

DRO: Doppler-Aware Direct Radar Odometry

Cedric Le Gentil*, Leonardo Brizi[†], Daniil Lisus*, Xinyuan Qiao*, Giorgio Grisetti[†] and Timothy D. Barfoot*

*Autonomous Space Robotics Lab

University of Toronto Institute for Aerospace Studies (UTIAS), Toronto, Ontario, Canada

Email: {cedric.legentil; daniil.lisus; samxinyuan.qiao; tim.barfoot}@utoronto.ca

[†]Department of Computer, Control, and Management Engineering “Antonio Ruberti”

Sapienza University of Rome

Email: {brizi; grisetti}@diag.uniroma1.it

Abstract—A renaissance in radar-based sensing for mobile robotic applications is underway. Compared to cameras or lidars, millimetre-wave radars have the ability to ‘see’ through thin walls, vegetation, and adversarial weather conditions such as heavy rain, fog, snow, and dust. In this paper, we propose a novel SE(2) odometry approach for spinning frequency-modulated continuous-wave radars. Our method performs scan-to-local-map registration of the incoming radar data in a direct manner using all the radar intensity information without the need for feature or point cloud extraction. The method performs locally continuous trajectory estimation and accounts for both motion and Doppler distortion of the radar scans. If the radar possesses a specific frequency modulation pattern that makes radial Doppler velocities observable, an additional Doppler-based constraint is formulated to improve the velocity estimate and enable odometry in geometrically feature-deprived scenarios (e.g., featureless tunnels). Our method has been validated on over 250 km of on-road data sourced from public datasets (Boreas and MulRan) and collected using our automotive platform. With the aid of a gyroscope, it outperforms state-of-the-art methods and achieves an average relative translation error of 0.26% on the Boreas leaderboard. When using data with the appropriate Doppler-enabling frequency modulation pattern, the translation error is reduced to 0.18% in similar environments. We also benchmarked our algorithm using 1.5 hours of data collected with a mobile robot in off-road environments with various levels of structure to demonstrate its versatility. Our real-time implementation is publicly available: <https://github.com/utiasASRL/dro>.

I. INTRODUCTION

From air traffic control to weather forecasting via marine navigation, radar sensing is a cornerstone of modern societies. During the early days of mobile robotics, radars were used to detect specific landmarks and perform Simultaneous Localisation And Mapping (SLAM) [26]. While the advent of lidar technologies and the ubiquity of cameras drove the state estimation research community away from radar sensing [18], we are currently experiencing a ‘new wave’ of interest toward this modality [37]. This interest is partly driven by the robustness of radars with respect to adversarial weather conditions. In this paper, we present a novel algorithm for radar-based odometry and benchmark its performance across different datasets in both automotive and off-road robotic scenarios as shown in Fig. 1.

The working principle of radar sensing consists of emitting radio waves and analyzing the signal returned through various reflections of the waves from objects present in the environ-




Fig. 1. This paper presents a Direct Radar Odometry (DRO) method that provides state-of-the-art performance in a wide range of environments including unstructured ones (right) and geometrically degenerated ones such as featureless tunnels (left).

ment to collect geometric information about the surroundings. This work leverages data from millimeter wave (mmWave) radars that use electromagnetic waves with a length scale that is in the order of magnitude of one millimetre. Waves in such frequency bands can penetrate through plastic, textiles, glass, etc. Accordingly, in the context of mobile robotics, radar allows the system to sense its environment in a wide range of conditions, even under intensely adverse weather conditions such as fog, rain, and snow. Like any wave, the radar-emitted signals are subject to the Doppler effect when the relative velocity between the system and part of its environment is not zero. Thus, when using specific frequency patterns and processing, radar data can allow for the extraction of Doppler-based radial velocity information. These benefits represent a competitive advantage of radars compared to other modalities, such as vision or commonly used lidars. The proposed odometry method accounts for the Doppler-based distortion of the radar data and can integrate a Doppler-specific velocity constraint in its optimization, allowing it to robustly perform odometry even in geometrically feature-deprived environments, such as featureless tunnels (see Fig. 1 (left)).

Two types of radar sensors are typically used in robotics applications: ‘fixed’ phased-array radars and spinning radars.

In this work, we focus on spinning radars, as they provide a 360° Field-of-View (FoV) that contains dense intensity information from up to several hundred meters away. This makes spinning radar enticing for navigational tasks such as odometry, localization, and object tracking, where a view of the entire surroundings is required for robustness and safety. The most common spinning-radar-odometry approaches use the dense intensity information to extract a set of points or features [21, 3, 17, 9, 14, 47, 54]. These approaches work well, with point-based methods currently being the state-of-the-art (SOTA) [3, 46], but they require computationally costly data association, critical parameter tuning, and they ultimately discard a large amount of the gathered data in favour of sparse representations. Methods that have used the entire scan directly for odometry have done scan-to-scan correlation matching on the intensity image [52], sometimes requiring pre-processing using a deep learning network [10, 62]. However, these methods do not explicitly account for the motion and Doppler distortion of the radar data, and discretize the relative transform search space to lower the computational cost. Additionally, the need for a learned processing component limits the transferability of these approaches to different environments.

In this work, we present the first Direct Radar Odometry (DRO) framework that accounts for the continuous motion and Doppler distortion of the incoming data in a principled way. It can be aided by a yaw-axis gyroscope and does not rely on deep-learning techniques or hand-tuned discretization parameters. As demonstrated on more than 250 km of data, it is applicable to a wide range of environments. With the appropriate radar emission pattern, it can also perform in geometrically challenging scenarios that generally lead to degenerated state estimates. The main contributions are as follows:

- 1) The formulation of the first direct radar registration method that accounts for continuous motion and Doppler-based distortion in a gradient-based optimization that does not discretize the search space.
- 2) A novel way to leverage the Doppler effect for velocity estimation from spinning radar without assuming that consecutive radar beams correspond to the same objects.
- 3) Our method is capable of SOTA radar-odometry performance when used with a gyroscope.
- 4) A real-time GPU-based implementation.¹

II. RELATED WORK

A radar odometry system [37, 1, 61, 20], like other odometry systems, aims to estimate the ego-motion of the sensor, in this case, a radar, thus determining the trajectory of the platform on which the sensor is mounted. The motion estimation process typically involves formulating an optimization problem, where the objective is to find the sensor's state and an environment model that maximizes the probability

of observing the measured data. This principle is known as maximum-likelihood estimation.

In the last decade, radar odometry has been dominated by indirect approaches [1], which rely on selecting salient points or features from radar measurements. Early works, such as [19], started by leveraging computer vision approaches, such as SIFT [50] features, from radar Cartesian images, while more recent state-of-the-art methods, such as CFAR [3], select points based on each individual azimuths of a radar signal. There exist two main categories for feature extraction methods [53]: signal-based and spatial-based extractors. Signal-based methods identify points by applying thresholds to individual azimuths in the radar signal. A well-known example is the Constant False Alarm Rate (CFAR) method [29], which processes each azimuth of a radar scan using a sliding window. The power readings within the window are used to calculate a detection threshold. Many variants of CFAR exist [51, 57, 33, 65, 63, 60, 11]. Another signal-based extractor, K-strongest [2], selects the K strongest power returns above a static threshold along each azimuth. In contrast, spatial-based extractors rely on extracting features and descriptors from radar scans transformed into Cartesian images, some of those leverage computer vision algorithms, such as ORB [58] or SIFT [50].

Indirect odometry systems, extracting points from the radar data, rely on matching these features. The data association step is one of the most critical building blocks of these systems. Burnett et al. [13] introduced motion-compensated RANSAC to ensure robust association. Subsequently, they developed Hero [14], a method that leverages learned key-point locations, uncertainties, and descriptors to address outliers and facilitate matching within the ESGVI [8] batch state estimation framework. In CFAR [3], the authors extract the surface points and demonstrate that the point-to-line and point-to-distribution perform better than point-to-point, as these metrics mitigate the problem of erroneous associations between individual points. After an initial scan filtering step, Kung et al. [42] extract sets of local Gaussian distributions and perform point-to-distribution registration. Finally, some other systems exploit data-driven approaches to downscale noisy point contributions in the optimization problem [48] or to select others suitable for the data association [4].

Alternatively, direct methods do not need explicit data association. While this kind of method has been well studied for lidar [23, 25, 35] and monocular odometry [28, 30, 36, 24, 27], the literature on direct radar-based odometry is still somewhat limited and methods are slower than feature-based methods. Barnes et al. [10] proposed a brute-force approach that samples and selects rotations to maximize the dense correlation between radar scans. Similarly, Checchin et al. [22] and Park et al. [52] addressed the need for sampling by decoupling rotation and translation. They utilized the Fourier-Mellin Transform, maximizing phase correlation between log-polar and Cartesian images to estimate rotation and translation, respectively. A coarse-to-fine phase correlation refinement on Cartesian images was then applied to improve the translation

¹<https://github.com/utiasASRL/dro>

estimate. While Barnes et al. [10] achieved remarkable performance in odometry estimation, their method suffered from high computational complexity. To address this, Weston et al. [62] replaced the brute-force rotation search with an FFT-based approach, significantly reducing computation time at the cost of slightly lower accuracy. Overall, existing direct methods either decouple the estimation of rotation and translation or sub-sample the problem within a discrete search space. These strategies can result in susceptibility to local minima and poor scalability across different resolutions. In contrast, our approach optimizes the locally continuous SE(2) radar trajectory without discretizing the search space.

For both direct and indirect methods, motion distortion can significantly degrade system performance. In the literature, this issue has been addressed using continuous-time trajectory estimation [5], [7, 6, 31] for lidar [12, 59] and rolling-shutter cameras [38]. This work applies similar concepts to radar data while also correcting Doppler distortion in a continuous-time optimization. The Doppler information channel provides valuable additional data that can be leveraged for odometry and ego-motion estimation. Several works have explored its integration with different models and estimation techniques to enhance motion estimation accuracy. Kellner et al. [39] demonstrated that automotive radar, combined with an Ackermann vehicle model, can estimate ego-motion by fitting least-squares to stationary point Doppler velocities identified via RANSAC. In [34], a gyroscope is used to alleviate the need for an Ackermann platform. Galeote-Luque et al. [32] extended this by using Doppler measurements for 3D odometry with RANSAC and a kinematic model. Kubelka et al. [41] showed that integrating high-quality Doppler data with an IMU outperforms ICP in feature-deprived environments, and Lisus et al. [49] implemented a continuous-time approach with a white-noise-on-acceleration prior. In [56], the authors extend the deep-learning-based preprocessing from [10] and trained an additional network for per-scan velocity regression. In contrast, our model-based approach does not require any training data and does not suffer from any generalization issues.

III. BACKGROUND

A. Gaussian process regression

Gaussian Process (GP) regression [55] is a probabilistic, non-parametric interpolation method. Let us consider a signal $h(\mathbf{x})$, with $\mathbf{x} \in \mathbb{R}^D$, modeled with a GP $h(\mathbf{x}) \sim \mathcal{GP}(0, k(\mathbf{x}, \mathbf{x}'))$. The function $k(\mathbf{x}, \mathbf{x}')$ is the so-called covariance kernel that characterizes the covariance between two instances of h (at $\mathbf{x}$ and $\mathbf{x}'$). Given noisy samples $y_i = h(\mathbf{x}_i) + \eta_i$, with $\eta_i \sim \mathcal{N}(0, \sigma)$ and $i = (1, \dots, N)$, GP regression aims at inferring the signal h and its variance at any new input location $\mathbf{x}^*$. The definition of a GP corresponds to the following multivariate normal distribution:

$$\begin{bmatrix} \mathbf{y} \\ h(\mathbf{x}^*) \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \mathbf{K}_{\mathbf{X}\mathbf{X}} + \sigma^2 \mathbf{I} & \mathbf{k}_{\mathbf{X}\mathbf{x}^*} \\ \mathbf{k}_{\mathbf{x}^*\mathbf{X}} & k(\mathbf{x}^*, \mathbf{x}^*) \end{bmatrix} \right), \quad (1)$$




Fig. 2. Impact of different frequency modulation patterns (top) and Doppler effect on the data collected with spinning FMCW radars (bottom).

where $\mathbf{k}_{\mathbf{x}^*\mathbf{X}} = (\mathbf{k}_{\mathbf{X}\mathbf{x}^*})^\top = [k(\mathbf{x}^*, \mathbf{x}_1) \dots k(\mathbf{x}^*, \mathbf{x}_N)]$, $\mathbf{K}_{\mathbf{X}\mathbf{X}} = [\mathbf{k}_{\mathbf{X}\mathbf{x}_1} \dots \mathbf{k}_{\mathbf{X}\mathbf{x}_N}]$, and $\mathbf{y} = [y_1 \dots y_N]$. By conditioning (1) with respect to the noisy observations we obtain the inference of h at $\mathbf{x}^*$ as

$$\begin{aligned} \hat{h}(\mathbf{x}^*) &= \mathbf{k}_{\mathbf{x}^*\mathbf{X}} (\mathbf{K}_{\mathbf{X}\mathbf{X}} + \sigma^2 \mathbf{I})^{-1} \mathbf{y}, \\ \text{var}(\hat{h}(\mathbf{x}^*)) &= k(\mathbf{x}^*, \mathbf{x}^*) - \mathbf{k}_{\mathbf{x}^*\mathbf{X}} (\mathbf{K}_{\mathbf{X}\mathbf{X}} + \sigma^2 \mathbf{I})^{-1} \mathbf{k}_{\mathbf{X}\mathbf{x}^*}. \end{aligned} \quad (2)$$

B. Radar data and Doppler effect

Let us consider the type of radars commonly used in the automotive and robotic industries: Frequency-Modulated Continuous-Wave (FMCW) radars. As their name suggests, these radars are continuously emitting frequency-modulated electromagnetic waves. Thus, instead of directly measuring the time between the emission and reception of waves, the distance to an electromagnetic reflector corresponds to a frequency difference Δf between the signal transmitted and received. For the sake of simplicity, the following explanation will consider a single reflector at a certain distance r from the radar. In reality, the FMCW radars provide a reflection score/value for many ‘range bins’ along the emission beam.

Considering linear modulation of the frequency with a slope s as shown in Fig. 2, the distance r between the radar and reflector is

$$\tilde{r} = \frac{c \Delta f}{2s}, \quad (3)$$

where c is the speed of light, and the divider 2 accounts for the fact that Δf represents twice the distance r due to the time for the wave to reach the reflector and return. This relationship is only true ($\tilde{r} = r$) if the relative velocity between the sensor and the reflector is null. In the presence of relative motion, the measured Δf is impacted by the Doppler effect that compresses or expands the electromagnetic wave, thus resulting in $\Delta f = \Delta f_t + \Delta f_d$, with Δf_t being the shift corresponding to the waves travelling time, and Δf_d the Doppler-induced frequency shift. This shift Δf_d is proportional to the signal frequency $\frac{1}{\lambda_t}$ and the velocity along the emission beam, denoted ‘radial velocity’ u , as

$$\Delta f_d = \frac{2u}{\lambda_t}. \quad (4)$$

Accordingly, the ‘range’ measurement provided by (3) is perturbed by the Doppler effect

$$\tilde{r} = \frac{c\Delta f_t}{2s} + \frac{c\Delta f_d}{2s} = r + r_d(u). \quad (5)$$

When considering a frequency modulation with a *sawtooth* pattern and multiple measurements of the same reflector, we cannot dissociate Δf_t and Δf_d from Δf . However, if we use a triangle pattern consisting of a succession of ‘up-chirps’ and ‘down-chirps’ as illustrated in Fig. 2, two measurements of the same reflector allow for the recovery of the radial velocity:

$$\tilde{r}_\uparrow - \tilde{r}_\downarrow = \frac{c\Delta f_t}{2s} + \frac{c\Delta f_d}{2s} - \frac{-c\Delta f_t}{-2s} - \frac{c\Delta f_d}{-2s} \quad (6)$$

where the components in Δf_t cancel out resulting in

$$\Delta \hat{r} = \tilde{r}_\uparrow - \tilde{r}_\downarrow = \frac{2uc}{s\lambda_y} = \beta u. \quad (7)$$

In reality, FMCW radars do not only expect one return per beam but collect ‘reflectivity’ information through a wide range of discrete ranges through a fast-Fourier transform of the received signal. Thus, the data from such radars can be interpreted as intensity/reflectivity images, in which rows correspond to different beam azimuths and columns to different ranges. Fig. 2 illustrates the effect of the Doppler effect on FMCW data. The data collected with a moving radar using a triangular pattern presents an alternating shift between each row of the scans, while all the rows are shifted according to the sign of the radial velocity for sawtooth radars. Depending on the radar model, the modulation pattern can be changed or not in the firmware. Thus, not all spinning FMCW radars enable the extraction of Doppler-based velocities.

IV. METHOD

The proposed method addresses the issue of radar odometry optionally aided by a gyroscope. The main concept is the direct use of FMCW intensity data without any point cloud or feature extraction, or data association.

A. Problem statement and overview

Let us consider a 2D mmWave FMCW spinning radar and a rigidly mounted gyroscope. Without loss of generality, the extrinsic calibration (rotation matrix) between the two sensors is assumed to be known and will not appear in the following derivations. The gyroscope provides measurements $\tilde{\omega}_i$, at times t_i , of the system’s angular velocity around the spinning axis of the radar. The radar data is assumed to be collected in scans covering the 360° field-of-view. Each scan contains information collected from N beams, each associated with an azimuth α_n and timestamp t_n , and each beam consists of the electromagnetic reflection intensity ψ_{nm} at M different ranges $\hat{r}_m$. Accordingly, a radar scan can be interpreted as a polar intensity image or a set of tuples $\{\alpha_n, \hat{r}_m, \psi_{nm}\}$.

The rotation, position, and velocity of the radar reference frame with respect to an Earth-fixed frame are described with continuous functions $t \in \mathbb{R} \mapsto \mathbf{R}_W^t \in SO(2)$, $\mathbf{p}_W^t \in \mathbb{R}^2$, and $\mathbf{v}_W^t \in \mathbb{R}^2$, respectively. The details of the various motion

models used to obtain $\mathbf{R}_W^t$, $\mathbf{p}_W^t$, and $\mathbf{v}_W^t$ from a finite set of state variables $\mathcal{S}$ will be discussed in Section IV-D. The proposed method aims to estimate the state variables $\mathcal{S}$, thus the sensor’s trajectory, during the collection time of the last radar scan using the radar intensity data and optionally the angular velocity from the gyroscope. The estimation is done with the numerical optimization of a combination of objective functions as

$$\mathcal{S}^* = \underset{\mathcal{S}}{\operatorname{argmax}} \quad O_\psi + O_d, \quad (8)$$

with O_ψ an objective function that relates to the direct registration of the radar intensity data (cf., Section IV-B, and O_d to the Doppler effect (cf., Section IV-C). Fig. 3 provides an overview of the proposed method. As detailed later in this section, both O_ψ and O_d can be interpreted as cross-correlation scores between different parts of the present and past radar data.

B. Direct intensity objective function

This objective function is inspired by direct methods in traditional monocular state estimation. The goal is to estimate the continuous pose of the sensor to register/align the last radar scan with past radar data. In this work, we opted for a local map that is updated ‘on-the-fly’ as a way to combine past radar scans into a single image representation of the sensor’s environment with intensity in Cartesian coordinates (as opposed to raw radar data in polar coordinates). The motivation for using a local map instead of solely the last radar scan is the higher robustness to outliers such as moving vehicles or multipath echoes.

1) *Direct registration*: Let us assume the availability of an image-like local map $\mathcal{M}$ of the sensor’s environment expressed in the radar’s reference frame at time t_1 . The rows and columns of $\mathcal{M}$ image correspond to discrete Cartesian coordinates. The function $\Gamma_{\mathcal{B}}(\mathcal{M}, x, y)$ allows querying the local map intensity for any x and $y \in \mathbb{R}$ with bilinear interpolation.

To compute the direct registration objective between the incoming radar scan and the local map $\mathcal{M}$, we need to correct for the Doppler effect and the motion of the radar as

$$\begin{bmatrix} \tilde{x}_{nm} \\ \tilde{y}_{nm} \end{bmatrix} = \mathbf{R}_{t_1}^{t_n} \left(\left(\hat{r}_{nm} \pm \frac{\Delta r_n}{2} \right) \begin{bmatrix} \cos(\alpha_n) \\ \sin(\alpha_n) \end{bmatrix} \right) + \mathbf{p}_{t_1}^{t_n}, \quad (9)$$

with Δr_n the Doppler-induced shift (7) as a function of the body-centric velocity and the beam azimuth:

$$\Delta r_n = \beta \begin{bmatrix} \cos(\alpha_n) \\ \sin(\alpha_n) \end{bmatrix}^\top \mathbf{R}_W^{t_n \top} \mathbf{v}_W^{t_n}. \quad (10)$$

The sign $\pm$ is defined by the chirp direction for the azimuth α_n . Please note that this objective function accounts for the Doppler effect but does not require an FMCW radar with a triangular pattern. Once the radar scan is corrected using the current estimate of the trajectory, the objective function is simply the cross-correlation score,

$$O_\psi = \sum_{n=1}^N \sum_{m=1}^M \psi_{nm} \Gamma_{\mathcal{B}}(\mathcal{M}, \tilde{x}_{nm}, \tilde{y}_{nm}), \quad (11)$$

between the corrected scan and the local map.




Fig. 3. Overview diagram of the proposed direct radar odometry method. To leverage the Doppler-based velocity constraint, the radar must use a triangular frequency modulation pattern. Note that O_ψ accounts for motion and Doppler-induced distortion.

2) *On-the-fly local map*: The proposed local map consists of an image-like representation of the past radar data using a per-pixel low-pass filter to update $\mathcal{M}$ with the last radar scan and the corresponding state estimate. Concretely, given the convergence of the optimization problem (8), the radar scan is corrected to the first timestamp of the following scan in Cartesian coordinates (using (9) replacing $\mathbf{R}_{t_1}^{t_n}$ and $\mathbf{p}_{t_1}^{t_n}$ with $\mathbf{R}_{t_{N+1}}^{t_n}$ and $\mathbf{p}_{t_{N+1}}^{t_n}$, respectively). The corrected radar data is converted into an image-structure $\tilde{\mathcal{I}}_{\mathcal{M}}$ with the help of bilinear interpolation. Then the local map update is

$$\mathcal{M} \leftarrow (1 - \gamma)\mathcal{M} + \gamma\tilde{\mathcal{I}}_{\mathcal{M}}, \quad (12)$$

with $\gamma \in [0, 1]$ a user-defined parameter ($\gamma = 0.1$ in our implementation). Note that after receiving the very first radar frame, the local map is directly initialized with $\tilde{\mathcal{I}}_{\mathcal{M}}$.

C. Doppler-based objective function

As presented in Section III-B, an FMCW radar using a triangular modulation pattern (alternating up-chirp and down-chirp for each collected beam) allows for observability of the radial velocity along that beam if the consecutive beams observe the same part of the environment. Unfortunately, this is not quite true in the context of mobile radar sensing, since the vehicle and the radar dish move in between consecutively recorded returns. We propose to split the incoming radar scan into two image-like structures based on chirp direction and use GP regression (2) to ‘infill’ the missing rows in both images. This is illustrated in Fig. 4. We denote the up-chirp and down-chirp infilled images as $\mathcal{I}^\uparrow = \{\alpha_n, \hat{r}_m, \psi_{nm}^\uparrow\}$ and $\mathcal{I}^\downarrow = \{\alpha_n, \hat{r}_m, \psi_{nm}^\downarrow\}$, respectively. Thanks to the GP regression step, the part of the environment observed in each row of $\mathcal{I}^\uparrow$, be it observed or interpolated, is the same as in the corresponding row of $\mathcal{I}^\downarrow$.

We show in Fig. 5 that when the sensors are static, both $\mathcal{I}^\uparrow$ and $\mathcal{I}^\downarrow$ overlap well. However, the Doppler shift is clearly visible in the difference between $\mathcal{I}^\uparrow$ and $\mathcal{I}^\downarrow$ when the sensors move. The objective function O_d aims to constrain the estimated velocity $\mathbf{v}(t)$ using $\mathcal{I}^\uparrow$ and $\mathcal{I}^\downarrow$. First, we must compute the ‘range shift’, with (10), between each row of the two images as a function of the sensor velocity. Then, we define the function $\Gamma_{\mathcal{L}}(\mathcal{I}^\uparrow, \alpha_n, r)$ that queries the intensity value of

$\mathcal{I}^\uparrow$ for any value of $r \in \mathbb{R}$ using linear interpolation for each of the image’s row. The objective function is a cross-correlation score between $\mathcal{I}^\downarrow$ and the Doppler-rectified version of $\mathcal{I}^\uparrow$:

$$O_d = \sum_{n=1}^N \sum_{m=1}^M \psi_{nm}^\downarrow \Gamma_{\mathcal{L}}(\mathcal{I}^\uparrow, \alpha_n, \hat{r}_m + \Delta r_n). \quad (13)$$

The objective function O_d is maximised when the body-centric velocity estimate $\mathbf{R}(t_n)^\top \mathbf{v}(t_n)$ corresponds to the actual velocity of the sensor.

D. Motion models

In this section, we present different ways to obtain the continuous trajectory functions $\mathbf{R}_W^t$, $\mathbf{p}_W^t$, and $\mathbf{v}_W^t$ from the state variables $\mathcal{S}$. As $\mathbf{R}_W^t$ is in $SO(2)$, the rotation can easily be represented with a simple rotation angle $\theta(t)$ as

$$\mathbf{R}_W^t = \mathbf{R}_W^{t_1} \begin{bmatrix} \cos(\theta(t)) & \sin(\theta(t)) \\ -\sin(\theta(t)) & \cos(\theta(t)) \end{bmatrix}, \quad (14)$$

with $\mathbf{R}_W^{t_1}$ being the starting rotation obtained from the state estimate of the previous radar scan. We introduce two different models to obtain $\theta(t)$ in Subsections IV-D1, and IV-D2. Regarding the translation and velocity, we use a simple constant body-centric velocity model in Subsection IV-D3. For convenience, we define $\mathcal{S} = \{\mathcal{S}_R, \mathbf{v}_b\}$ with $\mathcal{S}_R$ the rotational state variables, and $\mathbf{v}_b$ the body-centric velocity.

1) *Constant angular velocity*: In the absence of a gyroscope, the proposed method can rely on the assumption of constant angular velocity ω for the duration of any radar scan with $\theta(t) = \omega(t - t_1)$. Therefore, the state variable is $\mathcal{S}_R = \{\omega\}$.

2) *Gyro preintegration*: Inspired by [45], when using a gyroscope, the change of orientation of the sensor suite is defined as a function of the angular rate measurements. To account for the frequency discrepancy between the radar azimuths and the gyroscope measurements, it leverages continuous preintegration concepts drawn from [44] and [43]. For the sake of simplicity and lightweight computation, we use a simple piece-wise linear model

$$\theta(t) = \theta(t_i) + \frac{(\omega_{i+1} + \omega_i)(t - t_i)}{2(t_{i+1} - t_i)} + \omega_i(t - t_i), \quad (15)$$




Fig. 4. Illustration of the proposed infilling. The raw Doppler distorted radar data (left) is split between up and down-chirp azimuths. Gaussian Process (GP) regression is used to interpolate every second azimuth (middle two, the interpolated portion is in grey). Both $\mathcal{I}^\uparrow$ and $\mathcal{I}^\downarrow$ virtually observe the same geometry without assuming that two consecutive azimuths of the raw data observe the same part of the environment (right, with blue positive and red negative).




Fig. 5. Illustration of the GP-based infill highlighting the Doppler effect on radar data collected with a triangular frequency modulation pattern (the blue arrow is the sensor velocity).

with $\omega_i = \tilde{\omega}_i - b$, $t \in [t_i, t_{i+1}]$. We decided not to include the gyroscope bias b in the estimated state. In our experiments, biases range from $\approx 5e^{-5}$ to $\approx 4e^{-3}$ corresponding to a maximum of 0.057° drift over the 250 ms of a scan. With the typical 0.9° radar angular resolution, the signal-to-noise ratio is too high to enable robust bias estimation in a ‘frame-to-frame’ approach like ours. This is especially true as the proposed objective functions are not probabilistic, thus impeding the integration of additional residuals/objective functions from other modalities and the probabilistic prior/previous knowledge of the bias. Accordingly, when using a gyroscope, the rotational state is empty $\mathcal{S}_R = \emptyset$. However, we use a separate online bias estimation strategy based on a simple yet effective heuristic: when the vehicle’s velocity is zero, so is the angular rate. Accordingly, the bias estimate is initialized by averaging the first Q raw gyroscope measurements collected when the velocity estimate is under 5 cm/s. Then it is updated with a low-pass filter whenever the velocity estimate is under 5 cm/s.

3) *Constant body-centric velocity*: This model assumes a constant body-centric velocity $\mathbf{v}_b \in \mathbb{R}^2$ during any radar scan. The velocity in the global frame is $\mathbf{v}_W^t = \mathbf{R}_W^t \mathbf{v}_b$. The position is

$$\mathbf{p}_W^t = \mathbf{p}_W^{t_1} + \left(\int_{t_1}^t \mathbf{R}_W^s ds \right) \mathbf{v}_b, \quad (16)$$

with $\mathbf{p}_W^{t_1}$ the position resulting from the state estimation during the previous radar scan. The integral in (16) is computed analytically when using the constant angular velocity assumption. Otherwise, when using a gyroscope, it is numerically preintegrated using a piece-wise linear model similar to (15) upon each of the matrix elements.

V. EXPERIMENTS

A. Implementation

The proposed method has been implemented in Python using PyTorch for matrix and algebra operations. Accordingly, it can easily be run on either a GPU or a CPU. Note that we have implemented (8) with analytical Jacobians and our solver without relying on PyTorch’s automatic differentiation and integrated solvers. When using a gyroscope and the Doppler-based velocity constraints, our method converges in 13.5 iterations on average.

1) *Gaussian process regression*: GP regression is known to suffer cubic computational complexity due to the matrix inversion in (2). With around 2 million points in each radar scan, naive GP regression to generate $\mathcal{I}^\uparrow$ and $\mathcal{I}^\downarrow$ is intractable. Fortunately, efficient computation strategies are enabled by the grid-pattern (image-like) nature of both the radar data and the inference locations in the proposed GP-based infill step. By only considering the data in a $U \times V$ neighbourhood around each inference location, $\mathbf{k}_{\mathbf{x}^* \times \mathbf{x}} (\mathbf{K}_{\mathbf{x}\mathbf{x}} + \sigma^2 \mathbf{I})^{-1}$ can be precomputed and the multiplication with $\mathbf{y}$ is a simple convolution operation. Accordingly, with a GPU implementation, the inference of $\mathcal{I}^\uparrow$ and $\mathcal{I}^\downarrow$ is extremely efficient.

Similarly to [49], each row of the GP-inferred image is independently filtered to attenuate multipath and specular noise. First, the intensity standard deviation is computed and any range bin with a value inferior to twice the deviation is nullified. Then the intensity values are normalized so that the maximum value equals one. Per-azimuth Gaussian blur is applied to smooth the discontinuities introduced in the previous steps. Finally, each intensity value is cubed to augment contrast.

2) *Optimisation*: The optimization problem (8) is solved with a gradient ascent algorithm with an update step

$$\mathcal{S} \leftarrow \mathcal{S} + c \frac{\nabla(O_d + O_\psi)}{\|\nabla(O_d + O_\psi)\|}, \quad (17)$$




Fig. 6. Illustration of a challenging situation with a line of vehicles coming the other way while driving over a skyway (very few static elements visible in the radar data).

where ∇ is the gradient operator, and c is a constant equal to 0.1 at the start of the optimization and halved for every non-ascending step. Note that the objective functions are element-wise products of intensity values. Thus, the computations linked to intensity values that are equal to zero do not contribute to the optimization problem and can be omitted to lower the computational burden.

3) *Robust weighting*: The proposed objective functions O_d and O_ψ are generally robust to outliers such as moving vehicles and specular noise. However, in very rare occasions (2 frames among 42k in our automotive experiments), a very high proportion of consistent outliers can lead to erroneous state estimates. A typical example is illustrated in Fig. 6 when the sensing platform is moving over a bridge/skyway with a very small amount of static reflectors in the surroundings, and a line of vehicles comes the other way. It can make the velocity estimate suddenly jump unrealistically. Such scenarios are easily detected with a simple threshold on the sensor’s acceleration from one scan estimate to the next. When detected, we re-run the optimization weighting the objective functions O_ψ and O_d as

$$O_\bullet = \sum_{n=1}^N \sum_{m=1}^M \rho_{nm}^\bullet \psi_{nm} \Gamma(\cdot, \cdot, \cdot), \quad (18)$$

$$\rho_{nm}^\bullet = (|\psi_{nm} - \Gamma(\cdot, \cdot, \cdot)| - 1)^6,$$

with $\bullet = d$ or ψ . Intuitively, the weighting ρ acts as an outlier-rejection mechanism by down-weighting the objective components that are too dissimilar in terms of intensity.

4) *Velocity bias in Doppler constraint*: Similarly to [49], we have empirically observed a velocity bias when performing Doppler-only velocity estimation. This bias mostly concerns the lateral velocity of the sensing platform. We believe this

bias is partly explained by the disparity between the actual measurement process and the simple measurement model that assumes the information contained in a row of the radar data corresponds to a single azimuth/timestamp: in reality, the radar’s scanning pattern is continuous. Thus, every measurement from a row could correspond to a different azimuth and timestamp. We think that the motion model’s ‘inaccuracy’ leads to a velocity bias that appears as environment-dependent due to the ignored correlation between the data range and azimuth/timestamp. The authors of [49] addressed this issue by fitting a linear model to the velocity error using held-out data with accurate velocity ground-truth. Then, the linear model is used to correct the velocity estimates directly. In this work, we propose to leverage the non-slip kinematics constraint to estimate and compensate for the lateral component of the velocity bias online without a ‘calibration’ phase.

Using the angular rate and the extrinsic calibration between the radar and the vehicle’s rear axle, the velocity estimated based solely on the Doppler-based constraint O_d is projected into the axle reference frame. The y component is projected back to the radar frame to update a low-pass filter that tracks the lateral velocity bias. The filter’s output is then used to correct the radial velocities when computing the shifts (10) for Doppler distortion correction. This estimation and correction process is only applied when O_d is part of the optimization problem (8). It is important to note that this mechanism does not strictly enforce the non-slip constraint thanks to the low-pass filter. Accordingly, the vehicle can be subject to momentary side slips without altering the state estimates.

B. Metrics

To benchmark our algorithm, we use well-established metrics in the odometry literature. For datasets with ground-truth orientation, the KITTI relative translation and rotation errors are adopted. Succinctly, for one pose estimate, the estimated trajectory is aligned with the ground-truth. Then, after a certain distance travelled from the aligned poses (segments varying from 100 to 800 m), the estimated and ground-truth poses are used to compute position and orientation errors relative to the segment lengths. This process is repeated every 5 poses of the estimated trajectory and the results are reported as the average relative errors in % and $^\circ/100\text{m}$. In the absence of ground-truth orientation, we adapted the Relative Position Error (RPE) introduced in [64] to SE(2) trajectories. It corresponds to the position Root Mean Squared Error (RMSE) between aligned segments of the ground-truth and the estimated trajectory. The length of segments [50 m, 100 m, 150 m, 200 m]. Similarly to the KITTI metric, we display the RMSE error as a percentage of the segment lengths.

C. Doppler-enabled automotive odometry

In this section, we provide a detailed analysis of our method’s performance on a dataset collected with our automotive sensing platform. Note that the radar used in this set of experiments uses a triangular frequency modulation pattern,

thus enabling the use of the Doppler-based velocity constraints O_d in (8).

1) *Dataset description*: The goal of this set-up is to replicate experiments from [49]. To do so, we have collected data with an automotive platform equipped with a Navtech RAS6 radar, a Silicon Sensing DMU41 Inertial Measurement Unit (IMU), a Velodyne Alpha Prime lidar, and an accurate Applanix RTK-GNSS/INS solution for ground-truth with post-processing. All the sensors have been accurately synchronized. The data used for benchmarking has been collected repeatedly four times along four routes (16 sequences in total) displaying an increasing level of difficulty: *Suburbs*, *Highway*, *Tunnel*, and *Skyway*. The first two environments provide a relatively large amount of geometric constraints via the presence of many buildings and other human-made structures in the vehicle’s surroundings. In *Tunnel* and *Skyway*, the geometric features are either degenerate (not constraining all the movement’s axes) or close to nonexistent, respectively. In these scenarios, only Doppler-induced information can lead to usable odometry performances. The skyway scenario is especially challenging as for a couple of kilometres the radar does not observe any significant static features except for sparse lamp posts and weak returns from the side barriers (the radar is mounted on the top of the vehicle and the barriers are relatively short, thus the bulk of the emission cone passes over the barriers). Accordingly, the ratio of moving to static elements in the radar’s FoV is very high. The length and velocity characteristics of each sequence type are given in Table I along with the results discussed later in Section V-C3. Fig. 1 (left) shows our collection vehicle while recording a *Tunnel* sequence.

2) *Baselines*: To benchmark the proposed algorithm, we have replicated four baselines from [15] and [49], three based on radar data and one on lidar. The first baseline, denoted CT-R, corresponds to the odometry component of the publicly available ‘teach and repeat’ framework [15]. It consists of an ICP-based continuous-time scan registration using point clouds extracted from the raw radar data. The trajectory estimates are computed via a sliding window optimization based on a continuous-time GP state representation. Note that the velocity estimates are used to compensate for the Doppler-induced shift of the radar points. The second baseline is DG and replicates the work from [49] on per-azimuth radial velocity extraction, followed by a RANSAC-based outlier rejection and robust optimization. With the constant-velocity assumption, it allows for the estimation of the vehicle’s velocity for each radar scan. Combined with the gyroscope’s angular rate, it enables association-free radar odometry, even in geometrically feature-deprived environments. Despite not extracting any point cloud out of the radar data, one can consider the radial velocity as ‘features’ extracted in a ‘front-end’ before an optimization-based state estimation step. Third is CT-RDG, the integration of the per-scan velocity estimates from [49] and the gyroscope measurements within CT-R. Finally, the lidar baseline denoted CT-LG consists of the continuous-time ICP from [15] with additional gyroscope constraints. The various parameters of

the baseline, as well as the radar-specific value of β , have been tuned using an extra sequence of each type.

3) *Results*: Table I shows the average odometry error (KITTI metrics) obtained with each baseline and our algorithm in the four environments. While a thorough ablation study is conducted in the following subsection, here, we consider three variations of DRO that all rely on the gyroscope integration for the orientation estimation. The first variant, denoted DRO-D, solely uses the Doppler-based velocity constraint O_d to estimate the linear dynamics of the system. The second one, DRO-G, uses solely the direct registration objective function O_ψ . Note that DRO-G still leverages Doppler compensation of the radar data as part of O_ψ . Ultimately, DRO-GD combines both O_d and O_ψ as presented in (8).

Overall, DRO-GD significantly outperforms every other radar methods. While the performance gap between DRO-GD and DRO-G is minimal in the well-structured suburban environment, the difference increases as the data becomes more challenging. The best radar baseline, CT-RDG, shows errors around twice as large as ours throughout all the sequences. When considering ‘Doppler-only’ approaches, DG [49] and DRO-D perform similarly in *Suburbs* and *Highway*. In the more challenging environments, DRO-D displays a noticeable advantage. All these observations suggest that the direct nature of the proposed approach (considering all the radar information in a single optimization without any feature extraction) and our novel GP-based infill for Doppler data provide a greater level of robustness. Without the Doppler-based velocity constraint, DRO-G is not able to reliably perform odometry in challenging scenarios, similar to CT-R. Despite not completely failing on all the *Tunnel* sequences (3 out of 4 runs present a translation error under 2.5%), the Doppler-compensation that happens in O_ψ is not sufficient for robust estimation in feature-deprived environments. Only the lidar baseline outperforms DRO-G in some feature-dense environments (*Suburbs*). However, CT-LG poorly performs in feature-deprived sequences (*Tunnel* and *Skyway*) while DRO-GD keeps a similar level of performance.


Visualizations of the raw data, velocity estimate and local map are provided in the supplementary materials. Fig. 7 provides a sample of the estimated trajectories for each sequence type. It is interesting to note that despite relatively similar KITTI metrics, CT-RDG and DRO-DG result in fairly different trajectories, with DRO-DG being much closer to the ground-truth. Part of the explanation is that the KITTI metrics only consider segments up to 800m where the small orientation drift of CT-RDG has a weak impact on the position at the end of each segment. However, along the full trajectory length, the orientation drift leads to large position errors that are visually noticeable. This observation supports the fact that for kilometre-long odometry, the direct integration of the gyroscope data can provide better orientation estimates than SOTA odometry.

4) *Ablation study*: We have conducted a thorough ablation study to determine which of the presented mechanisms impacts the algorithm’s accuracy the most. Due to the extremely


TABLE I

AVERAGE RELATIVE POSE ACCURACY OF THE PROPOSED METHOD AND SEVERAL BASELINES ON OUR DATASET (BEST RADAR-BASED METHOD IN BOLD).


Sequence type (length, avg. / max. vel.)	CT-R [15]	DG [49]	CT-RDG [49]	CT-LG [15]	DRO-D (ours)	DRO-G (ours)	DRO-GD (ours)
Suburbs (4×7.9 km, 8.1 / 18.6 m/s)	1.22 / 0.37	0.60 / 0.04	0.32 / 0.04	0.15 / 0.07	0.72 / 0.02	0.18 / 0.02	0.19 / 0.02
Highway (4×9.3 km, 11.1 / 27.0 m/s)	1.92 / 0.47	0.74 / 0.04	0.49 / 0.04	0.19 / 0.06	0.84 / 0.02	0.16 / 0.02	0.24 / 0.02
Tunnel (4×1.9 km, 8.9 / 28.2 m/s)	35.90 / 0.96	1.35 / 0.04	0.70 / 0.04	3.35 / 0.11	0.54 / 0.02	7.12 / 0.02	0.34 / 0.02
Skyway (4×11.1 km, 18.2 / 31.7 m/s)	35.24 / 1.77	1.17 / 0.02	0.86 / 0.02	16.98 / 0.24	0.71 / 0.01	45.76 / 0.01	0.40 / 0.01

KITTI odometry metric reported as XX / YY with XX [%] and YY [$^\circ$ /100 m] the translation and orientation errors, respectively.


(a) Suburbs



(b) Highway



(c) Tunnel



(d) Skyway

Fig. 7. Trajectory estimate samples for the different sequence types.

TABLE II
POSE ACCURACY ABLATION STUDY: EACH ROW CORRESPONDS TO A MARGINAL DIFFERENCE WITH DRO-GD.

Method	Suburbs	Highway	Tunnel	Vel. er.
DRO-GD (ours)	0.19 / 0.02	0.24 / 0.02	0.34 / 0.02	0.119
No gyr. bias	0.27 / 0.03	0.31 / 0.04	0.42 / 0.04	0.119
No vel. bias	0.22 / 0.02	0.29 / 0.02	0.31 / 0.02	0.120
No local map	0.35 / 0.02	0.43 / 0.02	0.40 / 0.02	0.107
No gyroscope	1.00 / 0.24	1.56 / 0.38	2.92 / 0.99	0.132
KITTI odometry metric reported as XX / YY with XX [%] and YY [°/100 m] the trans. and rot. errors, respectively.				Avg RMSE [m/s]

challenging nature of the Skyway sequences, we have chosen to leave them out of this study, as other factors, such as the high proportion of outliers, can have a greater impact on the results. Table (II) shows the relative pose error of the proposed DRO-GD when various elements are removed from the pipeline. For the ‘no local map’ row, DRO-GD is run replacing the proposed local map updated on the fly with simply the last motion and Doppler-corrected scan. As expected, removing the gyroscope² corresponds to the largest performance degradation, leading to accuracy slightly better than CT-R in Table I (also without a gyroscope) for feature-dense environments. Estimating the system’s orientation in a continuous manner using only radar data is challenging due to the low angular resolution of the sensor (0.9°) and the beam width (1.8°), even when leveraging a motion prior like CT-R. Interestingly, in more challenging scenarios like the *Tunnel* sequences, the drop in our method’s accuracy is not as large as for CT-R thanks to the proposed Doppler-based velocity constraint. Regardless, the velocity RMSE is not significantly impacted by the absence of a gyroscope. As for the other variations, the difference with DRO-GD is limited, and on average, our algorithm still outperforms the best radar baseline CT-RDG.

In Fig. 8 we study the sensitivity of DRO-GD with respect to the γ parameter used in the local map update. One can see that the method’s accuracy is not very sensitive to γ . We also analyze the estimated side velocity bias and empirical evidence that, given the assumption made in our formulation, it depends on the observed environment in Appendix B. Overall, this ablation study demonstrates that the proposed direct radar registration is sound and robust.

D. Public automotive benchmarks

We have benchmarked the proposed method against state-of-the-art frameworks over the Boreas and MulRan datasets. In both cases, the radar uses a sawtooth modulation pattern. Thus, none of the results presented in this section include the Doppler-based objective function O_d .

²For the gyroscope-less operation, we correct the estimated state with an angular velocity bias calibrated using held-off data ($0.034^\circ/\text{s}$ for *Suburbs* and *Highway*, and $0.057^\circ/\text{s}$ for *Tunnel*). We have also applied the robust weighting when the change in orientation between consecutive scans is not realistic. Without the bias correction, we obtain relative translation errors of 1.37%, 1.61%, and 3.67% for *Suburbs*, *Highway*, and *Tunnel*, respectively.




Fig. 8. Sensitivity analysis of DRO-GD with respect to the local map update parameter γ .

1) *Boreas dataset*: The Boreas dataset [16] was collected by driving a vehicle equipped with a sensor suite composed of a 3D lidar, a 2D spinning radar (sawtooth pattern), a front-facing camera, and an RTK-GNSS/INS solution for ground-truthing. The dataset consists of 44 sequences collected over one year along more than 350 km of roads in Toronto, Canada. Accordingly, the data corresponds to various weather conditions ranging from clear summer sky to heavy snow and rain. A specificity of Boreas is that most sequences follow a unique route in a suburban environment, allowing one to test an algorithm’s robustness with respect to changing conditions. The Boreas dataset also offers a leaderboard for odometry benchmarking that leverages 13 of the sequences above. Note that the GNSS/INS ground-truth is not publicly available for these sequences.

After testing our algorithm on 6 sequences with ground-truth, we submitted our results to the leaderboard with DRO-G. Our algorithm significantly outperforms all the other methods. We have reported the four best radar-based methods in Table III. The core principle of CFEAR [2] and STEAM-RIO++ [17] is similar as both methods extract a point cloud from the incoming radar data before performing registration to multiple previous scans. A difference is that CFEAR first performs motion distortion correction based on the previous estimate and uses point-to-distribution residuals, while STEAM-RIO++ leverages a GP motion prior, point-to-point constraints and IMU data. CFEAR++ [46] builds atop CFEAR by embedding the use of the IMU and some semantic information.

Our results confirm the ability of DRO-G to provide state-of-the-art performances without requiring a radar with a triangular modulation pattern and O_d as long as the environment contains sufficient geometric cues. In Appendix A, we provide a table that details the error obtained for each of the 13 sequences. As expected, there seems to be no correlation between the weather conditions and the translation error. An interesting observation is that the orientation estimates of DRO-G using solely the gyroscope integration outperform the other inertial-aided methods. This means that the radar constraints in CFEAR++ and STEAM-RIO++ play a negative role on the orientation estimates.

2) *MulRan dataset*: The MulRan dataset [40] has been recorded with a lidar, 2D spinning radar, and an IMU mounted atop a car while driving in 4 different locations in Daejeon

TABLE III
BOREAS DATASET SE(2) ODOMETRY LEADERBOARD.

Method	Trans. err. [%]	Rot. err. [°/100m]
DRO-G (ours)	0.26	0.05
CFEAR++ [46]	0.51	0.14
CFEAR [3]	0.61	0.21
STEAM-RIO++ [17]	0.62	0.18

and Sejong, South Korea. This dataset originally targeted the task of place recognition, but the availability of a 6-DoF trajectory ground-truth also enables the benchmarking of odometry and SLAM algorithms. However, unlike in the previous experiments, the ground-truth is not obtained with an RTK-GNSS but via the optimization of a pose graph that leverages wheel odometry, an optical-fibre gyroscope, and a VRS-GPS. Additionally, the quality of the IMU is significantly lower than in the Boreas and our automotive datasets.

Table IV shows the KITTI odometry errors of the proposed method and variations of CFEAR as reported in [3]. As introduced in the previous subsection, CFEAR relies on a scan-to-local-keyframes registration. C-FEAR-3 and C-FEAR-3-50 refer to [3] using a different number of past keyframes in the registration process: 4 and 50, respectively. The table shows the ability of DRO-G to outperform SOTA methods even with a low quality IMU. The level of rotational accuracy highlights the soundness of the proposed integration-only orientation and reiterates the observation from the Boreas leaderboard that radar data can have a detrimental impact on the system’s rotation estimate compared to pure integration and a simplistic bias estimation strategy.

Surprisingly, both CFEAR and our method display errors significantly larger with MulRan data than with the Boreas dataset. We believe that the nature of the ground-truth generation process leads to noisier poses, but also that extrinsic calibration and synchronization are suboptimal. For example, we observed that the difference between consecutive IMU data timestamps presents a variation of $\pm 35\%$ around the nominal period of 10ms and that there is a constant offset with respect to the other sensors. In our experiments, we corrected the IMU synchronization by removing 50ms to each inertial timestamp of all sequences. We also adjusted the radar heading by 0.172° (without the heading correction DRO-G leads to errors of 1.46% and $0.38^\circ/100\text{m}$). For completeness, without the gyroscope bias estimation, we obtain errors of 6.00% and $1.57^\circ/100\text{m}$ on average. Due to the lower quality of the dataset, we believe no strong conclusion can be drawn from this experiment other than the demonstration that DRO-G does not require a high-end gyroscope to perform at the level of SOTA methods.

E. Off-road navigation

To demonstrate the versatility of the proposed method, we collected a series of radar-inertial data sequences with an off-road mobile robot.

TABLE IV
AVERAGE ERROR PER SEQUENCE TYPE OF THE MULRAN DATASET.

Sequences	CFEAR-3 [3]	CFEAR-3-50 [3]	DRO-G (ours)
KAIST	1.65 / 0.70	1.53 / 0.68	1.44 / 0.42
DCC	1.75 / 0.49	1.58 / 0.50	1.54 / 0.41
Riverside	1.46 / 0.51	1.39 / 0.50	1.38 / 0.31
Average	1.62 / 0.57	1.50 / 0.56	1.45 / 0.38

KITTI odometry metric reported as XX / YY with XX [%] and YY [°/100 m] the translation and orientation errors, respectively.

1) *Dataset description*: Our off-road dataset comprises 16 test sequences divided into four locations ranging progressively from highly structured to unstructured environments. The robot’s maximum velocity is approximately 1.2 m/s. The sensor suite comprises a Navtech RAS3 radar and an Ouster lidar with its embedded IMU. The robot is driven over the same loop three or four times for each environment. The easiest one, *Parking*, takes place in a parking lot with many artificial geometric features such as buildings, fences, and parked cars. The second location *Industrial* takes place around a large industrial building with part of the trajectory going over a grass patch. This environment contains a combination of geometric structures (fences and buildings) and sparse open areas (trees and bushes) at the back of the industrial building. For the third trajectory, *Grass*, the robot moves over a grass field and occasionally can observe a metallic fence. The last and most challenging environment, *Woods*, is entirely off-road in the woods. Fig. 1 (right) shows our Clearpath Warthog robot within the *Woods* environment.

2) *Baseline*: In this experiment, we benchmark DRO-G against a version of CT-R [15] with additional inertial constraints as in [54]. We denote this ICP-based baseline as CT-RG. For each CT-RG run, the gyroscope measurements have been corrected with a constant bias estimated from five seconds of data where the robot is static. CT-RG’s parameters have been tuned using held-out *Parking* sequences.

3) *Results*: We show the RPE obtained with and without gyroscope bias estimation in Table V. In structured environments (*Parking* and *Industrial*), both the direct and point-cloud-based methods perform similarly. The discrepancy between the two methods increases when dealing with more challenging environments. Overall, DRO-G maintains a satisfactory level of accuracy across all scenarios, demonstrating that the level of structure of the environment does not seem to impact the accuracy significantly. It also shows that direct radar registration is robust to the nature of the environment without requiring any tedious parameter tuning. Fig. 9 provides a visualization of estimated trajectories in each environment. Table V and Fig. 9 also show the efficacy of the proposed gyroscope bias estimation (DRO-G no-bias-est. vs DRO-G). Unlike the automotive scenario, where the gyroscope bias is nearly negligible, the lower quality of the lidar’s IMU has a significant impact on the results if the bias is not accounted for. Both methods perform poorly without any bias correction.




Fig. 9. Visualization of trajectory estimates in the off-road scenarios (aligned with the ground-truth for visualization using the first 20% of the trajectory).

TABLE V
AVERAGE RPE IN VARIOUS ENVIRONMENTS SPANNING OVER DIFFERENT LEVELS OF STRUCTURE.

Sequence type	CT-RG[54] no-bias-est	CT-RG[54]	DRO-G no-bias-est.	DRO-G (ours)
Parking (4×368 m)	3.45	1.03	3.44	1.04
Industrial (4×390 m)	3.63	0.77	3.57	0.62
Grass (4×175 m)	5.63	1.05	5.29	0.67
Woods (4×421 m)	4.70	1.30	3.42	0.71

Relative position error as a percentage [%] of the distance travelled.

F. Computation time

All the experiments in this paper were performed using a laptop equipped with an Intel i7-13850HX CPU and an Nvidia RTX 5000 Mobile GPU. On our automotive dataset, the computation of DRO-GD took on average 89 ms per radar frame using the GPU. This is far below the real-time limit of 250 ms (4 Hz radar). When considering solely the Doppler-based velocity objective function, DRO-D required only 34 ms per scan. For comparison, we also changed the target device to CPU in Pytorch. We obtained averages of 738 ms and 59 ms for DRO-GD and DRO-D, respectively. On the Boreas and MulRan datasets the GPU computations needed 102 ms and 95 ms per scan, respectively. The off-road experiments were the ones that required the most computation time with 134 ms per frame. To summarise, all our experiments have run in real-time. However, the current implementation is somewhat modular, allowing for different combinations of objective function and motion models. We believe that specializing the

code and leveraging a compiled programming language would lead to significant improvement in the proposed method's efficiency especially for CPU-based operations.

VI. LIMITATIONS

Currently, the proposed framework strongly relies on the gyroscope data as the orientation is obtained via direct integration of the angular velocity. We demonstrated that it was sufficient to obtain SOTA performance in both automotive and off-road scenarios with high-quality and low-quality IMUs. However, significant degradation of the inertial data (saturation, long dropouts, erroneous measurements, etc), while unlikely, would have a catastrophic impact on the estimated state. On a similar topic, the proposed approach to estimate the gyroscope bias requires the vehicle/robot to be static occasionally. Further work is needed to introduce a principled approach to gyroscope bias estimation in DRO.

Our method and experiments focused on spinning FMCW radars. Thus, our findings are limited to this particular type of sensor. However, we believe that DRO applies to other radar types (e.g. phased-arrays) both in 2D and 3D given a sufficient FoV and access to the intensity returns (unfortunately most off-the-shelf automotive radars do not provide this information). A practical limitation of using 3D data is the potential computational load increase if the radar provides much more information than the radars used in this work.

VII. CONCLUSION

We have introduced DRO: a direct method for radar-based odometry aided by a gyroscope. Unlike most of the top-

performing methods in the literature, DRO does not extract point clouds or features in the incoming radar scans but directly leverages the intensity information of the data to perform scan-to-local-map registration. The orientation is obtained by integrating the gyroscope measurements, and the body-centric velocity is estimated via the maximization of one or two objective functions depending on the radar's frequency modulation. With a triangular emission pattern, the Doppler-induced shifts in the data make the radial velocities observable, thus enabling odometry in geometrically challenging environments such as featureless tunnels. Regardless of the radar's signal shape, DRO is the first direct method that accounts in a principled way for both the Doppler-induced and motion distortion of the radar data in the estimation process. The real-time GPU-based implementation of the proposed method was shown to outperform existing methods on public benchmarks and datasets we have collected both in automotive and off-road environments. Future works include the integration of a principled way to estimate the gyroscope bias and perform SLAM by estimating the full trajectory and including loop-closure constraints.

ACKNOWLEDGMENTS

This paper was partially supported by an Ontario Research Fund - Research Excellence grant. The authors would like to thank Aoran Jiao for his help in collecting the off-road datasets.

REFERENCES

- [1] Nader J. Abu-Alrub and Nathir A. Rawashdeh. Radar odometry for autonomous ground vehicles: A survey of methods and datasets. *IEEE Trans. on Intelligent Vehicles (T-IV)*, 9(3):4275–4291, 2024. doi: 10.1109/TIV.2023.3340513.
- [2] Daniel Adolfsson, Martin Magnusson, Anas Alhashimi, Achim J. Lilienthal, and Henrik Andreasson. Cfar radarodometry - conservative filtering for efficient and accurate radar odometry. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 5462–5469, 2021. doi: 10.1109/IROS51168.2021.9636253.
- [3] Daniel Adolfsson, Martin Magnusson, Anas Alhashimi, Achim J. Lilienthal, and Henrik Andreasson. Lidar-level localization with radar? the cfar approach to accurate, fast, and robust large-scale radar odometry in diverse environments. *IEEE Trans. on Robotics (TRO)*, 39(2): 1476–1495, 2023. doi: 10.1109/TRO.2022.3221302.
- [4] Roberto Aldera, Daniele De Martini, Matthew Gadd, and Paul Newman. Fast radar motion estimation with a learnt focus of attention using weak supervision. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pages 1190–1196, 2019. doi: 10.1109/ICRA.2019.8794014.
- [5] Sean Anderson and Timothy D Barfoot. Ransac for motion-distorted 3d visual sensors. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 2093–2099. IEEE, 2013.
- [6] Sean Anderson and Timothy D. Barfoot. Full steam ahead: Exactly sparse gaussian process regression for batch continuous-time trajectory estimation on se(3). In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 157–164, 2015. doi: 10.1109/IROS.2015.7353368.
- [7] Sean Anderson, Timothy D. Barfoot, Chi Hay Tong, and Simo Särkkä. Batch nonlinear continuous-time trajectory estimation as exactly sparse gaussian process regression. *Autonomous Robots*, 39(3):221–238, 2015. ISSN 0929-5593. doi: 10.1007/s10514-015-9455-y.
- [8] Timothy D Barfoot, James R Forbes, and David J Yoon. Exactly sparse gaussian variational inference with application to derivative-free batch nonlinear state estimation. *Intl. Journal of Robotics Research (IJRR)*, 39(13):1473–1502, 2020. doi: 10.1177/0278364920937608.
- [9] Dan Barnes and Ingmar Posner. Under the radar: Learning to predict robust keypoints for odometry estimation and metric localisation in radar. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2020.
- [10] Dan Barnes, Rob Weston, and Ingmar Posner. Masking by Moving: Learning Distraction-Free Radar Odometry from Pose Information. In *Conference on Robot Learning (CoRL)*, 2019. URL <https://arxiv.org/pdf/1909.03752>.
- [11] S. Blake. Os-cfar theory for multiple targets and nonuniform clutter. *IEEE Trans. on Aerospace and Electronic Systems (T-AES)*, 24(6):785–790, 1988. doi: 10.1109/7.18645.
- [12] Michael Bosse and Robert Zlot. Map matching and data association for large-scale two-dimensional laser scan-based slam. *Intl. Journal of Robotics Research (IJRR)*, 27:667–691, 06 2008. doi: 10.1177/0278364908091366.
- [13] K. Burnett, Angela P. Schoellig, and Timothy D. Barfoot. Do we need to compensate for motion distortion and doppler effects in spinning radar navigation? *IEEE Robotics and Automation Letters (RA-L)*, 6(2):771–778, 2021. doi: 10.1109/lra.2021.3052439. URL <https://doi.org/10.1109/lra.2021.3052439>.
- [14] Keenan Burnett, David J Yoon, Angela P Schoellig, and Timothy D Barfoot. Radar Odometry Combining Probabilistic Estimation and Unsupervised Feature Learning. In *Proc. of Robotics: Science and Systems (RSS)*, 2021.
- [15] Keenan Burnett, Yuchen Wu, David J Yoon, Angela P Schoellig, and Timothy D Barfoot. Are We Ready for Radar to Replace Lidar in All-weather Mapping and Localization? *IEEE Robotics and Automation Letters (RA-L)*, 7(4):10328–10335, 2022.
- [16] Keenan Burnett, David J Yoon, Yuchen Wu, Andrew Z Li, Haowei Zhang, Shichen Lu, Jingxing Qian, Wei-Kang Tseng, Andrew Lambert, Keith YK Leung, Angela P Schoellig, and Timothy D Barfoot. Boreas: A multi-season autonomous driving dataset. *Intl. Journal of Robotics Research (IJRR)*, 42(1-2):33–42, 2023. doi: 10.1177/02783649231160195.

- [17] Keenan Burnett, Angela P. Schoellig, and Timothy D. Barfoot. Continuous-time radar-inertial and lidar-inertial odometry using a gaussian process motion prior. *IEEE Trans. on Robotics (TRO)*, 2024.
- [18] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J.J. Leonard. Past, present, and future of simultaneous localization and mapping: Towards the robust-perception age. *IEEE Trans. on Robotics (TRO)*, 32(6):1309–1332, 2016.
- [19] Jonas Callmer, David Törnqvist, Fredrik Gustafsson, Henrik Svensson, and Pelle Carlbom. Radar slam using visual features. *EURASIP Journal on Advances in Signal Processing*, 2011, 12 2011. doi: 10.1186/1687-6180-2011-71.
- [20] Sarah H. Cen and Paul Newman. Precise ego-motion estimation with millimeter-wave radar under diverse and challenging conditions. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pages 6045–6052, 2018. doi: 10.1109/ICRA.2018.8460687.
- [21] Sarah H. Cen and Paul Newman. Radar-only ego-motion estimation in difficult settings via graph matching. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pages 298–304, 2019. doi: 10.1109/ICRA.2019.8793990.
- [22] Paul Checchin, Franck Gérossier, Christophe Blanc, Roland Chapuis, and Laurent Trassoudaine. Radar scan matching slam using the fourier-mellin transform. In *Springer Tracts in Advanced Robotics*, volume 62, pages 151–161, 01 2009. ISBN 978-3-642-13407-4. doi: 10.1007/978-3-642-13408-1_14.
- [23] Bartolomeo Della Corte, Igor Bogoslavskyi, C. Stachniss, and Giorgio Grisetti. A general framework for flexible multi-cue photometric point cloud registration. *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pages 1–8, 2017. URL <https://api.semanticscholar.org/CorpusID:10485366>.
- [24] Amaël Delaunoy and Marc Pollefeys. Photometric bundle adjustment for dense multi-view 3d modeling. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1486–1493, 2014. doi: 10.1109/CVPR.2014.193.
- [25] Luca Di Giammarino, Leonardo Brizi, Tiziano Guadagnino, Cyrill Stachniss, and Giorgio Grisetti. Md-slam: Multi-cue direct slam. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 11047–11054, 2022. doi: 10.1109/IROS47612.2022.9981147.
- [26] M.W.M.G. Dissanayake, P. Newman, S. Clark, H.F. Durrant-Whyte, and M. Csorba. A solution to the simultaneous localization and map building (slam) problem. *IEEE Trans. on Robotics and Automation*, 17(3):229–241, 2001. doi: 10.1109/70.938381.
- [27] Jakob Engel, Vladlen Koltun, and Daniel Cremers. Direct sparse odometry. *IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI)*, PP, 2016. doi: 10.1109/TPAMI.2017.2658577.
- [28] Jakob J. Engel, Thomas Schöps, and Daniel Cremers. Lsd-slam: Large-scale direct monocular slam. In *Proc. of the Europ. Conf. on Computer Vision (ECCV)*, 2014. URL <https://api.semanticscholar.org/CorpusID:14547347>.
- [29] Harold M. Finn. Adaptive detection mode with threshold control as a function of spatially sampled clutter level estimates. *Rca Rev.*, 29:414–465, 1968. URL <https://api.semanticscholar.org/CorpusID:208092087>.
- [30] Christian Forster, Matia Pizzoli, and Davide Scaramuzza. Svo: Fast semi-direct monocular visual odometry. *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pages 15–22, 2014. URL <https://api.semanticscholar.org/CorpusID:206850490>.
- [31] Paul Furgale, Chi Hay Tong, Timothy D. Barfoot, and Gabe Sibley. Continuous-time batch trajectory estimation using temporal basis functions. *Intl. Journal of Robotics Research (IJRR)*, 34(14):1688–1710, 2015. doi: 10.1177/0278364915585860. URL <https://doi.org/10.1177/0278364915585860>.
- [32] Andres Galeote-Luque, Vladimír Kubelka, Martin Magnusson, Jose-Raul Ruiz-Sarmiento, and Javier Gonzalez-Jimenez. Doppler-only single-scan 3d vehicle odometry. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pages 13703–13709, 2024. doi: 10.1109/ICRA57147.2024.10611199.
- [33] P.P. Gandhi and S.A. Kassam. Analysis of cfar processors in nonhomogeneous background. *IEEE Trans. on Aerospace and Electronic Systems (T-AES)*, 24(4):427–445, 1988. doi: 10.1109/7.7185.
- [34] Reza Ghabcheloo and Shadman Siddiqui. Complete odometry estimation of a vehicle using single automotive radar and a gyroscope. In *Mediterranean Conference on Control and Automation (MED)*, pages 855–860, 2018. doi: 10.1109/MED.2018.8442474.
- [35] Luca Giammarino, Emanuele Giacomini, Leonardo Brizi, Omar Salem, and Giorgio Grisetti. Photometric lidar and rgb-d bundle adjustment. *IEEE Robotics and Automation Letters (RA-L)*, PP:1–8, 07 2023. doi: 10.1109/LRA.2023.3281907.
- [36] Mona Gridseth and Timothy Barfoot. Towards direct localization for visual teach and repeat. In *16th Conference on Computer and Robot Vision (CRV)*, pages 97–104, 2019. doi: 10.1109/CRV.2019.00021.
- [37] Kyle Harlow, Hyesu Jang, Timothy D. Barfoot, Ayoun Kim, and Christoffer Heckman. A new wave in robotics: Survey on recent mmwave radar applications in robotics. *IEEE Trans. on Robotics (TRO)*, 40:4544–4560, 2024. doi: 10.1109/TRO.2024.3463504.
- [38] Johan Hedborg, Per-Erik Forssén, Michael Felsberg, and Erik Ringaby. Rolling shutter bundle adjustment. *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1434–1441, 2012. URL <https://api.semanticscholar.org/CorpusID:15591674>.
- [39] Dominik Kellner, Michael Barjenbruch, Jens Klapstein, Jürgen Dickmann, and Klaus Dietmayer. Instantaneous

- ego-motion estimation using doppler radar. In *Intl. IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 869–874, 2013. doi: 10.1109/ITSC.2013.6728341.
- [40] Giseop Kim, Yeong Sang Park, Younghun Cho, Jinyong Jeong, and Ayoung Kim. Mulran: Multimodal range dataset for urban place recognition. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pages 6246–6253, 2020. doi: 10.1109/ICRA40945.2020.9197298.
- [41] Vladimír Kubelka, Emil Fritz, and Martin Magnusson. Do we need scan-matching in radar odometry? In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pages 13710–13716, 2024. doi: 10.1109/ICRA57147.2024.10610666.
- [42] Pou-Chun Kung, C. Wang, and Wen-Chieh Lin. A normal distribution transform-based radar odometry designed for scanning and automotive radars. *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pages 14417–14423, 2021. URL <https://api.semanticscholar.org/CorpusID:232233173>.
- [43] Cedric Le Gentil and Teresa Vidal-Calleja. Continuous latent state preintegration for inertial-aided systems. *Intl. Journal of Robotics Research (IJRR)*, 42(10):874–900, 2023. doi: 10.1177/02783649231199537. URL <https://doi.org/10.1177/02783649231199537>.
- [44] Cedric Le Gentil, Teresa Vidal-Calleja, and Shoudong Huang. Gaussian process preintegration for inertial-aided state estimation. *IEEE Robotics and Automation Letters (RA-L)*, 5(2):2108–2114, 2020. doi: 10.1109/LRA.2020.2970940.
- [45] Cedric Le Gentil, Raphael Falque, and Teresa Vidal-Calleja. Real-time truly-coupled lidar-inertial motion correction and spatiotemporal dynamic object detection. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 12565–12572, 2024. doi: 10.1109/IROS58592.2024.10802233.
- [46] Siru Li, Yushuai Chen, Ziyang Hong, and Liang Hu. Cfear++: An efficient radar-inertial semantic odometry. *ICRA Workshop, Radar in Robotics: Resilience from Signal to Navigation*, 2024.
- [47] Hyungtae Lim, Kawon Han, Gunhee Shin, Giseop Kim, Songcheol Hong, and Hyun Myung. ORORA: Outlier-robust Radar Odometry. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pages 2046–2053, 2023. doi: 10.1109/ICRA48891.2023.10160997.
- [48] Daniil Lisus, Johann Laconte, Keenan Burnett, and Tim D. Barfoot. Pointing the way: Refining radar-lidar localization using learned icp weights. *ArXiv*, abs/2309.08731, 2023. URL <https://api.semanticscholar.org/CorpusID:262043844>.
- [49] Daniil Lisus, Keenan Burnett, David J. Yoon, Richard Poulton, John Marshall, and Timothy D. Barfoot. Are doppler velocity measurements useful for spinning radar odometry? *IEEE Robotics and Automation Letters (RA-L)*, 10(1):224–231, 2025. doi: 10.1109/LRA.2024.3505821.
- [50] David G. Lowe. Distinctive image features from scale-invariant keypoints. *Intl. Journal of Computer Vision (IJCV)*, 60:91–110, 2004. doi: 10.1023/B:VISI.0000029664.99615.94. URL <http://dx.doi.org/10.1023/B:VISI.0000029664.99615.94>.
- [51] José Machado Fernández, Norelys Mojena-Hernández, and Jesús Bacallao Vidal. Evaluation of cfar detectors performance. *ITECKNE Innovación e Investigación en Ingeniería*, 14:170, 11 2017. doi: 10.15332/iteckne.v14i2.1772.
- [52] Yeong Sang Park, Young-Sik Shin, and Ayoung Kim. Pharao: Direct radar odometry using phase correlation. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pages 2617–2623, 2020. doi: 10.1109/ICRA40945.2020.9197231.
- [53] Elliot Preston-Krebs, Daniil Lisus, and Timothy Barfoot. The finer points: A systematic comparison of point-cloud extractors for radar odometry. *arXiv preprint arXiv:2409.10491*, 09 2024. doi: 10.48550/arXiv.2409.12256.
- [54] Xinyuan Qiao, Alexander Krawciw, Sven Lilge, and Timothy D Barfoot. Radar teach and repeat: Architecture and initial field testing. *arXiv preprint arXiv:2409.10491*, 2024.
- [55] C E Rasmussen and C K I Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- [56] Fraser Rennie, David Williams, Paul Newman, and Daniele De Martini. Doppler-Aware Odometry from FMCW Scanning Radar. In *Intl. IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 5126–5132, 2023.
- [57] Hermann Rohling. Radar cfar thresholding in clutter and multiple target situations. *IEEE Trans. on Aerospace and Electronic Systems (T-AES)*, AES-19(4):608–621, 1983. doi: 10.1109/TAES.1983.309350.
- [58] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. Orb: An efficient alternative to sift or surf. In *Proc. of the IEEE Intl. Conf. on Computer Vision (ICCV)*, pages 2564–2571, 2011. doi: 10.1109/ICCV.2011.6126544.
- [59] Omar Ashraf Ahmed Khairy Salem, Emanuele Giacomini, Leonardo Brizi, Luca Di Giammarino, and Giorgio Grisetti. Enhancing lidar performance: Robust deskewing exclusively relying on range measurements. In *AIxIA 2023 - Advances in Artificial Intelligence*, Berlin, Heidelberg, 2023. Springer-Verlag. ISBN 978-3-031-47545-0. doi: 10.1007/978-3-031-47546-7_21. URL https://doi.org/10.1007/978-3-031-47546-7_21.
- [60] M.E. Smith and P. K. Varshney. Vi-cfar: a novel cfar algorithm based on data variability. *Proceedings of the 1997 IEEE National Radar Conference*, pages 263–268, 1997. URL <https://api.semanticscholar.org/CorpusID:140690935>.
- [61] Arthur Venon, Yohan Dupuis, Pascal Vasseur, and Pierre Merriaux. Millimeter wave fmcw radars for perception,

recognition and localization in automotive applications: A survey. *IEEE Trans. on Intelligent Vehicles (T-IV)*, 7: 1–1, 09 2022. doi: 10.1109/TIV.2022.3167733.

- [62] Rob Weston, Matthew Gadd, Daniele De Martini, Paul Newman, and Ingmar Posner. Fast-mbym: Leveraging translational invariance of the fourier transform for efficient and accurate radar odometry. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, pages 2186–2192, 2022. doi: 10.1109/ICRA46639.2022.9812063.
- [63] Renli Zhang, Weixing Sheng, and Xiaofeng Ma. Improved switching cfar detector for non-homogeneous environments. *Signal Processing*, page 35–48, 2013. doi: 10.1016/j.sigpro.2012.06.015.
- [64] Zichao Zhang and Davide Scaramuzza. A tutorial on quantitative trajectory evaluation for visual(-inertial) odometry. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2018.
- [65] Lei Zhao, Weixian Liu, Xin Wu, and Jeffrey S Fu. A novel approach for cfar processors design. In *Proceedings of the 2001 IEEE Radar Conference*, pages 284–288. IEEE, 2001.

APPENDIX A

RESULTS DETAILS ON BOREAS DATASET

This appendix presents the detailed results obtained on the 13 sequences of the Boreas dataset leaderboard in Table VI.

TABLE VI
PER-SEQUENCE ERRORS OF DRO-G ON THE BOREAS DATASET
LEADERBOARD.

Sequence ID	Weather conditions	Trans. err. [%]	Rot. err. [°/100m]
2020-12-04-14-00	☁☼	0.24	0.05
2021-01-26-10-59	☁☼☼	0.20	0.04
2021-02-09-12-55	☁☼	0.22	0.02
2021-03-09-14-23	☼	0.23	0.06
2021-04-22-15-00	☁	0.25	0.05
2021-06-29-18-53	☁☼	0.36	0.02
2021-06-29-20-43	☁☼	0.44	0.04
2021-09-08-21-00	●	0.30	0.02
2021-09-09-15-28	☼	0.53	0.19
2021-10-05-15-35	☁	0.12	0.02
2021-10-26-12-35	☁	0.13	0.01
2021-11-06-18-55	●	0.32	0.10
2021-11-28-09-18	☁☼☼	0.10	0.01

☁: Overcast, ☼: Snow coverage, ☼☼: High snow coverage
☼☼☼: Snowing, ☼: Sun, ☁☼: Rain, ☼: Dusk, ●: Night

APPENDIX B

LATERAL VELOCITY BIAS ESTIMATE

This appendix discusses the lateral velocity bias estimation introduced in Section V-A4. Fig. 10 presents the lateral velocity estimates obtained across the different sequences used in Section V-C. The value of the bias is shown as a function of the vehicle’s position along the trajectory. Note that the *Highway*




Fig. 10. Lateral velocity bias estimate obtained by running DRO-GD on our automotive dataset. The bias is expressed as a function of the distance travelled in the sequence.

and *Tunnel* sequences are split into two plots as they have been collected separately in one direction and the other. As one can see, there is a clear correlation between the location of the vehicle and the value of the velocity bias estimate. This observation supports our hypothesis that the Doppler-based velocity bias is impacted by the structure of the environment. Note that the extrinsic calibration between the radar and the vehicle rear axle has been thoroughly verified, and thus, the miscalibration hypothesis has been discarded. Further work will be required to obtain a deeper characterization of this bias and a clearer understanding of the physical phenomenon behind it.