

Extrinsic Calibration of 2D Millimetre-Wavelength Radar Pairs Using Ego-Velocity Estimates

Qilong Cheng[†], Emmett Wise[†], and Jonathan Kelly

Abstract—Correct radar data fusion depends on knowledge of the spatial transform between sensor pairs. Current methods for determining this transform operate by aligning identifiable features in different radar scans, or by relying on measurements from another, more accurate sensor. Feature-based alignment requires the sensors to have overlapping fields of view or necessitates the construction of an environment map. Several existing techniques require bespoke retroreflective radar targets. These requirements limit both where and how calibration can be performed. In this paper, we take a different approach: instead of attempting to track targets or features, we rely on ego-velocity estimates from each radar to perform calibration. Our method enables calibration of a subset of the transform parameters, including the yaw and the axis of translation between the radar pair, without the need for a shared field of view or for specialized targets. In general, the yaw and the axis of translation are the most important parameters for data fusion, the most likely to vary over time, and the most difficult to calibrate manually. We formulate calibration as a batch optimization problem, show that the radar-radar system is identifiable, and specify the platform excitation requirements. Through simulation studies and real-world experiments, we establish that our method is more reliable and accurate than state-of-the-art methods. Finally, we demonstrate that the full rigid body transform can be recovered if relatively coarse information about the platform rotation rate is available.

Index Terms—Calibration & Identification, Radar, Robot Sensing Systems, Sensor Fusion

I. INTRODUCTION

Millimetre-wavelength radar has proven to be a valuable sensing modality for autonomous vehicles (AVs) because radar is relatively robust to inclement weather and is able to provide velocity information [1]. However, the field of view of most radar sensors is limited and radars are known to produce noisy range and range-rate measurements. To provide full coverage of the environment and to ensure redundancy, AVs often aggregate data from multiple radars. Data fusion, in turn, requires accurate knowledge of the spatial transform(s) between the sensors. The process of determining the sensor-to-sensor transform is known as extrinsic calibration. While the sensors on board many AVs are factory-calibrated prior to deployment, the spatial transform may change during operation for a variety of reasons (e.g., collisions, wear and tear, etc.). An ability to perform online, in-situ extrinsic calibration is therefore important for safety and reliability.

[†] denotes equal contribution.

All authors are with the Space & Terrestrial Autonomous Robotics Systems Laboratory at the University of Toronto Institute for Aerospace Studies, Toronto, Canada. Jonathan Kelly is a Vector Institute Faculty Affiliate. <first>. <last>@robotics.utias.utoronto.ca




Fig. 1: Illustration of our 2D radar-to-radar extrinsic calibration problem. Radars a and b measure the range, azimuth, and range-rate to targets in the environment. Assuming that the targets are stationary in a fixed, world reference frame, we can estimate the ego-velocity of each radar. Our algorithm fuses the ego-velocity estimates from radars a and b to estimate the rotation, \mathbf{R} , and translation direction θ_t between the two sensors.

Existing extrinsic calibration methods for 2D radar pairs either align clouds of 3D points derived from radar measurements [2] or involve other (non-radar) sensors that have better accuracy (e.g., lidar units) [3]–[5]. The method in Olutomilayo et al. [2] relies on specialized trihedral radar retroreflectors to ensure sufficient environmental structure for calibration and to simplify the data association problem, for example. However, these retroreflectors are bespoke and are unavailable outside of the laboratory, limiting the possibility for calibration in the field, during operation.

In this paper, we study the radar-to-radar extrinsic calibration problem and develop a calibration approach that does not require specialized targets or other sensors. To avoid the challenges of data association, we instead propose a method that relies on instantaneous ego-velocity estimates from each radar only, as shown in Figure 1. More specifically, we determine the rotation angle and translation axis (i.e., unit vector) between the sensors. Radar data fusion, especially radar ego-velocity fusion, is most sensitive to these parameters. Under nominal operating conditions, the distance

between the sensors is unlikely to change appreciably from specifications. In contrast, the orientation of each radar can easily be altered (e.g., by a minor impact) and is very difficult to measure by hand. To the best of the authors’ knowledge, this is the first work to explore extrinsic calibration of sensor pairs that provide velocity estimates only. That is, no direct information about the rotation or rotation rate of the sensor platform is considered (instead, rotation must be inferred from the velocities). Herein, we:

- 1) show that the yaw angle and the direction of translation between pairs of coplanar 2D radar units can be determined from ego-velocity estimates only;
- 2) formulate a batch solver for the calibration problem and prove that the parameters are identifiable, given sufficient excitation;
- 3) carry out simulation studies to analyze the sensitivity of our method to varying levels of measurement noise;
- 4) confirm via real-world experiments that our approach is more accurate and reliable than two state-of-the-art methods; and
- 5) demonstrate that the full spatial transform can be recovered when an additional, coarse source of information about the platform rotation rate is available.

The remainder of the paper is organized as follows. In Section II, we review existing radar extrinsic calibration algorithms. Section III formulates our batch optimization problem. We prove that the calibration parameters are identifiable and establish the necessary trajectory excitation requirements in Section IV. In Section V, we show, through simulations and real-world experiments, that our algorithm is more reliable and accurate at estimating the yaw angle and the translation axis between radar pairs than two state-of-the-art-methods. Finally, we summarize our work and discuss future research directions in Section VI.

II. RELATED WORK

In this section, we survey pairwise extrinsic calibration methods where one (or both) of the sensors in the pair is a mm-wavelength radar. Sections II-A and II-B detail target-based and target-free extrinsic calibration algorithms, respectively, that rely on feature detection and matching. In Section II-C, we review extrinsic calibration methods that relate the instantaneous radar ego-velocity to the motion of the second sensor.

A. Target-Based Methods

Most existing radar-camera extrinsic calibration algorithms estimate the projective transform (homography) between the horizontal 2D radar sensing plane and the camera image plane. Due to the sparse and noisy nature of radar measurements, these methods rely on specialized trihedral retroreflectors that produce point-like detections in the radar and camera data while simplifying the correspondence problem [6]–[9]. Additionally, although 2D radars are incapable of estimating elevation, the sensors do often detect off-plane targets, and these detections bias the homography estimate. Since signal returns from targets on the radar horizontal plane

are stronger than those from off-plane targets, Sugimoto et al. [6] filter on-plane targets by maximizing the measured radar cross section (RCS) of the targets.

The most common error metric for radar-to-sensor extrinsic calibration is a form of ‘reprojection’ error, which defines the misalignment between identifiable targets (objects) viewed by both sensors. For example, Olutomilayo et al. [2] estimate the 2D transform between the radar frame and a vehicle coordinate frame by aligning radar measurements of stationary retroreflectors with a known map in the vehicle coordinate frame. The approaches of El Natour et al. [10], Domhoff et al. [11], and Peršić et al. [12] treat all radar measurements as lying on spherical arcs with constant range and azimuth (i.e., the measurements vary only in elevation). To estimate the 3D transform between sensor pairs, the arcs are aligned with the measurements from a second sensor. In order to account for the elevations of the retroreflector targets relative to the horizontal radar sensing plane, these methods introduce additional calibration constraints by designing specific target arrangements [10], [11] or explicitly modelling the radar-target interactions [12]. All of these techniques require the sensor pair to simultaneously view one or more specialized targets, so the sensors must share overlapping fields of view. Our approach does not require specialized infrastructure or a shared field of view, allowing for calibration of a wider range of sensor configurations and in more environments.

B. Target-Free Methods

In contrast to methods that rely on specialized radar retroreflectors, target-free or ‘targetless’ algorithms estimate the radar-to-sensor transform by aligning identifiable environment features (observed by both sensors). Schöller et al. [13] train a neural network to correct an inaccurate rotation estimate between a 2D radar-camera pair using raw camera and radar vehicle detection data. Peršić et al. [14] align tracked objects to determine the yaw angle between 2D radar-camera and radar-lidar pairs. Due to the challenge of tracking environmental features consistently across radar scans, these methods only calibrate the rotation between the sensors. Burnett et al. [3] estimate the transform between a 2D radar-lidar pair, where both sensors have a 360° field of view. The method in [3] aligns measured radar and lidar point clouds, which requires a large number of jointly-observed features. Heng et al. [15] estimate the extrinsic calibration parameters between 3D radar-lidar pairs by constructing a lidar point cloud map and localizing the 3D radar units within the map. While this approach could possibly be used for 2D radars, the method requires the known poses of the vehicle and the construction of dense map. Our approach does not require tracking of environmental features, which simplifies the calibration process.

C. Ego-Velocity Methods

Ego-velocity methods estimate the transform by minimizing the error between radar ego-velocity estimates and the motion of another sensor [1], [16]. By minimizing the lateral

velocity error between a radar and an inertial measurement unit (IMU), Kellner et al. [17] estimate the rotation angle between a radar-IMU pair, but their scheme requires accurate knowledge of the translation between the sensors. Doer et al. [5] and Wise et al. [4] extend this approach to 3D radar-IMU and radar-camera extrinsic calibration. To date, each ego-velocity method relies on one sensor that provides rotation information (e.g., angular velocity or $SO(2)$ measurements relative to an inertial frame). Herein, we determine the subset of extrinsic calibration parameters that are identifiable without angular velocity measurements.

III. PROBLEM FORMULATION

A. Notation

In this paper, Latin and Greek letters, such as a and α , represent scalar variables. Lowercase (e.g., \mathbf{h} and $\boldsymbol{\sigma}$) and uppercase (e.g., $\boldsymbol{\Theta}$ and \mathbf{C}) boldface characters are reserved for vectors and matrices, respectively. A Cartesian reference frame is identified by $\underline{\mathcal{F}}_{\rightarrow}$. The translation vector from $\underline{\mathcal{F}}_{\rightarrow a}$ to $\underline{\mathcal{F}}_{\rightarrow b}$, expressed in $\underline{\mathcal{F}}_{\rightarrow a}$, is denoted by \mathbf{t}_a^{ba} . The function $\mathbf{R}(\theta)$ maps $\theta \in \mathbb{R}$ to an element of $SO(2)$; for example, $\mathbf{R}(\theta_{ab})$ defines the rotation from $\underline{\mathcal{F}}_{\rightarrow b}$ to $\underline{\mathcal{F}}_{\rightarrow a}$. We use \mathbf{I}_n to denote the n -by- n identity matrix. The operator \times is the cross product operator. The unary operator \wedge acts on $r \in \mathbb{R}$ to produce

$$r^\wedge = \begin{bmatrix} 0 & -r \\ r & 0 \end{bmatrix}. \quad (1)$$

B. Radar Ego-Velocity Estimation

Let the static world and moving radar frames be $\underline{\mathcal{F}}_{\rightarrow w}$ and $\underline{\mathcal{F}}_{\rightarrow r}$, respectively. At each time index j , the radar detects N stationary, environmental features in $\underline{\mathcal{F}}_{\rightarrow w}$. The resulting radar measurement is $\{[r_1 \ \theta_1 \ \dot{r}_1], \dots, [r_N \ \theta_N \ \dot{r}_N]\}$, where r_i , θ_i , and \dot{r}_i are, respectively, the range, azimuth, and range-rate (i.e., the Doppler velocity) of feature i in $\underline{\mathcal{F}}_{\rightarrow r}$. If we assume that each feature lies on the horizontal plane of the radar, then the range-rate of a feature is

$$\dot{r}_i = -[\sin(\theta_i) \ \cos(\theta_i)] \mathbf{h}_r^j, \quad (2)$$

where \mathbf{h}_r^j is the 2D radar ego-velocity at time j . A depiction of the relationship between the range-rate of stationary features and ego-velocity of the radar is shown in Figure 1.

Radar ego-velocity estimation can be cast as a linear least squares problem, where the measurement model is

$$\underbrace{\begin{bmatrix} -\dot{r}_1 \\ -\dot{r}_2 \\ \vdots \\ -\dot{r}_N \end{bmatrix}}_{\mathbf{y} = \mathbf{A}\mathbf{h}_r^j} = \begin{bmatrix} \sin(\theta_1) & \cos(\theta_1) \\ \sin(\theta_2) & \cos(\theta_2) \\ \vdots & \vdots \\ \sin(\theta_N) & \cos(\theta_N) \end{bmatrix} \mathbf{h}_r^j. \quad (3)$$

The resulting error equation is

$$\boldsymbol{\epsilon} = \mathbf{y} - \mathbf{A}\mathbf{h}_r^j, \quad (4)$$

and the radar ego-velocity estimation problem is

$$\min_{\mathbf{h}_r^j \in \mathbb{R}^2} \boldsymbol{\epsilon}^T \boldsymbol{\epsilon}. \quad (5)$$

As a result, the estimated ego-velocity at time j is

$$\mathbf{h}_r^{j*} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{y}, \quad (6)$$

with covariance

$$\boldsymbol{\Sigma}_r^j = \frac{(\boldsymbol{\epsilon}^T \boldsymbol{\epsilon})(\mathbf{A}^T \mathbf{A})}{N-2}. \quad (7)$$

We leverage RANSAC to remove outliers such as targets moving relative to $\underline{\mathcal{F}}_{\rightarrow w}$ and multipath radar reflections [18].

C. Radar Ego-Velocity Measurement Models

Let $\underline{\mathcal{F}}_{\rightarrow a}$ and $\underline{\mathcal{F}}_{\rightarrow b}$ be the reference frames of two rigidly attached radars that share and move along one horizontal sensing plane. The ego-velocity measurement models for radars a and b , at time j , are

$$\mathbf{h}_{r,a}^j = \mathbf{v}_{r,a}^j + \mathbf{n}_{r,a}^j, \quad (8)$$

$$\mathbf{n}_{r,a}^j \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{r,a}^j),$$

$$\mathbf{h}_{r,b}^j = \mathbf{R}(\theta_{ba})(\omega^j \wedge \mathbf{t}_a^{ba} + \mathbf{v}_{r,a}^j) + \mathbf{n}_{r,b}^j, \quad (9)$$

$$\mathbf{n}_{r,b}^j \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{r,b}^j),$$

where $\mathbf{v}_{r,a}$ is the ego-velocity of the radar, θ_{ba} is the rotation from radar a to radar b , ω^j is the angular velocity of the rigid body, and \mathbf{t}_a^{ba} is the translation from radar a to radar b , expressed in the reference frame of a . The vectors $\mathbf{n}_{r,a}^j$ and $\mathbf{n}_{r,b}^j$ are additive zero-mean Gaussian noise terms with covariances $\boldsymbol{\Sigma}_{r,a}^j$ and $\boldsymbol{\Sigma}_{r,b}^j$, respectively. The values of $\boldsymbol{\Sigma}_{r,a}^j$ and $\boldsymbol{\Sigma}_{r,b}^j$ are determined with use of Equation (7).

The error equations corresponding to the ego-velocity estimates are

$$\mathbf{e}_{r,a}^j = \mathbf{h}_{r,a}^j - \mathbf{v}_{r,a}^j,$$

$$\mathbf{e}_{r,a}^j \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{r,a}^j),$$

$$\mathbf{e}_{r,b}^j = \mathbf{h}_{r,b}^j - \mathbf{R}(\theta_{ba})(\omega^j \wedge \mathbf{t}_a^{ba} + \mathbf{v}_{r,a}^j), \quad (10)$$

$$\mathbf{e}_{r,b}^j \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{r,b}^j),$$

where $\mathbf{h}_{r,a}^j$ and $\mathbf{h}_{r,b}^j$ are the values from Equation (6) for radars a and b , respectively.

D. Batch 2D Radar to Radar Extrinsic Calibration

Given M pairs of synchronized radar measurements, the vector of parameters that we wish to estimate includes the ego-velocity of radar a from time 1 to M , the angular velocity of radar a from time 1 to M , the translation from radar a to b expressed in $\underline{\mathcal{F}}_{\rightarrow a}$, and the rotation from radar $\underline{\mathcal{F}}_{\rightarrow a}$ to $\underline{\mathcal{F}}_{\rightarrow b}$,

$$\mathbf{x}^T = \left[\mathbf{v}_{r,a}^1{}^T \ \omega^1 \ \dots \ \mathbf{v}_{r,a}^M{}^T \ \omega^M \ \mathbf{t}_a^{baT} \ \theta_{ba} \right]. \quad (11)$$

Our calibration problem is to solve

$$\min_{\mathbf{x}} \sum_{j=1}^M \mathbf{e}_{r,a}^j{}^T \boldsymbol{\Sigma}_{r,a}^{j-1} \mathbf{e}_{r,a}^j + \mathbf{e}_{r,b}^j{}^T \boldsymbol{\Sigma}_{r,b}^{j-1} \mathbf{e}_{r,b}^j. \quad (12)$$

E. Scale and Angular Velocity Indistinguishability

Unfortunately, the optimization problem defined by Equation (12) has infinitely many indistinguishable solutions.

Given any solution that minimizes Equation (12), another minimizer can be found by arbitrarily scaling $\omega^j \forall j = 1, \dots, M$ and \mathbf{t}_a^{ba} by $\gamma \in \mathbb{R}$ and $\frac{1}{\gamma}$, respectively. However, the problem can be made distinguishable with additional constraints.

To make the optimization problem identifiable (see Section IV), we constrain

$$\|\mathbf{t}_a^{ba}\|_2 = 1. \quad (13)$$

We enforce this constraint by setting

$$\mathbf{t}_a^{ba} = \begin{bmatrix} \cos(\theta_t) \\ \sin(\theta_t) \end{bmatrix}, \quad (14)$$

where θ_t is the angle from the x-axis of radar a to the line of possible translations between radars a and b . Since the angle to the line is periodic with period π , we bound $0 \leq \theta_t < \pi$. We denote the resulting unscaled angular velocity as ω_γ^j . Consequently, our vector of parameters for the optimization problem becomes

$$\mathbf{x}^T = \left[\mathbf{v}_{r,a}^{1T} \quad \omega_\gamma^1 \quad \dots \quad \mathbf{v}_{r,a}^{MT} \quad \omega_\gamma^M \quad \theta_t \quad \theta_{ba} \right]. \quad (15)$$

We substitute Equation (14) into Equation (12) and solve the problem using the Levenberg-Marquardt algorithm.

F. Problem Initialization

Since the two radar sensors provide no rotational information, we require a method to initialize the angular velocities that appear in Equation (15). To start, we determine θ_{ba} by finding the K pairs of ego-velocity estimates that have similar magnitudes. Using these ‘velocity pairs,’ we compute

$$\theta_{ba}^l = [0 \quad 0 \quad 1] \left(\begin{bmatrix} \mathbf{h}_{r,a}^l \\ 0 \end{bmatrix} \times \begin{bmatrix} \mathbf{h}_{r,b}^l \\ 0 \end{bmatrix} \right). \quad (16)$$

The initial θ_{ba} is the median of $\theta_{ba}^l \forall l = 1, \dots, K$. To initialize θ_t , we use

$$\begin{aligned} \mathbf{b}^j &= [b_x^j \quad b_y^j]^T = \mathbf{R}(\theta_{ba})^T \mathbf{h}_{r,b}^j - \mathbf{h}_{r,a}^j, \\ \theta_t^j &= \arctan2 \left(\frac{b_y^j}{\|\mathbf{b}^j\|_2}, -\frac{b_x^j}{\|\mathbf{b}^j\|_2} \right). \end{aligned} \quad (17)$$

Each θ_t^j is mapped to the corresponding value within $[0, \pi)$ and the initial θ_t is the median of $\theta_t^j \forall j = 1, \dots, M$. By fixing θ_t and θ_{ba} to our initial estimates, Equation (12) becomes an unconstrained quadratic problem. We solve this problem to initialize ω_γ^j and $\mathbf{v}_{r,a}^j \forall j = 1 \dots M$.

IV. IDENTIFIABILITY

In this section, we prove that, given sufficient excitation of the system, the extrinsic calibration problem is identifiable. Since a problem that is locally weakly observable is also identifiable (in the batch setting), we use the rank criterion defined by Hermann and Krener in [19] in our proof. Section IV-A reviews the rank criterion. In Section IV-B, we demonstrate that our problem is locally weakly observable. Finally, we highlight important degenerate motions in Section IV-C that result in a loss of observability and potentially also identifiability.

A. The Observability Rank Criterion

Consider the system

$$S \begin{cases} \dot{\mathbf{x}} = \mathbf{f}_0(\mathbf{x}) + \sum_{j=1}^p \mathbf{f}_j(\mathbf{x}) u_j \\ \mathbf{y} = \mathbf{h}(\mathbf{x}) \end{cases}, \quad (18)$$

where \mathbf{x} is the state vector, $\mathbf{f}_0(\mathbf{x})$ is the drift vector field, $\mathbf{f}_j(\mathbf{x})$ is a vector field on the state manifold that is linear with respect to the control input u_j , \mathbf{y} is the measurement vector, and $\mathbf{h}(\mathbf{x})$ is the measurement model. Given the vector field $\mathbf{f}(\mathbf{x})$, we can compute the Lie derivative of \mathbf{h} with respect to \mathbf{f} , which is defined as

$$L_{\mathbf{f}} \mathbf{h}(\mathbf{x}) = \nabla_{\mathbf{f}} \mathbf{h}(\mathbf{x}) = \frac{\partial \mathbf{h}(\mathbf{x})}{\partial \mathbf{x}} \mathbf{f}(\mathbf{x}). \quad (19)$$

The n^{th} Lie derivative of h with respect to \mathbf{x} along vector field \mathbf{f} is defined as,

$$L_{\mathbf{f}}^n \mathbf{h}(\mathbf{x}) = \frac{\partial L_{\mathbf{f}}^{n-1} \mathbf{h}(\mathbf{x})}{\partial \mathbf{x}} \mathbf{f}(\mathbf{x}), \quad (20)$$

where $L^0 \mathbf{h}(\mathbf{x}) = \mathbf{h}(\mathbf{x})$.

The Lie derivatives can be vertically stacked to form the observability matrix \mathbf{O} . From Hermann and Krener [19], a system is locally weakly observable at \mathbf{x} if the matrix \mathbf{O} is full column rank at \mathbf{x} .

B. Observability of Extrinsic Calibration

Let $\mathbf{h}_{r,a}$ and $\mathbf{h}_{r,b}$ be 2D radar ego-velocity measurements. We define the state at timestep t_j as

$$\mathbf{x}^T(t_j) = [\omega_\gamma(t_j) \quad \alpha_\gamma(t_j) \quad \theta_t \quad \theta_{ba}], \quad (21)$$

where $\alpha_\gamma(t_j)$ is the instantaneous angular acceleration.¹ We assume the vehicle follows the constant angular acceleration model given by

$$\begin{aligned} \dot{\alpha}_\gamma(t_j) &= n_\alpha, \\ n_\alpha &\sim \mathcal{N}(0, \sigma_\alpha^2), \end{aligned} \quad (22)$$

where n_α is an additive zero-mean Gaussian noise term with variance σ_α^2 . Since the motion is noiseless in this analysis (i.e., $\dot{\alpha}_\gamma(t_j) = 0$), the motion model is

$$\dot{\mathbf{x}}^T(t_j) = [\alpha_\gamma(t_j) \quad \mathbf{0}_{3 \times 1}^T]. \quad (23)$$

We can substitute Equation (8) into Equation (9), which simplifies the measurement model to

$$\mathbf{h}_{r,b}^j = \mathbf{R}(\theta_{ba}) \left(\omega_\gamma(t_j) \begin{bmatrix} \cos \theta_t \\ \sin \theta_t \end{bmatrix} + \mathbf{h}_{r,a}^j \right). \quad (24)$$

The observability matrix of this system can be written as

$$\mathbf{O} = \begin{bmatrix} \nabla_{\mathbf{x}} L^0 \mathbf{h}_{r,b}^j \\ \nabla_{\mathbf{x}} L^1 \mathbf{h}_{r,b}^j \end{bmatrix}, \quad (25)$$

which is full column rank² except when the sensor platform motion is degenerate, as discussed below.

¹The analysis can, in fact, be simplified by removing the angular acceleration state; we use this formulation, specifically, in Section IV-C.

²The rank of \mathbf{O} can be determined using a symbolic math package. We omit the full proof for brevity.

C. Degeneracy Analysis

The system is unobservable (and potentially unidentifiable) when \mathbf{O} does not have full column rank. The determinant of the observability matrix is

$$\det(\mathbf{O}) = \begin{bmatrix} 0 & 0 & \alpha_\gamma(t_j) \end{bmatrix} \left(\begin{bmatrix} \mathbf{h}_{r,a}^j \\ 0 \end{bmatrix} \times \begin{bmatrix} \cos \theta_t \\ \sin \theta_t \\ 0 \end{bmatrix} \right), \quad (26)$$

which is rank-deficient when $\det(\mathbf{O}) = 0$. As a result, the system must have nonzero angular acceleration, α_γ , and nonzero ego-velocity, \mathbf{h}_a . Additionally, the direction of ego-motion must not align with the sensor translation axis.

V. EXPERIMENTS

To verify the performance of our algorithm, we conducted a series of simulated and real-world experiments. In Section V-A we show, using simulated data, that our algorithm is robust to realistic levels of radar measurement noise and that it yields an improved ego-velocity estimate. In Section V-B, we compare our approach to two state-of-the-art methods on the publicly-available Endeavour dataset.³

A. Simulation Studies

We performed a series of simulation studies to evaluate the robustness of our algorithm to measurement noise. We varied the simulation duration and the level of noise and generated 100 randomized trials with each pair of settings. Each simulation ranged in duration from 15 s to 120 s; the simulated sensor platform followed a periodic, nominal (noise-free) trajectory with sufficient excitation for our calibration problem (see Figure 2). The radar ego-velocity estimates for radars a and b were computed using the ground truth linear and angular velocities of the platform along the trajectory. Ego-velocity measurements from radars a and b were then corrupted with zero-mean Gaussian noise ($\Sigma_{r,a}^j = \Sigma_{r,b}^j = \sigma_r^2 \mathbf{I}_2$), where the standard deviation of the noise (σ_r) ranged from 0.05 m/s to 0.2 m/s. Based on our experiments (discussed in more detail in Section V-B), we found the real-world measurement noise to be at the lower end of this range.

The error distribution of the estimated calibration parameters is shown in Figure 3. For most noise levels and durations, our estimated translation direction and rotation angle are, respectively, within 2° and 3° of the ground truth. Figure 4 shows that the median of the estimated ego-velocity errors for radars a and b are both 4 cm/s lower than the raw estimates. Importantly, this improvement can be achieved without the need for additional rotation information.

B. Real-World Experiments

We demonstrate the reliability of our method and compare to two state-of-the-art algorithms on the Endeavour dataset. Post-hoc extrinsic calibration for this dataset is challenging because the environments contain no trihedral reflectors. We demonstrate that the lack of trihedral reflectors has a

³The dataset is available at: https://gloryhry.github.io/2021/06/25/Endeavour_Radar_Dataset.html




Fig. 2: Top: ego-velocity of radar a over 15 s. Bottom: angular velocity of radar a over 15 s. Both plots show the full period of the velocity functions.

negligible impact on our method, but is detrimental to the method in Olutomilayo et al. [2]. Additionally, we show that the parameters estimated by our method result in smaller velocity errors than the parameters estimated by two state-of-the-art methods. The first technique follows the approach in Olutomilayo et al. [2]. To build the required map for this method, we collate measurements from one radar while the vehicle is stationary. The second method is similar to the approach in Burnett et al. [3]. The parameters estimated by this method are included in the Endeavour dataset.⁴ Next, we demonstrate that normal driving motions provide sufficient excitation to calibrate radar pairs that have translation axes which align with the forward direction of the vehicle. Finally, we show that, when a source of angular velocity information is available, the scale of the translation between the radar pair can be estimated.

⁴The radar to lidar extrinsic calibration code for the Endeavour dataset can be found at: https://github.com/gloryhry/radar_lidar_static_calibration

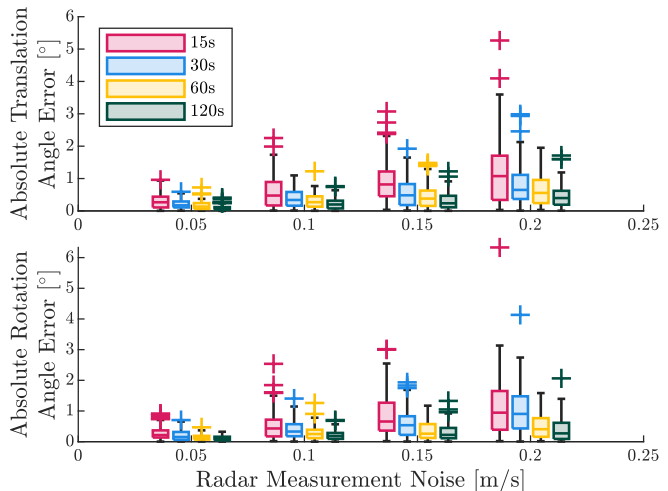


Fig. 3: Absolute translation direction and rotation angle error for our algorithm at varying levels of measurement noise and simulation durations. The translation direction is the angle from the x-axis of radar a to the line of indistinguishable translations between radars a and b .

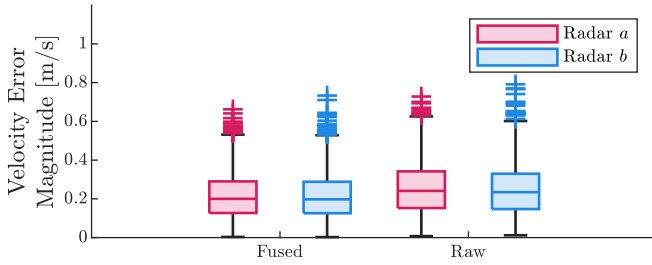


Fig. 4: Raw and fused ego-velocity estimate errors of radars *a* and *b* for an experiment that is 120 s in duration with a measurement noise level of 0.2 m/s.

The Endeavour dataset was collected from a small shuttle bus driving around three different loops in a campus setting. The dataset contains two runs for each loop, where each run is roughly 10 minutes long. The shuttle bus has a BDStar Navigation Npos320 RTK GNSS, four Velodyne VLP-16 lidars, and five Continental ARS430 radars, which operate at 100 Hz, 10 Hz, and 14 Hz, respectively. The radar labelled *Near_5* is mounted on the front bumper and observes the environment in front of the vehicle. Radar pairs *Near_3*–*Near_1* and *Near_4*–*Near_2* observe the environment surrounding the sides of the vehicle. Radars *Near_3* and *Near_1* are mounted on the front and back driver’s side of the vehicle, while radars *Near_4* and *Near_2* are mounted on the front and back passenger side of the vehicle.⁵

Before applying our method to the Endeavour dataset, we tuned our RANSAC-based ego-velocity estimator, synchronized the radar measurement timestamps and removed zero-velocity measurement pairs. For RANSAC, the inlier and outlier thresholds were set to 40% of the number of measured reflections and 0.025 m/s, respectively. These thresholds were determined using the radar and GNSS velocity data from *East_2*. To temporally synchronize the data streams, we aligned the radar measurement timestamps using linear interpolation. Finally, we removed ego-velocity measurement pairs with magnitudes less than 0.05 m/s to improve the signal-to-noise ratio in the calibration problem.

Estimating the transforms for Olutomilayo et al. [2] required two pre-processing steps. First, we identified stationary radar measurements using the RTK GNSS data and removed points that were observed less than five times. Next, we expressed (using the Endeavour parameters) the radar point clouds in a common reference frame and associated points that were within a 10 cm threshold. Finally, the extrinsic transforms from Olutomilayo et al. [2] and Endeavour were chained together to compute the rotation angles and translation axes relative to *Near_5*.

Table I shows that our parameters are within 3° of the provided parameters while the parameters from Olutomilayo et al. [2] deviate very significantly. This deviation is due to the narrow overlap between the fields of view of some of the radar pairs, which results in sparse overlapping point clouds. Often, this systematic issue results in the data collection

⁵Additional detailed information is available at: https://gloryhry.github.io/2021/06/25/Endeavour_Radar_Dataset.html

TABLE I: θ_t and θ_{ba} parameters estimated by each method on *East_2*. Radar *a* is *Near_5* for every parameter given.

Method	Near_1		Near_2		Near_3		Near_4	
	θ_t	θ_{ba}	θ_t	θ_{ba}	θ_t	θ_{ba}	θ_t	θ_{ba}
Endeavour	1.71	-1.57	1.44	1.59	2.73	-1.58	0.33	1.56
Olutomilayo [2]	1.36	-1.60	1.41	1.58	2.16	-1.61	0.39	1.56
Ours	1.74	-1.58	1.41	1.61	2.79	-1.58	0.35	1.56

*All angles are in radians.

runs having insufficient information for the method in [2] to operate properly (see Table II).

We use the mean velocity error magnitude to evaluate the accuracy of the calibration parameters given in Table I. The velocity error of a radar measurement pair, $\mathbf{h}_{r,a}^j$ and $\mathbf{h}_{r,b}^j$, is $\mathbf{e}_{r,b}^j$ from Equation (10), where $\mathbf{v}_a^j = \mathbf{h}_{r,a}^j$, θ_{ba} is the estimated rotation, θ_t is the estimated translation axis, and ω_γ^j is the value that minimizes the magnitude of $\mathbf{e}_{r,b}^j$. Table III shows the mean velocity error magnitude for each run, radar pair, and set of parameters. Our parameters yield lower velocity errors in almost all cases, reducing the mean error magnitude for the *Near_5*–*Near_4* radar pair by over 1 cm/s.

Due to the configuration of radar pairs *Near_3*–*Near_1* and *Near_4*–*Near_2*, the ego-motion of the vehicle driving forward aligns with the translation axes of these pairs, which, in theory, should make the calibration data poorly conditioned. However, these radar pairs are mounted on the periphery of the vehicle, so any angular velocity induces unaligned ego-motion measurements. Carrying out calibration on data from pairs *Near_3*–*Near_1* and *Near_4*–*Near_2*, with initial calibration parameters greater than 20° from the Endeavour values, results in estimated parameters that are consistently within 3° of the Endeavour values; this indicates that the problem is not poorly conditioned.

By including a third sensor that is able to measure angular velocity, we can estimate the (metric) scale of the translation between the radars, without requiring the exact extrinsic transform of the third sensor to be known. The magnitude of the angular velocity of a rigid body is the same for all points on the body, allowing us to match the unscaled radar estimate to the angular velocity of the third sensor. For example, assuming that the z-axis of an on-board GNSS receiver is roughly perpendicular to the sensing plane of the radar units, we can apply constant-acceleration smoothing and linear interpolation of the GNSS pose measurements to estimate angular velocity. We tried this approach on the

TABLE II: Identifiability of Olutomilayo et al. [2] for each run and radar pair in the Endeavour dataset. Our algorithm is identifiable for each run and radar pair.

Radar Pairs	Data Collection Run					
	East1	East2	Mid1	Mid2	West1	West2
<i>Near_1</i> – <i>Near_3</i>	✓	✓	✓	✓	✓	✓
<i>Near_2</i> – <i>Near_4</i>	✓	✓	✓	✓	✓	—
<i>Near_3</i> – <i>Near_5</i>	—	✓ [†]	—	—	—	—
<i>Near_4</i> – <i>Near_5</i>	—	✓	—	—	—	—

[†] This problem is only identifiable using features that appear in less than 5% of measurements.

TABLE III: Endeavour mean velocity error magnitude for each run excluding *East2*, which was used to estimate the calibration parameters. The errors presented below show how consistent the estimated parameters are when explaining the velocity vector field of a system with a unit moment arm.

Data	Near_5–Near_1			Near_5–Near_2			Near_5–Near_3			Near_5–Near_4		
	Endeavour	Olutomilayo [2]	Ours	Endeavour	Olutomilayo [2]	Ours	Endeavour	Olutomilayo [2]	Ours	Endeavour	Olutomilayo [2]	Ours
Mid1	0.0218	0.0591	0.0175	0.0232	0.0242	0.0173	0.0185	0.0928	0.0173	0.0411	0.0334	0.0184
Mid2	0.0215	0.0580	0.0170	0.0228	0.0235	0.0166	0.0195	0.0885	0.0288	0.0289	0.0240	0.0140
East1	0.0206	0.0657	0.0152	0.0208	0.0212	0.0157	0.0152	0.0846	0.0139	0.0305	0.0249	0.0121
West1	0.0206	0.0574	0.0175	0.0247	0.0249	0.0179	0.0151	0.1323	0.0146	0.0354	0.0285	0.0149
West2	0.0210	0.0558	0.0176	0.0259	0.0266	0.0192	0.0189	0.1320	0.0178	0.0515	0.0411	0.0224

*All values are in m/s.

Endeavour dataset. After removing measurement pairs with angular velocity magnitudes less than 0.1 rad/s, we computed the translation estimates. Table IV shows that, in most cases, the metric translation values recovered by our algorithm are closer to the ground-truth Endeavour dataset values than those estimated by Olutomilayo et al. [2]. While the sign of the translation may still be positive or negative (i.e., one z-axis could be inverted), this information can be easily determined from a rough model of the system.

VI. CONCLUSION

In this paper, we presented a 2D radar-to-radar extrinsic calibration algorithm that uses radar ego-velocity data only. We proved that the yaw angle and the axis of translation between the sensors can be identified given sufficient excitation. Using simulations, we demonstrated that our calibration method is robust to varying levels of radar measurement noise and that we are able to improve the raw radar ego-velocity estimates. Finally, we showed, using data from a vehicle, that our algorithm was more reliable and accurate than a state-of-the-art method.

There are multiple potential directions for future research. Our approach could be extended to pairs of 3D radar sensors, similar to those discussed in Wise et al. [4]. Another possibility is to perform temporal calibration using ego-velocity estimates, which could simplify the estimation problem for some systems.

REFERENCES

- [1] D. Kellner, M. Barjenbruch, J. Klappstein, J. Dickmann, and K. Dietmayer, “Instantaneous ego-motion estimation using doppler radar,” in *16th Int. IEEE Conf. on Intelligent Transportation Systems (ITSC)*, 2013, pp. 869–874.
- [2] K. T. Olutomilayo, M. Bahramgiri, S. Nooshabadi, and D. R. Fuhrmann, “Extrinsic calibration of radar mount position and orientation with multiple target configurations,” *IEEE Trans. Instrumentation and Measurement*, vol. 70, pp. 1–13, 2021.
- [3] K. Burnett et al., “Boreas: A multi-season autonomous driving dataset,” *Intl. J. Robotics Research (IJRR)*, vol. 42, no. 1-2, pp. 33–42, 2023.
- [4] E. Wise, J. Peršić, C. Grebe, I. Petrović, and J. Kelly, “A continuous-time approach for 3D radar-to-camera extrinsic calibration,” in *IEEE Intl. Conf. Robotics and Automation (ICRA)*, 2021, pp. 13 164–13 170.
- [5] C. Doer and G. F. Trommer, “Radar inertial odometry with online calibration,” in *European Navigation Conf. (ENC)*, 2020, pp. 1–10.
- [6] S. Sugimoto, H. Tateda, H. Takahashi, and M. Okutomi, “Obstacle detection using millimeter-wave radar and its visualization on image sequence,” in *Int. Conf. Pattern Recog. (ICPR)*, 2004, pp. 342–345.
- [7] T. Wang, N. Zheng, J. Xin, and Z. Ma, “Integrating millimeter wave radar with a monocular vision sensor for on-road obstacle detection applications,” *Sensors*, vol. 11, no. 9, pp. 8992–9008, 2011.
- [8] D. Y. Kim and M. Jeon, “Data fusion of radar and image measurements for multi-object tracking via Kalman filtering,” *Information Sciences*, vol. 278, pp. 641–652, 2014.
- [9] J. Kim, D. S. Han, and B. Senouci, “Radar and vision sensor fusion for object detection in autonomous vehicle surroundings,” in *4th Int. Conf. Ubiquitous and Future Networks (ICUFN)*, 2018, pp. 76–78.
- [10] G. El Natour, O. Ait Aider, R. Rouveure, F. Berry, and P. Faure, “Radar and vision sensors calibration for outdoor 3D reconstruction,” in *IEEE Intl. Conf. Robotics and Automation (ICRA)*, 2015, pp. 2084–2089.
- [11] J. Domhof, J. F. P. Kooij, and D. M. Gavrilu, “An extrinsic calibration tool for radar, camera and lidar,” in *IEEE Intl. Conf. Robotics and Automation (ICRA)*, 2019, pp. 8107–8113.
- [12] J. Peršić, I. Marković, and I. Petrović, “Extrinsic 6DoF calibration of a radar–lidar–camera system enhanced by radar cross section estimates evaluation,” *Robotics and Autonomous Systems*, vol. 114, pp. 217–230, 2019.
- [13] C. Scholler et al., “Targetless rotational auto-calibration of radar and camera for intelligent transportation systems,” in *2019 IEEE Intelligent Transportation Systems Conf. (ITSC)*, 2019, pp. 3934–3941.
- [14] J. Peršić, L. Petrović, I. Marković, and I. Petrović, “Online multi-sensor calibration based on moving object tracking,” *Advanced Robotics*, vol. 35, no. 3-4, pp. 130–140, 2021.
- [15] L. Heng, “Automatic targetless extrinsic calibration of multiple 3D lidars and radars,” in *IEEE/RSJ Intl. Conf. Intelligent Robots and Systems (IROS)*, 2020, pp. 10 669–10 675.
- [16] C. C. Stahoviak, “An instantaneous 3D ego-velocity measurement algorithm for frequency modulated continuous wave (FMCW) Doppler radar data,” Master’s thesis, University of Colorado at Boulder, 2019.
- [17] D. Kellner, M. Barjenbruch, K. Dietmayer, J. Klappstein, and J. Dickmann, “Joint radar alignment and odometry calibration,” in *18th Int. Conf. Information Fusion (FUSION)*, 2015, pp. 366–374.
- [18] M. A. Richards, J. A. Scheer, and W. A. Holm, Eds., *Principles of Modern Radar: Basic principles*. Inst. of Eng. and Technol., 2010, vol. 1.
- [19] R. Hermann and A. Krener, “Nonlinear controllability and observability,” *IEEE Trans. Automatic Control*, vol. 22, no. 5, pp. 728–740, 1977.

TABLE IV: Estimated translation magnitude for each method on *East2*. Radar *a* is *Near_5* for every parameter given. The bolded values are the translation magnitudes closest to the Endeavour parameters.

Method	Near_1	Near_2	Near_3	Near_4
Endeavour	5.68	5.82	0.83	0.86
Olutomilayo [2]	4.87	6.04	0.47	0.96
Ours	6.08	5.54	0.77	0.95

*All values are in m.