org/ol/abstract.cfm?URI=ol-33-2-156

[41] T. Fujisawa and M. Ikehara, "High-accuracy image rotation and scale estimation using radon transform and sub-pixel shift estimation," *IEEE Access*, vol. 7, pp. 22 719–22 728, 2019.

[42] D. Sundararajan, *Digital Image Processing: A Signal Processing and Algorithmic Approach*, 10 2017.

[43] P. J. Rousseeuw and C. Croux, "Alternatives to the median absolute deviation," *Journal of the American Statistical Association*, vol. 88, no. 424, pp. 1273–1283, 1993. [Online]. Available: https://www.tandfonline.com/doi/abs/10.1080/01621459.1993.10476408

[44] C. Leys, C. Ley, O. Klein, P. Bernard, and L. Licata, "Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median," *Journal of Experimental Social Psychology*, vol. 49, no. 4, pp. 764–766, 2013. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0022103113000668

[45] T. C. Beers, K. Flynn, and K. Gebhardt, "Measures of Location and Scale for Velocities in Clusters of Galaxies—A Robust Approach," , vol. 100, p. 32, Jul. 1990.

[46] H. S. Stone, B. Tao, and M. McGuire, "Analysis of image registration noise due to rotationally dependent aliasing," *Journal of Visual Communication and Image Representation*, vol. 14, no. 2, pp. 114–135, 2003. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1047320303000026

[47] R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard, "G2o: A general framework for graph optimization," in *2011 IEEE International Conference on Robotics and Automation*, 2011, pp. 3607–3613.

[48] G. Grisetti, R. Kümmerle, C. Stachniss, and W. Burgard, "A tutorial on graph-based slam," *IEEE Intelligent Transportation Systems Magazine*, vol. 2, no. 4, pp. 31–43, 2010.

[49] F. Dellaert and M. Kaess, *Factor Graphs for Robot Perception*. Foundations and Trends in Robotics, Vol. 6, 2017. [Online]. Available: http://www.cs.cmu.edu/~kaess/pub/Dellaert17fnt.pdf

[50] D. Barnes, M. Gadd, P. Murcutt, P. Newman, and I. Posner, "The oxford radar robotcar dataset: A radar extension to the oxford robotcar dataset," *arXiv preprint arXiv: 1909.01300*, 2019. [Online]. Available: https://arxiv.org/pdf/1909.01300

[51] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *2012 IEEE Conference on Computer Vision and Pattern Recognition*, 2012, pp. 3354–3361.

[52] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of rgb-d slam systems," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012, pp. 573–580.

[53] Z. Zhang and D. Scaramuzza, "A tutorial on quantitative trajectory evaluation for visual(-inertial) odometry," in *2018 IEEE/RSJ International*

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Intelligent Vehicles. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TIV.2023.3324941

Conference on Intelligent Robots and Systems (IROS), 2018, pp. 7244–7251.

VIII. Biography Section

Daniel Louback S. Lubanco Daniel Louback S. Lubanco received the B.Sc. in Control and Automation Engineering from the Instituto Federal Fluminense, Brazil in 2018 and the M.Sc. in Automotive Mechatronics and Management from the University of Applied Sciences Upper Austria, Austria in 2020. Since 2019, he is with the Linz Center of Mechatronics and since 2022 he is also with the University of Applied Sciences Upper Austria and the Johannes Kepler University. He is currently working towards the PhD degree in Engineering Sciences - Mechatronics at the Johannes Kepler University, Austria. His research interests include radar signal processing, robotics and image processing.

Ahmed Hashem Ahmed Hashem is a PhD candidate at the Johannes Kepler University, Linz, Austria. In the research institute for Communications Engineering and RF-systems, he is developing autonomy and navigation algorithms for small aircraft and autonomous drones in GPS denied environments using Synthetic Aperture Radar. His research interests include: autonomy systems and algorithm development, environment perception for self-driving cars, drones and robots, state-estimation for autonomous vehicles and robots, computer vision, and Synthetic-Aperture Radar (SAR).

Markus Pichler-Scheder Markus Pichler-Scheder (Member, IEEE) was born in Linz, Austria, in 1976. He received the Dipl.-Ing. degree in mechatronics and the Dr.techn. degree from Johannes Kepler University, Linz, in 2002 and 2007, respectively. He was with the Institute for Communications and Information Engineering, University of Linz, Austria, and joined the Linz Center of Mechatronics GmbH, Austria, in 2002. He has also been a lecturer for signal processing, systems theory, and communications at the Johannes Kepler University and the University of Applied Sciences Upper Austria. He has authored or co-authored over 60 publications. His research interests include systems design, signal generation, digital signal processing, and estimation theory for radar and positioning systems.

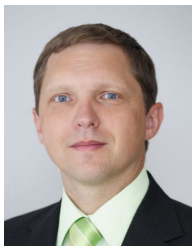

Andreas Stelzer Andreas Stelzer (M'00) received the Diploma Engineer degree in electrical engineering from the Technical University of Vienna, Vienna, Austria, in 1994, and the Dr. techn. degree (Ph.D.) in mechatronics (with honors sub auspiciis praesidentis rei publicae) from the Johannes Kepler University (JKU) Linz, Austria, in 2000. In 2003, he became Associate Professor with the Institute for Communications Engineering and RF Systems, Johannes Kepler University Linz. Since 2008, he has been a key researcher for the Austrian Center of Competence in Mechatronics (ACCM), where he is responsible for numerous industrial projects. In 2007 he was granted a Christian Doppler Research Laboratory for Integrated Radar Sensors and since 2011 he is full Professor at the Johannes Kepler University Linz, heading the Department for RF-Systems. He is co-founder of company Inras, meanwhile taken over by Joby-Austria, where radar sensors for urban air mobility are developed. Since 2020 he is also head of the joint Linz Institute of Technology (LIT) and Silicon-Austrian-Labs (SAL) Millimeter-Wave Lab working on combined sensing and communication applications for future 6G. He has authored or coauthored over 430 journal, conference and workshop contributions. His research is focused on microwave sensor systems for industrial and automotive applications, integrated radar sensor concepts, SiGe based circuit design, microwave packaging in eWLB, RF and microwave subsystems, surface acoustic wave (SAW) sensor systems and applications, as well as digital signal processing for sensor signal evaluation. Dr. Stelzer is a member of the Austrian ÖVE. He has served as an associate editor for the IEEE MICROWAVE AND WIRELESS COMPONENTS LETTERS. He was chair of MTT-27 Wireless-Enabled Automotive and Vehicular Applications. He was recipient of several awards including the 2008 IEEE Microwave Theory and Techniques Society (IEEE MTT-S) Outstanding Young Engineer Award, the 2011 IEEE Microwave Prize, and the Best Paper Award of the International Journal of Microwave and Wireless Technologies (IJMWT) 2016. Furthermore, he was co-recipient of the 2012 European Conference on Antennas and Propagation (EuCAP) Best Measurement Paper Prize, the 2012 Asia Pacific Conference on Antennas and Propagation (APCAP) Best Paper Award, the 2011 German Microwave Conference (GeMiC) Best Paper Award, as well as the EEEfCOM Innovation Award and the European Microwave Association (EuMA) Radar Prize of the European Radar Conference (EuRAD) 2003. He is a member of the IEEE MTT, IM, and CAS Societies and he served as IEEE Distinguished Microwave Lecturer for the period 2014 to 2016 and was chair of IEEE International Conference on Microwaves for Intelligent Mobility (ICMIM) 2020.

Reinhard Feger Reinhard Feger was born in Kufstein, Austria, in 1980. He received the Dipl.-Ing. (M.Sc.) degree in mechatronics and the Dr. Techn. (Ph.D.) degree in mechatronics from Johannes Kepler University Linz, Linz, Austria, in 2005 and 2010, respectively. In 2005, he joined the Institute for Communications and Information Engineering, Johannes Kepler University Linz, as a Research Assistant. In 2007, he became a member of the Christian Doppler Laboratory for Integrated Radar Sensors, Johannes Kepler University Linz. He is currently an Associate Professor with the Institute for Communications Engineering and RF-Systems, Johannes Kepler University Linz. His research topics are radar signal processing, as well as radar system design for industrial and automotive radar sensors. Dr. Feger was recipient of the 2011 Microwave Prize and the 2011 German Microwave Conference Best Paper Award. In 2012, he received the Best Measurement Paper Prize at the European Conference on Antennas and Propagation.

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

This article has been accepted for publication in IEEE Transactions on Intelligent Vehicles. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TIV.2023.3324941

Thomas Schlechter Thomas Schlechter (M'07) was born in Germany in 1979. In 2007 he earned a Dipl.-Ing. degree in electrical engineering, electronics and information technology and in 2010 a master's degree in economical sciences, psychology, and slavonic studies, both from the FAU Erlangen-Nuremberg. By the year 2012 he was working towards his PhD in the field of embedded systems and signal processing at the University of Klagenfurt, where he also received another master's degree in information management in 2016, both with distinction.

After some years in industry, he joined the University of Applied Sciences Upper Austria as Professor for Automotive Mechatronics in 2018. He published around 30 conference papers and journal articles and is author of more than 50 patents, delivering an H-index of 6. His research interest covers multiple fields, including mechatronics, algorithms, and especially communication theory and technology. As a special field of interest, cyber security topics may be mentioned as well.

Prof. Schlechter has been part of the organizing committee of the ICMRE conference since 2019 and the ISM conference until 2021.

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/