Extrinsic 6DoF Calibration of a Radar – LiDAR – Camera System Enhanced by Radar Cross Section Estimates Evaluation

Juraj Peršić[a], Ivan Marković[a], Ivan Petrović[a],

[a]University of Zagreb Faculty of Electrical Engineering and Computing,
Department of Control and Computer Engineering,
Laboratory for Autonomous Systems and Mobile Robotics,
Unska 3, HR-10000, Zagreb, Croatia,
`juraj.persic@fer.hr, ivan.markovic@fer.hr, ivan.petrovic@fer.hr`

Abstract

Autonomous navigation of mobile robots is often based on information from a variety of heterogeneous sensors; hence, extrinsic sensor calibration is a fundamental step in the fusion of such information. In this paper, we address the problem of extrinsic calibration of a radar – LiDAR – camera sensor system. This problem is primarily challenging due to sparse informativeness of radar measurements. Namely, radars cannot extract rich structural information about the environment, while their lack of elevation resolution, that is nevertheless accompanied by substantial elevation field of view, introduces uncertainty in the origin of the measurements. We propose a novel calibration method which involves a special target design and two-step optimization procedure to solve the aforementioned challenges. First step of the optimization is minimization of a reprojection error based on an introduced point-circle geometric constraint. Since the first step is not able to provide reliable estimates of all the six extrinsic parameters, we introduce a second step to refine the subset of parameters with high uncertainty. We exploit a pattern discovered in the radar cross section estimation that is correlated to the missing elevation angle. Additionally, we carry out identifiability analysis based on the Fisher Information Matrix to show minimal requirements on the dataset and to verify the method through simulations. We test the calibration method on a variety of sensor configurations and address the problem of radar vertical misalignment. In the end, we show via extensive experiment analysis that the proposed method is able to reliably estimate all the six parameters of the extrinsic calibration.

Keywords:
Sensor calibration, Radar, LiDAR, Camera, Radar cross section

1. Introduction

Autonomous systems navigate through the environment based on the information they gather from sensors. They have to solve many task such as simultaneous localization and mapping, detection and tracking of moving objects, etc., based on the available information from a variety of sensors. Commonly used proprioceptive sensors in robotics can include global positioning system, inertial measurement units, and wheel encoders, while extroceptive sensors include LiDARs, cameras, sonars, and radars. Appropriateness of a sensor is dependent on the application, because these sensors utilize different physical phenomena, leading to different sets of advantages and disadvantages. Therefore, to achieve a robust, full-stack autonomy, information from the aforementioned sensors is often fused.

The fundamental step in sensor fusion is sensor calibration, commonly divided to intrinsic and extrinsic calibration. The former provides internal parameters of an individual sensor related to its working principle, while the latter represent spatial displacement between a pair of sensors. The calibration can tackle both parameter groups at the same time or assume that sensors are already intrinsically calibrated and proceed with the extrinsic calibration. On the one hand, methods for finding intrinsic parameters do not share much similarities for different types of sensors since they are related to the working principle of the sensor. On the other hand, parametrization of extrinsic calibration, i.e., homogeneous transform, can always be expressed in the same manner, regardless of the sensors involved in it. Nevertheless, solving the extrinsic calibration requires finding correspondences in the data acquired by the sensors which can be challenging since different types of sensors measure different physical quantities. The calibration approaches can be target-based or targetless. In the case of target-based calibration, correspondences originate from a specially designed target, while targetless methods utilize environment features perceived by both sensors. Registration of structural correspondences can be avoided by motion-based methods, which leverage motion estimated by individual sensors for calibration.

Cameras and LiDARs are rich sources of information, commonly used in robotics, which often require precise calibration. Therefore, extensive research has been devoted to calibration of these sensors within all aforementioned calibration approaches. Target-based camera calibration approaches, based on pioneering work [1, 2], typically involve planar targets with known pat-

terns such as checkerboard [3] or a grids of circles [4]. Novel calibration target is presented in [5] where authors use a noise-like pattern with many features of varying scales. It is suitable for both intrinsic and extrinsic calibration of multiple cameras with no or little filed of view (FoV) overlap. LiDAR calibration also uses flat surfaces as calibration targets. For instance, intrinsic calibration of LiDARs is achieved by placing the LiDAR inside a box [6] or by observing planar wall [7], while extrinsic calibration of multiple 2D LiDARs was found by the aid of a corner structure [8]. Extrinsic target-based calibration between LiDARs and cameras has also received significant research attention, while the common targets are planes covered with a pattern suited for camera detection. Widely adopted and extended method presented in [9] introduced point-plane geometric constraint initially designed for 2D LiDAR – camera calibration. Proposed approach was also applied in the calibration of a 3D LiDAR and a camera [10]. Further improvements were made by decoupling rotation from translation in the optimization procedure [11]. To reduce the labour requirements, authors in [12] extended the method with global correspondence registration which allows for multiple plane observations in a single shot. The same constraint was used in [13] where instead of checkerboard pattern, AprilTag fiducial markers were used [14]. Additionally, they extended the extrinsic calibration with estimation of intrinsic LiDAR parameters. AprilTag markers and the same geometric constraint were also used in [15] as a part of multi-sensor graph based calibration. Besides commonly used point-plane constraint, 3D LiDAR-camera pair was calibrated based on the point-point correspondences. In [16] authors used a target with circular holes for localization, while in [17] authors extracted centerline and edge features of a V-shaped planar target to improve 2D LiDAR-camera calibration.

Radars are frequently used in automotive applications for detection and tracking of multiple objects due to their low price and robustness. Since radars cannot provide rich information about the detections, automotive systems often fuse radars with cameras [18, 19] or LiDARs [20, 21] to perform advanced tasks, e.g., object classification [22, 23]. Although sensor fusion requires precise calibration, extrinsic radar calibration has not gained much research attention. Existing calibration methods are all target-based since, for all practical means and purposes, the targetless methods are hardly feasible due to limited resolution of current automotive radar systems, as the radar is virtually unable to infer the structure of the detected objects and extract features such as lines or corners. Current radars have no elevation resolution while the information about the detected objects they provide contains range, azimuth angle, radar cross section (RCS) and range-rate based on the Doppler effect. Although having no elevation resolution, radars have substantial elevation FoV which makes the extrinsic calibration challenging due to the uncertainty in the origin of the measurements. Concerning automotive radars, common operating frequencies (24 GHz and 77 GHz) result with reliable detections of conductive objects, such as plates, cylinders and corner retroreflectors, which are then used in intrinsic and extrinsic calibration methods [24]. In [25] authors used a metal plate as the target for radar – camera calibration assuming that all radar measurements originate from a single ground plane, thereby neglecting the 3D nature of the problem. The calibration is then found by optimizing a homography transformation between the ground and image plane. Later, a similar approach was adopted by using thin metal poles as calibration targets [18]. Contrary to previous examples, 3D nature of the problem was taken into account by moving a corner retroreflector within the FoV and manually searching for detection intensity maximums [26]. Authors assumed that detections lie on the radar plane (zero elevation plane in the radar coordinate frame) and used the points to optimize a homography transform between the radar and camera. The drawback of this method is that the maximum intensity search is prone to errors, since the returned intensity depends on a number of factors, e.g., target orientation and radar antenna radiation pattern, which is usually designed to be as constant as possible in the nominal FoV.

Even though current automotive radars cannot provide 3D information about the targets (the missing elevation angle), accurate 6DoF extrinsic calibration involving a more informative sensor, e.g., LiDAR or camera, can also be especially useful for detecting *vertical misalignment*. Namely, radars should be mounted on the vehicle so that the radar and the ground plane are aligned. Vertical misalignment is loosely defined as an angular deviation between these two planes, while typical commercial radars allow the misalignment for up to a few degrees (e.g. Delphi ESR allows $\pm 1°$). With greater misalignment, radar range and detection probability are decreased, as less energy is radiated in the direction of interest. To the best of the authors' knowledge, existing related work does not address the vertical misalignment problem nor are the existing calibration methods accurate enough to provide reliable misalignment assessment. However, several misalignment detection procedures are patented [27, 28, 29], thus confirming the importance of the aforementioned issue.

Sensor calibration approaches should ideally address the aspects of identifiability, i.e., give answers if and to what extent in terms of uncertainty, one can estimate the parameters of the addressed calibration problem. Furthermore, minimal requirements on the dataset can also give practical advice on the experiment design and are also useful for robust estimation techniques (e.g. RANSAC), where the time cost of the estimation depends on the minimal size of the dataset. Some methods approach the identifiability question from the geometric viewpoint, while others from the framework of nonlinear observability or through statistical tools such as Fisher Information Matrix (FIM). In [13] authors calibrated a 3D LiDAR – camera pair by examining how the geometric point-plane constraints react in the scenarios in which they observe one, two, or three planes with linearly independent normals. Nonlinear observability analysis developed in [30] is a convenient tool for cases where system dynamics are exploited in the calibration, such as visual-inertial odometry combined with extrinsic calibration, as demonstrated in [31] and [32]. Authors in [8] presented a solution which uses corner structures to perform extrinsic calibration of multiple 2D LiDARs. To show identifiability requirements, they relied on the FIM rank to show that the problem becomes identifiable when at least three perpendicular planes

are observed. FIM was also used in motion-based calibration [33] to detect unobservable directions in parameter space from the available data.

In this paper we present a novel target-based calibration method for extrinsic 6DoF calibration of 3D LiDAR – radar and camera – radar sensor pairs. By using FIM based statistical analysis, we also address the questions of parameter identifiability, estimation uncertainty, and the choice of transform parametrization. The proposed method involves a special calibration target design whose properties enable accurate cross-sensor correspondence localization and registration. Afterwards, these correspondences are used in two consecutive optimization steps: reprojection error based optimization and RCS enhanced optimization. When combined, the steps are able to accurately estimate all the 6DoF of the extrinsic calibration. The current paper draws upon our earlier work [34], where the target design and preliminary results of 3D LiDAR – radar calibration were presented. We extend this work with novel contributions by adding camera in the optimization framework, performing FIM based identifiability and estimation uncertainty analysis, introducing improved RCS enhanced optimization step, and correspondingly reporting extended experimental analysis for both sensor pairs with two radars from different manufacturers to demonstrate the validity of the proposed method.

The paper is organized as follows. Section 2 elaborates the calibration method including calibration target design and data correspondence registration. Section 3 explains two steps of the optimization: reprojection error optimization and RCS optimization. Section 4 gives insight on the theoretical background used in the identifiability analysis and the tools used in the FIM analysis. Section 5 provides details on the results of the identifiability analysis, the setup and the results of the real-world experiments. In the end, Section 6 concludes the paper.

2. Target based correspondence registration

The proposed method is based on observing a calibration target placed at a range of different heights and positions, both within and outside of the nominal radar FoV. The final goal of the calibration is to estimate relative displacements between the radar, LiDAR and camera coordinate frames, i.e., $\mathcal{F}_r$, $\mathcal{F}_l$, and $\mathcal{F}_c$, respectively. In the present paper, we will designate both the 3D LiDAR and camera as 3D sensors, in the sense that they can both infer the 3D position of a known target from measurements. The method further assumes that the 3D sensor's FoV exceeds radar's vertical FoV, which is the case in most applications. Given that, when it is not necessary to differentiate between the two 3D sensor coordinate frames, we will designate the 3D sensor frame as $\mathcal{F}_s$. Additionally, due to challenges associated with radars, such as ghost measurements from multipath propagation and low angular resolution, data collection has to be performed outdoors at a set of distances ranging from 2–10 m with enough clear space around the target.

2.1. Calibration Target Design

Calibration target design for radar – LiDAR calibration was developed within our previous work [34], where we gave de-

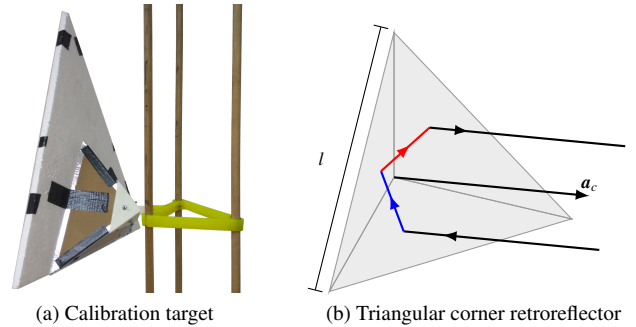(a) Calibration target (b) Triangular corner retroreflector

Figure 1: Constructed calibration target and the illustration of the working principle of the triangular trihedral corner retroreflector

tailed remarks considering the design. However, for completeness, in this section we provide essential information necessary for the rest of the paper.

Properties of a well-designed target are (i) ease of detection and (ii) high localization accuracy for all the three sensors. For the radar, a target with a high RCS provides good detection rates. Formally, RCS of an object is defined as the area of a perfectly conducting sphere whose echo strength would be equal to the object strength [24]. Consequently, it is a function of object size, material, shape, and orientation.

We proposed a complementary target design which consist of a styrofoam triangle covered by a checkerboard-like pattern and a triangular corner retroreflector. Since the styrofoam is mostly made out of air (98%), it is virtually invisible to the radar, while its flat shape enables precise localization within the point cloud. Furthermore, its triangular shape solves localization ambiguity issues existing with common rectangular targets caused by the finite LiDAR resolution, as shown in [35] and [36]. On the other hand, the triangular corner retroreflector, which consists of three orthogonal flat metal triangles, has good detection and localization properties with the radar. It has an interesting property that any ray reflected from all three sides is returned in the same direction as illustrated in Fig. 1b. Due to this property, regardless of the incident angle, many rays are returned to their source, i.e., the radar, which leads to a high and orientation-insensitive RCS. When the retroreflector axis, $\mathbf{a}_c$, points directly to the radar, it reaches its maximum RCS value:

$$\sigma_c = \frac{\pi l^4}{3\lambda^2}, \tag{1}$$

where l is the hypotenuse of the retroreflector's side and λ is radar's operating wavelength. Furthermore, authors in [37] show that all the rays which go through multiple reflections travel the same length as the ray which is reflected directly from the corner centre, thus providing good localization accuracy. Lastly, target stand is designed to have RCS as small as possible, while it allows adjusting of target's height and orientation. The constructed radar calibration target and an illustration of the working principle is shown in Fig. 1a.

2.2. Correspondence registration

Correspondence registration procedure from our previous work [34] is expanded with checkerboard detection in the im-
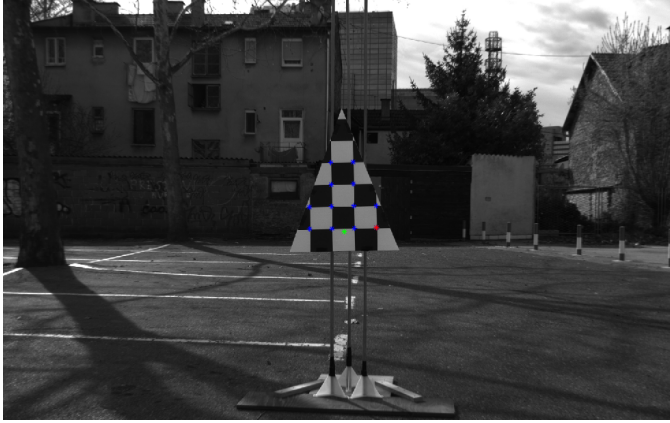
Figure 2: Front view of the calibration target at the experiment site with detected corners (blue and origin red) and estimated position of the retroreflector origin (green).

ages. It starts with the detection and localization of a target in the LiDAR point cloud or camera image. Once we obtain the 3D location of the retroreflector origin, the rest of the method is equal for the camera–radar and LiDAR–radar calibration. Method for the target localization within the point cloud is explained in [34], while the image procedure is given in the sequel.

The intrinsic calibration of a camera, modelled as a pinhole camera with radial distortion, is found using the Kalibr toolbox [38]. In the sequel we perform all the steps on the rectified images. To estimate the position of a corner origin in the image, we use the toolbox developed in [12], which was able to effectively find the corners in our cluttered environment shown in Fig. 2. The size of the checkerboard corners was selected to present a compromise between the number of points on the target and the ability to be detected at larger distances. In the end, we opted for the size of 0.1 m. However, since our target did not have a rectangular form, we had to adapt the toolbox to accept non-square patterns. After the corners of the pattern are found, to recover the pose of the triangle based on a known checkerboard configuration, we used the built-in Matlab function *extrinsics* , which is based on a closed form solution resulting with sufficient accuracy. Finally, as in the LiDAR's case, the position of the retroreflector origin $^c\mathbf{x}_c$ is calculated based on the pose of the checkerboard and known target configuration.

Radar reports data as a list of detected objects described by the measured azimuth $^r\phi_{r,i}$, range $^r r_{r,i}$ and RCS $\sigma_{r,i}$. The i-th object from the list is described by the vector $^r\boldsymbol{m}_i = [^r\phi_{r,i} \; ^r r_{r,i} \; ^r\sigma_{r,i}]$ in the radar coordinate frame, $\mathcal{F}_r : (^r x, ^r y, ^r z)$. The only structural property of detected objects is contained within the RCS, which is influenced by many other factors; hence, it is impossible to classify a detection as the retroreflector based solely on radar measurements. To find the matching object, a rough initial calibration is required, e.g., with a measurement tape, which is used to transform the estimated corner position from the 3D sensors coordinate frame, $\mathcal{F}_s : (^s x, ^s y, ^s z)$, to the radar frame $\mathcal{F}_r : (^r x, ^r y, ^r z)$, and eliminate all other objects that fall outside of a predefined distance threshold. The correspondence is accepted only if a single object is left.

Lastly, we form correspondence groups by observing the target at rest for a short period while the registered correspondences fill a correspondence group with pairs of vectors $^r\boldsymbol{m}_i$ and $^s\boldsymbol{x}_s$. Variances of the radar data ($^r\phi_{r,i}, ^r r_{r,i}, ^r\sigma_{r,i}$) within the group are used to determine the stability of the target. If any of the variances surpasses a preset threshold, the correspondence is discarded, since it is likely that the target detection was obstructed. Otherwise, the values are averaged. Hereafter, we will refer to the mean values of the groups as radar and 3D sensor measurements.

3. Two-step Optimization

In this section we provide insight on how the optimization is performed to obtain the 6DoF transformation between the radar and the 3D sensor. The optimization is divided in two steps which are based on different information provided by the radar. Namely, first step, i.e., reprojection error optimization, optimizes all six transformation parameters based on the comparison of 3D corner positions estimated by the 3D sensor, and range and azimuth information provided by the radar. On the other hand, second step, i.e., RCS optimization, uses information from the 3D sensor combined with the radar RCS estimate to refine only a subset of transformation parameters which could not be estimated reliably in the first step.

3.1. Reprojection error optimization

Reprojection error optimization is based on a point–circle geometric constraint, while the optimization parameter vector includes the translation and rotation parameters, i.e., $\boldsymbol{c}_r = [^r\boldsymbol{p}_s \; ^s_r\boldsymbol{\Theta}]$. For translation, we choose position of the 3D sensor in the $\mathcal{F}_r$, $^r\boldsymbol{p}_s = [^r p_{s,x} \; ^r p_{s,y} \; ^r p_{s,z}]^T$. For rotation, we choose Euler angles parametrization $^s_r\boldsymbol{\Theta} = [^s_r\theta_z \; ^s_r\theta_y \; ^s_r\theta_x]$ where rotation from $\mathcal{F}_r$ to $\mathcal{F}_s$ is given by:

$$^s_r R(^s_r\boldsymbol{\Theta}) =^s_2 R_x(^s_r\theta_x)^2_1 R_y(^s_r\theta_y)^1_r R_z(^s_r\theta_z). \qquad (2)$$

Although transformation can be expressed in multiple ways, the proposed choice is preferable due to its distribution of uncertainty caused by radar's inability to measure elevation angle. Further elaboration of the parametrization choice will be given in Sec. 4.2 with results in Sec. 5.1 which further confirm this assertion.

Figure 3 illustrates the calculation of the reprojection error for the i-th paired measurement. As discussed previously, radar provides measurements in spherical coordinates lacking elevation $^r\boldsymbol{s}_{r,i} = [^r r_{r,i} \; ^r\phi_{r,i} \; \sim]$, i.e., it provides an arc $^r a_{r,i}$ upon which the object potentially resides. On the other hand, 3D sensor provides a point in Euclidean coordinates $^s\boldsymbol{x}_{s,i}$. Using the current transformation estimate, 3D sensor measurement $^s\boldsymbol{x}_{s,i}$ is transformed into the radar coordinate frame:

$$^r\boldsymbol{x}_{s,i}(\boldsymbol{c}_r) = {}^s_r R^T(^s_r\boldsymbol{\Theta}) \, ^s\boldsymbol{x}_{s,i} + {}^r\boldsymbol{p}_s, \qquad (3)$$

and then $^r\boldsymbol{x}_{s,i}$ is converted to spherical coordinates $^r\boldsymbol{s}_{s,i} = [^r r_{s,i} \; ^r\phi_{s,i} \; ^r\psi_{s,i}]$. By neglecting the elevation angle $^r\psi_{s,i}$, we obtain the arc $^r a_{s,i}$ upon which 3D sensor measurement resides and can be compared to the radar's. Reprojection error $\epsilon_{r,i}$ is then
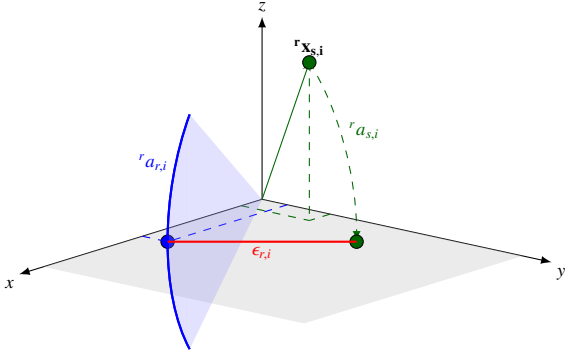
4

Figure 3: Illustration of the reprojection error calculation. Green: 3D sensor's measurement; blue: radar's; red: reprojection error.

defined as the Euclidean distance of points on the arc for which $^r\psi_{r,i} = {}^r\psi_{s,i} = 0°$:

$$\epsilon_{r,i}(\boldsymbol{c}_r) = \left\| \begin{bmatrix} {}^r r_{r,i} \cos({}^r\phi_{r,i}) \\ {}^r r_{r,i} \sin({}^r\phi_{r,i}) \end{bmatrix} - \begin{bmatrix} {}^r r_{s,i} \cos({}^r\phi_{s,i}) \\ {}^r r_{s,i} \sin({}^r\phi_{s,i}) \end{bmatrix} \right\|. \quad (4)$$

The estimate of the calibration parameters $\hat{\boldsymbol{c}}_r$ is obtained using the Levenberg-Marquardt (LM) algorithm, which minimizes the sum of squared reprojection errors from N measurements:

$$\hat{\boldsymbol{c}}_{\mathbf{r}} = \arg\min_{\boldsymbol{c}_r} \left(\sum_{i=1}^{N} \epsilon_{r,i}^2(\boldsymbol{c}_r) \right). \quad (5)$$

Reprojection error optimization yields unequally uncertain calibration parameters, in other words, some parameters are easier to estimate than the others. The lack of radar's elevation angle measurement leads to poor estimation of $^r p_{s,z}$, $^s_r\theta_y$ and $^s_r\theta_x$. A formal analysis of these properties based on FIM is carried out in Section 4.

3.2. RCS optimization

For the second optimization step, i.e., the RCS optimization, we propose a method that is based on the distribution of RCS across the 3D measurements. The idea of this step is to exploit patterns discovered in radar's RCS estimation; namely, RCS depends on the object properties and relative orientation with respect to the radar. The reason behind these patterns is that radars can only estimate the RCS based on the amplitude difference between the radiated and received electromagnetic energy. Ideally, radars would radiate with constant strength within the nominal FoV and zero outside of it; however, this is infeasible and leads to errors in RCS estimation. Using the retroreflector as a calibration target, we can assume that the RCS estimate is constant with respect to the object properties, since we use the same target in all the experiments, and with respect to the relative orientation, due to the retroreflector properties. However, radars emit the highest amount of radiation at the zero elevation angle, while the dependence between elevation angle and radiated energy, and thus RCS estimation, can be modelled as a curve. Since radars cannot distinguish objects at different elevation angles, they can neither compensate for the error in the RCS estimation. For the usual application, such as object tracking, this might not seem like an exploitable property, but

for our case of calibration with a target of a stable RCS, we can exploit this pattern of varying RCS with respect to elevation and enhance calibration results.

The results from the reprojection error optimization exhibit varying uncertainty among the calibration parameters, which was examined in the identifiability analysis (cf. Section 4). In the RCS optimization step, only the parameters with the highest uncertainty from the previous optimization step are refined. Given that, the RCS optimization parameter vector consist of a subset of transformation parameters and curve parameters

$$\boldsymbol{c}_\sigma = [{}^r p_{s,z}\ {}^s_r\theta_y\ {}^s_r\theta_x\ c_0, c_2],$$

while other extrinsic parameters are kept fixed. Through the empirical evaluation of the used radars, we have noticed that the RCS – elevation dependence follows a quadratic form; hence, we have modelled it as a second order polynomial without the linear term. In the experiments (cf. Section 5.2), the proposed model gave accurate and stable results for two automotive radars from different manufacturers. However, other radars might exhibit different patterns and the procedure could require a revision of the curve parametrization. To initialize curve parameters, a fair assumption is to assume that at the elevation angle zero, RCS is equal to the target maximum value defined in (1), while at the edge of the nominal FoV it reduces -3dBm. Due to the sufficiently good initialization of transformation parameters provided by the reprojection error optimization, the proposed curve initialization showed sufficient for converging. The proposed step can be seen as a combination of extrinsic and intrinsic radar calibration, where the estimated curve is merely a nuisance variable used to obtain an enhanced extrinsic calibration (since it is of no relevance to other radar applications). Another perspective on the idea behind the RCS optimization concept is to provide a replacement for the radar's lack of elevation measurements. The prerequisite for this method is a target with reliable and stable RCS with respect to its orientation, which in our case is ensured by the retroreflector properties.

The cost function for optimization is formed as follows. After transforming a 3D sensor measurement $^s\boldsymbol{x}_{s,i}$ to $\mathcal{F}_r$, the elevation angle $^r\psi_{s,i}$ in $\mathcal{F}_r$ is calculated. Afterwards, the expected RCS is obtained using

$$\hat{\sigma}_{s,i} = c_2\, {}^r\psi_{s,i}^2 + c_0. \quad (6)$$

Cost function is then given by the sum of squared distances between the expected and measured RCS, $\hat{\sigma}_{s,i}$ and $\sigma_{s,i}$, respectively:

$$\hat{\boldsymbol{c}}_\sigma = \arg\min_{\boldsymbol{c}_\sigma} \left(\sum_{i=1}^{N} \left(\sigma_{s,i} - \hat{\sigma}_{s,i}(\boldsymbol{c}_\sigma) \right)^2 \right). \quad (7)$$

In our previous work [34], we referred to the second optimization step as the FoV optimization. Although the presently proposed and previous approach exploit the same effect, the present one shows better results and has several advantages. First, for the FoV optimization, we have noticed that it works well with many measurements, while it becomes unstable with only few measurements. The problem with the FoV optimization is that the cost function focuses only on the measurements near the

nominal FoV border and ignores all the other measurements. Therefore, the proposed RCS optimization was designed so that it takes into account all the measurements. Second, FoV optimization requires predetermination of the nominal FoV, which can also affect calibration results. The nuisance parameter in the FoV optimization, i.e., the RCS threshold, requires more precise initialization than the nuisance parameters, i.e., curve parameters, in the RCS optimization.

4. Identifiability analysis

Extrinsic calibration methods typically involve minimization of a specific reprojection error depending on the type of the data provided by the sensors. This minimization will yield an estimate of the calibration parameters, but it would also be desirable if it could provide an assessment of the whole process – for example, by answering the following questions. What are the minimal conditions on the dataset to ensure identifiability of the parameters? How should the dataset be constructed to maximize the quality of the estimation? Does the chosen parametrization fit well with the optimization problem? In the sequel, we present theoretical background and experimental results that address the aforementioned questions for the calibration problem investigated in the present paper.

For dynamical systems the term *observability* is used within the context of a procedure assessing if system states can be estimated given a sequence of measurements. The term *identifiability* is used in conjunction with a procedure for estimating system parameters that are constant over time. However, the term observability is also often used within the context of estimating constant system parameters, due to commonly used tools in control theory and robotics. Nevertheless, in the present paper we use the term *identifiability*, since we believe that it more precisely describes the problem at hand. Given that, the objective of the identifiability analysis is to determine whether it is possible to correctly estimate parameters of a model based on the chosen criterion, e.g., the reprojection error, and available data. In some cases, it is possible to derive analytical solutions for such problems. However, when nonlinear transformations in the criterion grow in complexity, using methods such as those developed in [30] becomes impractical, if not infeasible. Since our reprojection error design, described in Sec. 3.1, involves heavy nonlinearities, we decided to adopt the statistical concept of FIM through which local identifiability of the system can shown. In the sequel, we provide the theoretical background on the FIM, followed by the description of the performed experiments that can be used to address identifiability, asses the parametrization and give general advice on the experiment setup.

4.1. Theoretical Background

Before approaching any identification problem, it is important to know if it is even possible to correctly estimate the parameters in a noise-free system. That is the intuitive purpose of the identifiability analysis. To approach it more formally, we first define our system as a nonlinear regression

$$\mathbf{Y} = \mathbf{H}(\mathbf{\Theta}, \mathbf{X}) + \epsilon, \tag{8}$$

where the response variable $\mathbf{Y} \in \mathbb{R}^{2N \times 1}$ represents radar measurements, the predictor variable $\mathbf{X} \in \mathbb{R}^{3N \times 1}$ represents LiDAR measurements, $\mathbf{\Theta} \in \mathbb{R}^d$ are parameters of the extrinsic calibration, $\epsilon \sim \mathcal{N}(0, \mathbf{Q}) \in \mathbb{R}^{2N \times 1}$ is additive zero-mean white noise, N is the number of measurements, and the nonlinear transformation $\mathbf{H}(\cdot)$ represents the reprojection function. We can notice that LiDAR measurements are modelled as noise-free. This may lead to slight imprecision in the simulation of the error; however, we are not here concerned with precise estimation of the error and covariance, but with the impact of the proposed reprojection error on the identifiability of the calibration parameters.

Identifiability can be a global or a local concept for a specific $\mathbf{\Theta_0}$ [39]. Since FIM cannot provide insights into global identifiability, we restrict our analysis to local identifiability. This is sufficient for our method, since we assume to have a rough initial estimate of the parameters, e.g., by hand measuring the displacements or from the project design. Now, we move on to more formally defining the local identifiability.

Definition 4.1. Local identifiability

The noise-free system is locally identifiable at $\mathbf{\Theta_0}$ if

$$\exists U_{\mathbf{\Theta_0}} \subset \mathbb{R}^d \text{ (open subset containing } \mathbf{\Theta_0})$$

$$\forall \mathbf{\Theta} \in U_{\mathbf{\Theta_0}}, \{\mathbf{\Theta} \neq \mathbf{\Theta_0}\} \Rightarrow \{\mathbf{H}(\mathbf{\Theta}, \mathbf{X}) \neq \mathbf{H}(\mathbf{\Theta_0}, \mathbf{X})\}.$$

In other words, for a different parameter set the nonlinear function cannot yield the same output. This is intuitively clear, since we would like to see a change in the response variable given the change in the parameter values. Another theoretical concept that we require for the present problem is the *score*.

Definition 4.2. Score funtion

The score function $\dot{\mathcal{L}}_{\mathbf{\Theta}}$ is the gradient of the log-likelihood function $\mathcal{L}(\mathbf{Y}; \mathbf{\Theta}, \mathbf{X})$ at $\mathbf{\Theta}$

$$\dot{\mathcal{L}}_{\mathbf{\Theta}} = \nabla_{\mathbf{\Theta}} \log \mathcal{L}(\mathbf{Y}; \mathbf{\Theta}, \mathbf{X}).$$

The score function can be seen as an indicator of how sensitive the likelihood functions is to the change in its parameters. Intuitively, this would mean that higher the sensitivity, the more easy it should be to estimate the parameter. An interesting notion that we will use is that FIM is defined as the covariance matrix of the score.

Informally, FIM tells how much information about the parameters is available in any direction of the parameter space from observing the sample. Since the expected value of the score is zero, FIM is a positive semi-definite matrix of size $d \times d$ whose elements can be computed as

$$[\mathcal{I}(\theta)]_{i,j} = \mathrm{E}_\theta \left[\left(\frac{\partial}{\partial \theta_i} \log \mathcal{L}(Y; \theta, X) \right) \left(\frac{\partial}{\partial \theta_j} \log \mathcal{L}(Y; \theta, X) \right) \right]. \tag{9}$$

Since we defined our problem as a nonlinear regression with additive white noise, our likelihood function is simply a well-known probability density function of a multivariate normal

distribution. For such cases, it can be shown that calculation of the FIM elements simplifies to [39]

$$[\mathcal{I}(\theta)]_{i,j} = \frac{\partial \mathbf{H}(\boldsymbol{\Theta}, \mathbf{X})}{\partial \theta_i} \mathbf{Q}^{-1} \frac{\partial \mathbf{H}(\boldsymbol{\Theta}, \mathbf{X})^{\mathrm{T}}}{\partial \theta_j}. \qquad (10)$$

As discussed in [40], such simplification is beneficial, especially for numerical accuracy, which can cause problems in complex nonlinear problems.

It is also worth mentioning some additional properties of FIM that are commonly used. First, the Cramér–Rao lower bound (CRLB), calculated as an inverse of the FIM, is used to express the lower bound on the variance of the estimated parameters. Second, if we draw independent identically distributed samples, likelihood function is simply the product of individual likelihoods, whereas log-likelihood turns into summation of the individual log-likelihoods. Due to linearity, this property also holds for FIM. Therefore, if we draw two data samples of the same random variable, maximum information expressed with FIM is doubled. Finally, in [41] it was shown that the local identifiability as defined in Def. 4.1 is equivalent to the regularity of the FIM. Therefore, if FIM is not of full rank, we conclude that the problem is not identifiable.

4.2. FIM tests

After the FIM has been evaluated at the estimated maximum likelihood estimate, we proceed with test which will give us insight into: (i) minimal requirements on the dataset which ensures identifiability of our problem, (ii) appropriateness of the parametrization, and (iii) general advice on the dataset collection. To evaluate our reprojection function, we will create synthetic datasets and test FIM behavior.

To show the minimal requirements on the dataset, we will apply the rank test of FIM. For the case of 3D point – point correspondences, at least three non-colinear points are required to estimate the 6D transformation between two coordinate frames [42]. However, the problem that we face is more complex, and our reprojection error is less informative, because we use point – circle correspondences. Therefore, it is a fair assumption to take three non-colinear, but coplanar points, as a starting dataset and expand it to find the minimal requirements. The FIM is computed for each dataset and based on its regularity, we infer on the identifiability. Furthermore, numerical inaccuracies and noise can result in an illusory full rank of FIM; therefore, it is advisable to examine the numerical rank of the matrix [43]. A convenient summary statistic is given by the matrix conditional number, i.e., the ratio of the biggest and the smallest singular value, where high values indicate degeneracy of the matrix.

In order to evaluate the choice of parametrization and to provide some practical advice on the dataset collection, we will also rely on the theory of optimal experiment design. The optimal experiment is the experiment that allows estimation of parameters without bias and with minimum variance with equal or less experiment data than any other non-optimal experiment. There exist many optimality criteria which a single experiment

can satisfy; however, we will use only the T-optimality criterion, which tries to maximize the trace of the FIM. It is convenient as it tells us that we can observe only the diagonal elements of the FIM, which actually represent informativeness of individual parameters. With this tool at our disposal, we are able to infer on how different datasets affect estimation of individual parameters.

Furthermore, extrinsic calibration seeks for a homogeneous transformation which can be parametrized in a number of ways. Translation can be expressed in any of the two coordinate frames, while orientation can expressed through multiple Euler angle parametrizations. Generally, it may seem counterintuitive that a certain parametrization of the transformation can be preferable to others. However, for our calibration method it is important due to the second optimization step – the RCS optimization. Namely, in that step we do not refine all the parameters estimated in the first step, the reprojection error optimization, but only the poorly estimated parameters (which we will be able to identify with our FIM tests). However, we justify locking the parameters that were well estimated by concentrating the information in them. Our aim is to show, through the FIM tests, that our parametrization has highest concentration of information in the locked parameters for a variety of sensor configurations. This result is a direct consequence of the radar's inability to measure the elevation angle.

5. Experiment

To test the proposed calibration method, we conducted both simulated and real-world data experiments. Through the simulations described in Section 5.1, based on the framework of FIM described in Section 4, we have tested the properties of designed reprojection error. Afterwards, in Sections 5.2 and 5.3 we describe the setup of the conducted experiment and the final results, respectively. Finally, in the Section 5.4, we present a real world application where our calibration method is used to find radar vertical misalignment.

5.1. Simulations

To test our method in simulations under various conditions, we have created a number of different synthetic datasets described with the labelled tuple $\mathcal{D}_{label} = (^{\mathbf{S}}\mathbf{X}, {}^{\mathbf{R}}\mathbf{Y}, \mathcal{T}, N, S)$ where:

- $^{\mathbf{S}}\mathbf{X}$ represents the measurement set originating from a 3D sensor in the sensor coordinate frame $\mathcal{F}_{\mathcal{S}}$,

- $^{\mathbf{R}}\mathbf{Y}$ represents the planar measurement set originating from the radar in the radar coordinate frame $\mathcal{F}_{\mathcal{R}}$,

- $\mathcal{T}$ represents the transformation between $\mathcal{F}_{\mathcal{S}}$ and $\mathcal{F}_{\mathcal{R}}$ which can be parametrized in different forms,

- N represents number of unique measurement points in the dataset,

- S represents number of samples of each unique point in the dataset.

7

For simulation purposes we have assumed a diagonal covariance matrix $\mathbf{Q} = \text{diag}(\sigma^2)$, where $\sigma^2 = 6.25 \times 10^{-4}\,\text{m}^2$. Furthermore, as datasets are comprised of a different number of unique measurements N, S compensates for the total number of used points. It allows a fair comparison of FIMs since amount of information is proportional to the number of points. The measurements are first given in radar's spherical coordinates $^r s_{r,i} = [^r r_{r,i}\ ^r \phi_{r,i}\ ^r \psi_{r,i}]$, with ϕ and ψ being azimuth and elevation, respectively. Afterwards, they are transformed into the 3D sensor frame $^\mathbf{S}\mathbf{X}$, and in radar's planar measurements in the zero-elevation plane $^\mathbf{R}\mathbf{Y}$ (cf. Sec.3.1).

Minimal requirements on the number of measurements is found by examining FIM singular values for the following marginal datasets: $\mathcal{D}_{3CP}$ consists of three ($N = 3$) coplanar, non-colinear points at $^r r_{r,i} = 5\,\text{m}$, $^r \phi_{r,i} = [-45, 0, 45]°$, $^r \psi_{r,i} = 0°$; $\mathcal{D}_{4CP}$ consists of four ($N = 4$) coplanar, non-colinear points at $^r r_{r,i} = 5\,\text{m}$, $^r \phi_{r,i} = [-45, -15, 15, 45]°$, $^r \psi_{r,i} = 0°$; $\mathcal{D}_{4nCP}$ consists of four ($N = 4$) non-coplanar, non-colinear points at $^r r_{r,i} = 5\,\text{m}$, $^r \phi_{r,i} = [-45, -45, 45, 45]°$, $^r \psi_{r,i} = [-5, 5, -5, 5]°$. Additionally, dataset $\mathcal{D}_{FoV}$ consists of $N = 300$ uniformly spread points through the FoV within the following range, azimuth and elevation intervals: $^r r_{r,i} = [4, 5]\,\text{m}$, $^r \phi_{r,i} = [-45, 45]°$, $^r \psi_{r,i} = [-5, 5]°$. It illustrates the upper bound on the achievable parameter informativeness. FIM analysis results for the four datasets are shown in Table 1, where we are striving to have the singular values as large as possible, since it suggests identifiability of the parameters. Note that at this point we are not concerning ourselves which exact parameters are identifiable, but only with if all the 6 parameters of the relative transformation between the coordinate frames are identifiable. By examining smallest singular values, we can see an evident increase ($\sim 10^4$) in the conditional number κ, i.e., the ratio between the largest and smallest singular value, when the non-coplanar point is added to the dataset (note the increase of the smallest singular value σ_6). Difference in the order of magnitude between the largest and smallest singular value for $\mathcal{D}_{4nCP}$ still exists, but unlike the other two datasets, this is not caused by the degeneracy of FIM, i.e., non-identifiability. It is caused by different scales of the parameters, i.e., Euler angles and translation, and uneven sensitivity in the parameters, which is further elaborated in the justification of parametrization choice. This conclusion was also confirmed by the optimization results, since regardless of how big of an S we chose, the reprojection error optimization was unable to converge close to parameter ground truth values for $\mathcal{D}_{3CP}$ and $\mathcal{D}_{4CP}$, while for $\mathcal{D}_{4nCP}$ and $\mathcal{D}_{FoV}$ it always converged successfully. Finally, dataset $\mathcal{D}_{FoV}$ shows that adding more unique points to the dataset does not present a significant impact on the singular values in terms of the identifiability. This brings us to the first important result of the identifiability analysis. To calibrate a radar and a 3D sensor, the previous analysis suggests that to have all the 6 parameters identifiable, the best course of action would be to have at least 4 non-coplanar non-colinear points in the dataset.

The second important result of the identifiability analysis is the justification of the parameter locking in the second optimization step, which, as we will see, is related to parametrization of the relative transformation between the two sensor frames. We have conducted four experiments which differ only in the poses between the sensor coordinate frames and the parametrization of the pertaining transformation. The dataset $\mathcal{D}_{rPs\_0}$ consists of a transformation $\mathcal{T}_{rPs\_0}$ that assumes the simplest case of no rotation and translation between the sensors, while the parametrization is the same as the one defined in Section 3.1 – translation defined as the position of the 3D sensor in $\mathcal{F}_\mathcal{R}$, i.e., radar's coordinate system. The dataset $\mathcal{D}_{sPr\_0}$ differs in the parametrization of the translation. Namely, it is defined as the position of the radar in $\mathcal{F}_\mathcal{S}$. The other two datasets, $\mathcal{D}_{rPs\_45}$ and $\mathcal{D}_{sPr\_45}$, share the same differences in the translation parametrization, but they also assume that there exists a difference in the pitch angle $^s_r\theta_y = 45°$ between the radar and 3D sensor. All the datasets use $N = 300$ unique ($S = 1$) uniformly distributed measurements within the following range, azimuth and elevation intervals: $^r r_{r,i} = [2, 8]\,\text{m}$, $^r \phi_{r,i} = [-75, 75]°$, $^r \psi_{r,i} = [-10, 10]°$.

By analyzing the results for this experiment, which are shown in Table 2, we can notice that the datasets with the same sensor poses, $\mathcal{D}_{rPs\_0}$ and $\mathcal{D}_{sPr\_0}$, but different translation parametrization, exhibit the same FIM results, which confirms uneven uncertainty, or equivalently, uneven informativeness in estimating each parameter. We can see that the yaw angle $^s_r\theta_z$ is significantly more informative than the other Euler angles. Similarly, translations in directions $^r p_{s,x}$ and $^r p_{s,y}$ are more informative compared to the direction $^r p_{s,z}$. For $\mathcal{D}_{rPs\_0}$ and $\mathcal{D}_{sPr\_0}$, the uncertainty is equivalent for directions $^s p_{r,x}$, $^s p_{r,y}$, and $^s p_{r,z}$ since the axes coincide due to the lack of rotation between the sensor frames. However, if we observe datasets $\mathcal{D}_{rPs\_45}$ and $\mathcal{D}_{sPr\_45}$, which include displacement in rotation, we can notice significant differences in FIM diagonal elements for the two translation parametrizations. Namely, when the translation is defined in $\mathcal{F}_\mathcal{R}$, informativeness remains the same as in the previous two cases. However, if the translation is expressed in $\mathcal{F}_\mathcal{S}$, we can notice that informativeness somewhat decreases in the $^s p_{r,x}$ direction, while it increases in the $^s p_{r,z}$ direction, leading to the same informativeness of the two directions.

The main cause for this uneven informativeness of the parameters is radar's inability to measure the elevation angle. To illustrate the assertion, we refer to Fig. 4. We observe the effect on a single measurement, $\mathbf{X} = [2\,\text{m}, 0°, 0°]$, for two cases: when the radar is translated along its $^r x$ and along its $^r z$ axis, yielding new measurements $\mathbf{X_x}$ and $\mathbf{X_z}$, respectively. The 3D sensor and the target from which the measurement originates are kept fixed. Measurement $\mathbf{X_x} = [1.8\,\text{m}, 0°, 0°]$ is acquired by translating the radar along the direction of $^r x$ for $\Delta^r p_{s,x} = 0.2\,\text{m}$, while the measurement $\mathbf{X_z} = [2.01\,\text{m}, 0°, 5.7°]$ is acquired by translating the radar along the direction of $^r z$ for $\Delta^r p_{s,z} = 0.2\,\text{m}$. The only difference that radar detects, in this case, is the change in the range measurement which is significantly smaller in the case of $\mathbf{X_z}$. To generalize, if the radar is displaced along its xy-plane, or rotates around its $^r z$ axis, it would produce significant changes in range or azimuth or both. Meanwhile, the elevation, which is unavailable, would not take away the information about the translation or rotation, which is a case for the changes in the other parameters.

Furthermore, Fig. 4 explains why parametrization $\mathcal{T}_{rPs\_45}$ is

Table 1: FIM singular values for the three datasets used for analyzing the minimum number and the distribution of points in the dataset

	ς_1	ς_2	ς_3	ς_4	ς_5	ς_6	κ
$\mathcal{D}_{3CP}$	1.17×10^7	5.07×10^5	1.84×10^5	8.83×10^3	4.38×10^3	1.58×10^{-1}	7.41×10^7
$\mathcal{D}_{4CP}$	1.17×10^7	5.05×10^5	1.57×10^5	8.98×10^3	3.69×10^3	6.47×10^{-1}	1.81×10^7
$\mathcal{D}_{4nCP}$	1.18×10^7	5.18×10^5	2.59×10^5	5.09×10^4	4.33×10^4	3.70×10^3	3.19×10^3
$\mathcal{D}_{FoV}$	1.01×10^7	4.79×10^5	8.83×10^5	2.15×10^4	4.86×10^3	1.29×10^3	7.83×10^3

Table 2: FIM's diagonal elements corresponding to the informativeness of individual parameters.

	$^s_r\theta_z$	$^s_r\theta_y$	$^s_r\theta_x$	p_x	p_y	p_z
$\mathcal{D}_{rPs\_0}$	1.37×10^7	5.81×10^4	8.28×10^4	4.79×10^5	4.80×10^5	4.81×10^3
$\mathcal{D}_{sPr\_0}$	1.37×10^7	5.81×10^4	8.28×10^4	4.79×10^5	4.80×10^5	4.81×10^3
$\mathcal{D}_{rPs\_45}$	1.37×10^7	5.28×10^4	6.87×10^6	4.78×10^5	4.81×10^5	4.74×10^3
$\mathcal{D}_{sPr\_45}$	1.37×10^7	5.28×10^4	6.87×10^6	2.41×10^5	4.81×10^5	2.41×10^5

preferred to $\mathcal{T}_{sPr\_45}$. Namely, in $\mathcal{T}_{rPs\_45}$, where ^rx coincides with the range, previously elaborated uncertainty has the most spread form. On the other hand, in $\mathcal{D}_{sPr\_45}$, this uncertainty is equally spread between ^sx and ^sz, which is confirmed by the FIM analysis in Tab. 2. Different distribution of uncertainty, due to different parametrization, would merely be a preference if we performed only reprojection error optimization, since it does not provide any more information or lead to better calibration. However, the second step tries to compensate for the lack of radar's elevation angle measurements based on the RCS estimation. Since RCS measurements are less reliable, we do not want to refine parameters which can be properly estimated through reprojection error. Therefore, it is desirable to separate reliable from unreliable parameters, as good as possible, which is the case when translation is given in $\mathcal{F}_\mathcal{R}$ as $^r\mathbf{p}_s$, while there is a rotation which coincides with the rotation around the ^rz axis.

Finally, based on the simulation results, we can give some general advice on the dataset collection. Despite the omnipresent rule, the more the merrier, we would like to emphasize the requirement on the observation of the target at a wide range of radar elevation angle. If we observe points only in the radar plane (zero elevation), we would obviously provide a degenerate case as shown in the minimal requirements tests. As we observe the target at a wider range of elevation angles, we are further away from the singularity. However, when we observe the target at greater elevation angles, there is a risk we might misinterpret the target stand for the actual target. This is best avoided by determining the target stand RCS and performing RCS thresholding. Better results of reprojection error optimization will provide better initial values for the RCS optimization.

5.2. Experiment setup

An outdoor experiment was conducted to test the proposed calibration method. A mobile robot Husky UGV, shown in Fig. 5 was used as the platform for data collection. It was equipped with a Velodyne HDL-32E 3D LiDAR, two short range radars from different manufacturers, namely the Continental SRR 20X and Delphi SRR2, and PointGrey camera sensor combined with Kowa lens with resolution 1920×1080 and
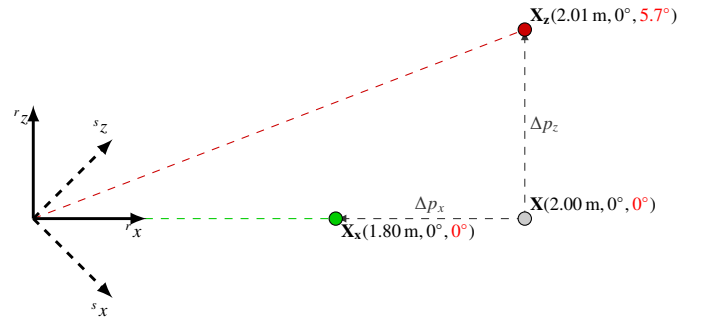

Figure 4: Illustration of unequal uncertainty in parametrization caused by radar's inability to measure elevation angle (indicated in red for reference). $\mathbf{X_x}$ and $\mathbf{X_z}$ show how a single radar measurement in spherical coordinates $\mathbf{X}$ changes when the radar is translated for 0.2 m along ^rx and ^rz, respectively.
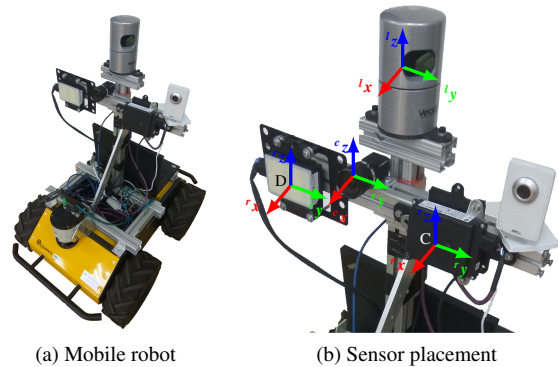

(a) Mobile robot   (b) Sensor placement

Figure 5: Mobile robot and sensors used in the experiment. Marks D and C stand for Delphi and Continental radar, respectively.

$\text{HFoV} \times \text{VFoV} = 60° \times 40°$.

Commercially available radars are sensors which provide high level information in the form of detected object list. Raw data, i.e., the return echo, is processed by proprietary signal processing techniques and is unavailable to the user. However, from the experiments conducted with both radars, we noticed that they follow the behaviour as expected from our calibration method. The only noticed difference is that the target stand without the target was completely invisible to the Continental radar, while the Delphi radar was able to detect it at closer ranges ($^rr_{r,i} < 5\,\text{m}$).

9

Table 3: Continental SRR 20X specifications

Continental SRR 20X	Value
HFoV × VFoV 150° × 12°	
Range Accuracy	0.2 m
Azimut Accuracy @ HFoV	±2°@±20°; ±4°@±60°; ±5°@±75°

Table 4: Delphi SRR2 specifications

Delphi SRR2	Value
HFoV × VFoV	150° × 10°
Range Accuracy	0.5 m (noise error); 2.5% (bias error)
Azimut Accuracy @ HFoV	±1°@±75°

Although the purpose of the experiment is evaluation of the proposed calibration method and not radar performance, we believe it is important to present results for two different radars since they exhibit slightly different behaviour as previously elaborated. Furthermore, RCS optimization uses a novel metric based on a pattern that may not be equal for all the radars. Therefore, success of the calibration using radars from two different manufacturers further confirms the validity of the proposed method.

Technical data of interest for the Continental and Delphi radars is given in Tables 3 and 4, respectively. Based on the analysis of the reprojection error for the Continental radar, radar measurements outside of the azimuth angle range of ±45° were excluded from the optimization, because they exhibited significantly higher reprojection error than those inside the range. For the Delphi radar, measurements outside of the azimuth angle range of ±60° were also excluded due to the observed increase in the reprojection error. The calibration target was composed of a retroreflector with side length $l = 0.32$ m with a maximum RCS of $\sigma_c = 18.75$ dBm^2. Based on the vertical resolution of Velodyne HDL-32E LiDAR (1.33°), we used a styrofoam triangle of height $h = 0.65$ m. This ensured extraction of at least two lines from the target in the experimental data, which is a prerequisite to unambiguously determine the pose. Data acquisition was done by driving a robot in the area of up to 7 m of distance from the target which was placed at 17 different heights ranging from ground level up to a 2 m height. For the Continental radar, 334 registered radar – LiDAR and 227 radar – camera corresponding measurements were collected. For the Delphi radar, 322 registered radar – LiDAR and 193 radar – camera corresponding measurements were collected.

5.3. Experimental results

In this section we present calibration results of four sensor combinations, i.e., two radars combined with the camera and 3D LiDAR. Since the method for the LiDAR – radar and camera – radar differs only in the step of target 3D localization, results of all the experiments are shown simultaneously for both pairs. We have noticed larger difference in calibration results when using different radars, as opposed to calibrating different sensor types with the same radar. Therefore, we first present calibration results for the Continental radar combined with both
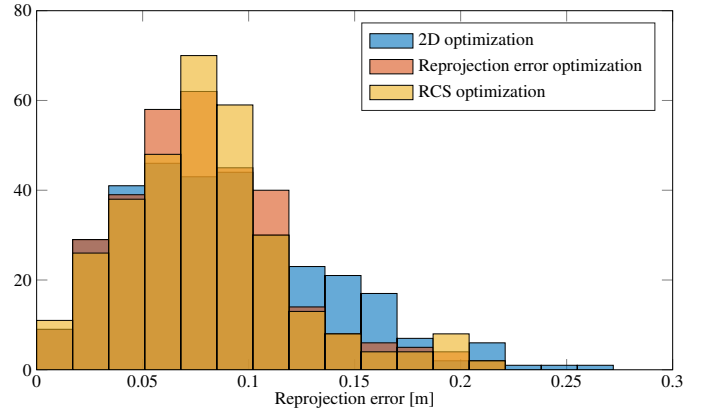

Figure 6: Histogram of reprojection errors for the two steps of the calibration and the 2D calibration for Continental radar – LiDAR calibration.

3D sensors, and then we show results of the calibration involving the Delphi radar.

To assess the quality of calibration results we conducted four experiments. First, we examined the distribution of the reprojection error after both optimization steps and compared it to a 2D optimization that minimizes reprojection error by optimizing only the calibration parameters with lower uncertainty, i.e., translation parameters $^r p_{l,x}$ and $^r p_{l,y}$, and rotation $^s_r \theta_z$. Second, we inspect FoV placement with respect to the distribution of RCS over the 3D sensor's data. Afterwards, we examine the correlation between RCS and the elevation angle. In the end, we run Monte Carlo simulations by random bootstrap resampling with replacement of the dataset, to examine reliability of the estimated parameters and potential overfitting of data.

5.3.1. Continental radar

We obtained the following results for the reprojection error optimization, $^l \hat{\boldsymbol{c}}_r$, RCS optimization $^l \hat{\boldsymbol{c}}_\sigma$, and the carefully hand measured translation, $^r \tilde{\boldsymbol{p}}_l$, for the Continental radar – LiDAR pair:

- $^l \hat{\boldsymbol{c}}_r = [-0.05\,\text{m}, -0.14\,\text{m}, 0.11\,\text{m}, -2.2°, 5.1°, -1.7°]$

- $^l \hat{\boldsymbol{c}}_\sigma = [0.20\,\text{m}, 4.8°, -0.8°, -0.13\,\text{dBm}^2\text{deg}^{-2}, 16.2\,\text{dBm}^2]$

- $^r \tilde{\boldsymbol{p}}_l = [-0.08\,\text{m}, -0.12\,\text{m}, 0.19\,\text{m}]^T$.

Furthermore, for the Continental radar – camera pair, we obtained the following results:

- $^c \hat{\boldsymbol{c}}_r = [0.04\,\text{m}, -0.15\,\text{m}, -0.08\,\text{m}, 0.1°, 5.9°, -2.3°]$

- $^c \hat{\boldsymbol{c}}_\sigma = [0.04\,\text{m}, 5.6°, -1.6°, -0.15\,\text{dBm}^2\text{deg}^{-2}, 16.2\,\text{dBm}^2]$

- $^r \tilde{\boldsymbol{p}}_c = [0.00\,\text{m}, -0.15\,\text{m}, 0.04\,\text{m}]^T$.

Figures 6 and 7 show distribution of the reprojection errors for LiDAR – radar and camera – radar calibrations, respectively. They are composed of three histograms, where we can see how the case of 2D calibration compares to the reprojection error of both steps of the proposed calibration. Besides neglecting three additional DoF, 2D reprojection error assumes that all the measurements reside in the same plane, thus reducing the
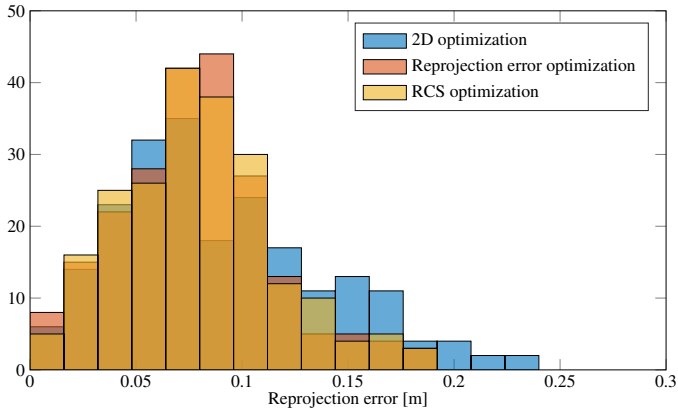
10

Figure 7: Histogram of reprojection errors for the two steps of the calibration and the 2D calibration for Continental radar – camera calibration.

original circle – point relationship to point – point distance. Although this 2D reprojection metric is not the same as the one used for our two steps of optimization, it is the only fair comparison since 2D optimization is based on minimizing it. We can observe that the 2D reprojection error has a larger number of point correspondences with higher reprojection error. These originate from the measurements that are further away from the radar plane because the circle – point relationship has a greater impact than the 2D optimization can explain. Therefore, we conclude that neglecting the 3D nature of the problem causes higher mean of the reprojection error which implies poor calibration. Furthermore, the RCS optimization is bound to degrade the overall reprojection error because it is not a part of the optimization criterion. However, resemblance between the distributions after the first and the second optimization steps implies low degradation. Finally, it can be seen from both the first and second optimization step, that the reprojection error is below the nominal range accuracy of the radar.

In Fig. 8, distribution of the RCS across LiDAR's data is shown, while we omit results for camera since they do not exhibit any significant difference. Measurements from 3D sensors are color-coded with the RCS of the paired radar measurement, while the pose of the radar's nominal FoV is illustrated with blue bounding pyramid. We can see that within the nominal FoV, target produces a strong, fairly constant reflections. As the elevation angle of the target leaves the radar's nominal FoV, the RCS decreases and this effect is the basis of the RCS optimization step.

To examine the effect of decrease in the target's RCS as a function of the elevation angle after both optimizations, we use Fig. 9 and Fig. 10 for LiDAR – radar and camera – radar results, respectively. Each figure shows elevation $^r\psi_{s,i}$ of each 3D sensor measurement transformed into the $\mathcal{F}_r$ and RCS of the paired radar measurement. Furthermore, intrinsic radar curve estimated by the RCS optimization is plotted. In the ideal case, i.e., if the transformation was correct and the axis of retroreflector always pointed directly to the radar, the data would lay on the curve which describes radar's radiation pattern with respect to the elevation angle. The dispersion around the curve is present in both steps due to imperfect directivity of the target

and measurement noise. We have evaluated directivity of the target towards the radar after the calibration and noticed that all the measurements differed less than 18° from the ideal directivity in the context of maximum response. According to experimental results in [24], such small angles do not reduce retroreflector's RCS significantly, which, combined with errors in directivity estimation, prevents us from performing directivity compensation. We assert that this is not crucial in our case. However, if the experiment is performed in such way that corner retroreflector orientation differs significantly from the ideal, we believe that RCS directivity compensation would be necessary. From the plots, we can notice that dispersion of the measurements after the reprojection error optimization is higher compared to the case of RCS optimization. This effect is caused due to the poor estimation of the parameters with higher uncertainty which are corrected by the RCS optimization.

In the end, we performed Monte Carlo analysis to test how sensitive our parameter estimates are to the available dataset. We performed random bootstrap resampling with replacement of our dataset. Optimization was performed in 1000 runs on different randomly sampled datasets from which we observe the estimated extrinsic calibration parameters. The results follow a Gaussian distribution whose estimated parameters are given in Table 5 for the LiDAR – radar calibration and the Table 6 for the camera – radar calibration. As expected, distributions of parameters $^rp_{s,x}$, $^rp_{s,y}$ and $^s_r\theta_z$ obtained by the reprojection error optimization have significantly lower variance than the other parameters. Figures 11 and 12 illustrate how the RCS optimization refines parameters $^rp_{s,z}$, $^s_r\theta_y$ and $^s_r\theta_x$. We can see significant decrease in variance, as well as the shift in the mean. For the purposes of distribution visualisation, we have reduced the bin size for RCS optimization to 10% compared to the results for the reprojection error optimization. Otherwise, all the results would fall within one bin due to the significantly lower variance. Estimation of the mean of $^rp_{s,z}$ using the reprojection error optimization is clearly further away from the measured value than the RCS optimization's estimate. Estimation of $^rp_{c,z}$ is fairly close to the measured value, while $^rp_{l,z}$ exhibits a slight bias of 2 cm from the measured value. The cause of the bias could be imprecise hand-measurement or the systematic errors in the LiDAR's estimates of the retroreflector's position. Furthermore, a bias compared to the hand-measured values is also visible in the estimation of both $^rp_{l,x}$ and $^rp_{c,x}$. We believe that it could originate from the bias in the radar's range measurements or the imprecision of the target design. However, when introduced to the reprojection error optimization as a parameter, it could not be distinguished between the translation parameters.

5.3.2. Delphi radar

From the experimental results for the Delphi radar, we noticed that estimation of the $^rp_{s,x}$ exhibited a noticeable offset from the measured value. Although the exact origin of this bias is uncertain, we hypothesize that most likely it is an effect caused by the target stand. Namely, since the Delphi radar is able to detect the stand, proprietary algorithms that determine the range of the observed object could infer that a target is at the greater range. To address this issue, the reprojection error estimation was expanded with estimation of the
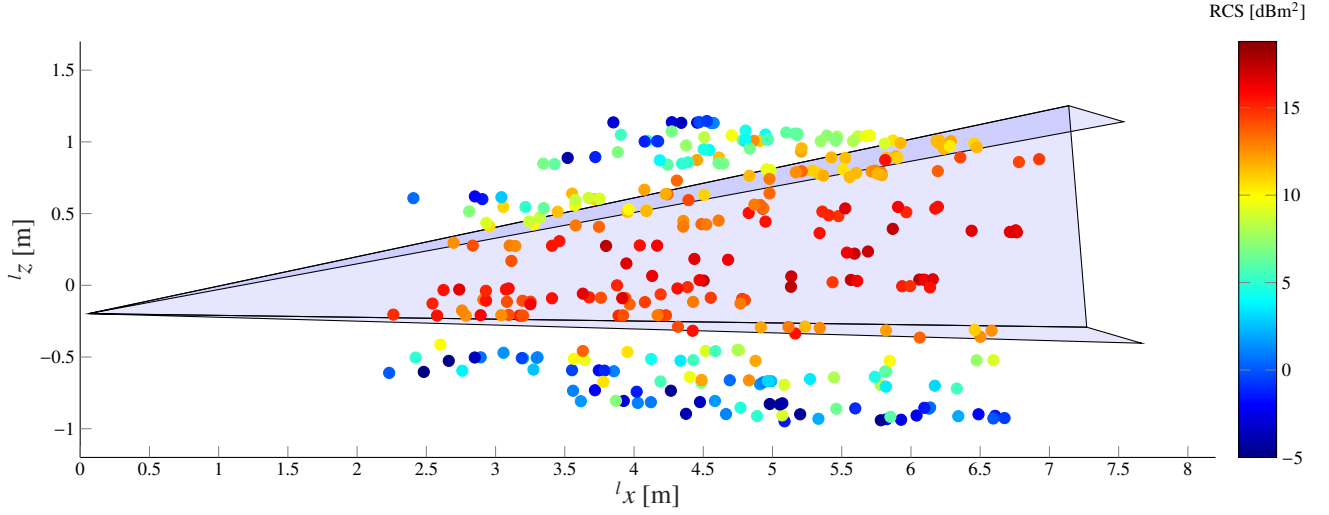
Figure 8: RCS distribution across LiDAR 3D data and placement of the Continental radar's FoV.
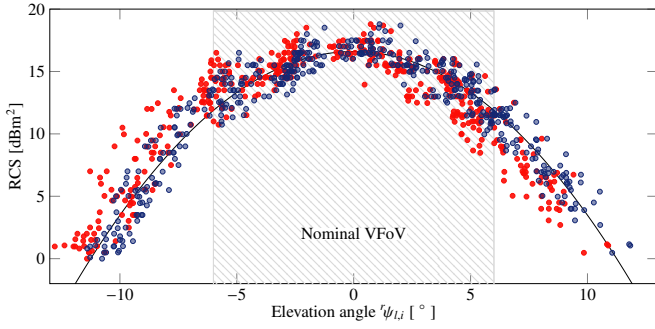

Figure 9: RCS distribution across radar's VFoV for Continental radar – LiDAR calibration. Red: reprojection error optimization; blue: RCS optimization.
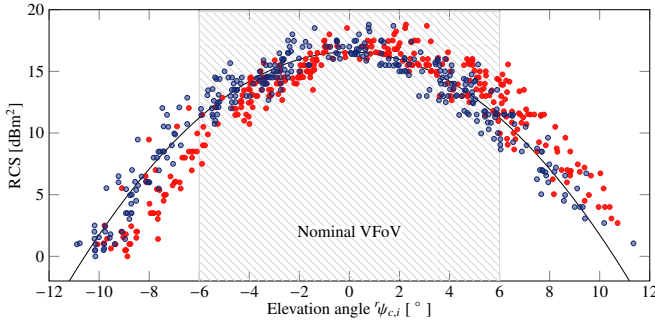

Figure 10: RCS distribution across radar's VFoV for Continental radar – camera calibration. Red: reprojection error optimization; blue: RCS optimization.
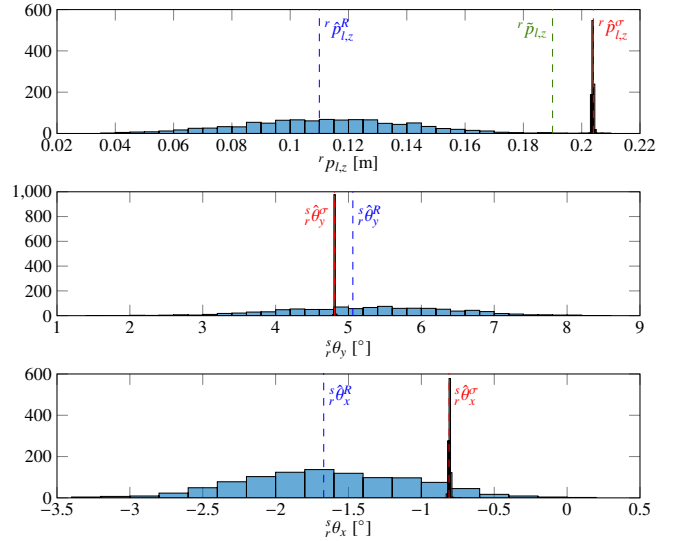

Figure 11: Monte Carlo analysis results for Continental radar – LiDAR calibration. Blue: calibration after reprojection error optimization; red: with RCS optimization.

Table 5: Monte Carlo analysis results for Continental radar – LiDAR calibration.

	Reprojection Error Optimization	RCS optimization
$^r p_{l,x}$ [m]	$\mathcal{N}(-0.050, 2.36 \times 10^{-5})$	
$^r p_{l,y}$ [m]	$\mathcal{N}(-0.134, 8.04 \times 10^{-5})$	
$^r p_{l,z}$ [m]	$\mathcal{N}(0.113, 8.46 \times 10^{-4})$	$\mathcal{N}(0.204, 1.48 \times 10^{-7})$
$^l_r\theta_z$ [°]	$\mathcal{N}(-2.21, 1.61 \times 10^{-2})$	
$^l_r\theta_y$ [°]	$\mathcal{N}(5.29, 1.29)$	$\mathcal{N}(4.81, 2.32 \times 10^{-5})$
$^l_r\theta_x$ [°]	$\mathcal{N}(-1.63, 3.39 \times 10^{-1})$	$\mathcal{N}(-0.81, 4.29 \times 10^{-5})$

range offset that is subtracted from the radar range measurements, $c_r = [^r \boldsymbol{p}_s \ \boldsymbol{\Theta}_s, \Delta^r r_r]$.

We obtained the following results for the reprojection error optimization, $^l\hat{\boldsymbol{c}}_r$, RCS optimization $^l\hat{\boldsymbol{c}}_\sigma$, and the carefully hand measured translation, $^r\tilde{\boldsymbol{p}}_l$, for the Delphi radar – LiDAR pair:

- $^l\hat{\boldsymbol{c}}_r = [-0.07\,\text{m}, 0.13\,\text{m}, 0.11\,\text{m}, -2.9°, 5.0°, 7.6°, 0.10\,\text{m}]$

- $^l\hat{\boldsymbol{c}}_\sigma = [0.21\text{m}, 2.0°, -0.2°, -0.25\,\text{dBm}^2\text{deg}^{-2}, 17.9\,\text{dBm}^2]$

- $^r\tilde{\boldsymbol{p}}_l = [-0.08\,\text{m}, 0.15\,\text{m}, 0.20\,\text{m}]^T$.

Furthermore, for the Delphi radar – camera pair, we obtained the following results:

- $^c\hat{\boldsymbol{c}}_r = [0.02\,\text{m}, 0.11\,\text{m}, -0.01\,\text{m}, -0.1°, 4.7°, 7.3°, 0.11\,\text{m}]$

- $^c\hat{\boldsymbol{c}}_\sigma = [0.02\,\text{m}; 2.6°, -0.8°, -0.27\,\text{dBm}^2\text{deg}^{-2}, 17.4\,\text{dBm}^2]$

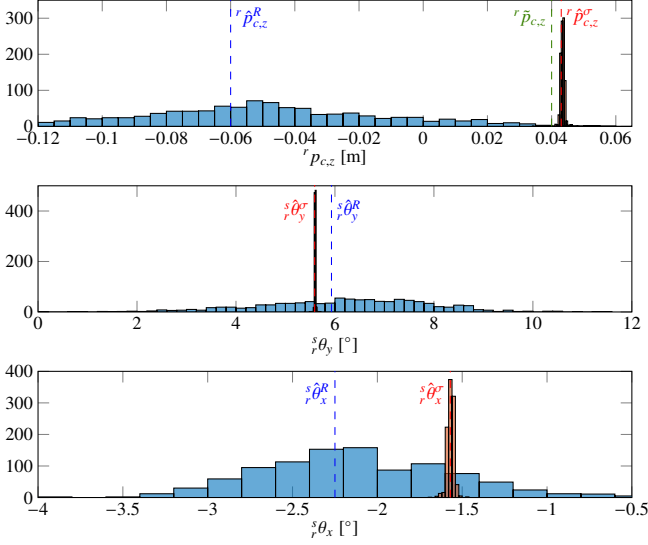- $^r\tilde{\boldsymbol{p}}_c = [0.00\,\text{m}, 0.12\,\text{m}, 0.05\,\text{m}]^T$.

Figure 12: Monte Carlo analysis results for Continental radar – camera calibration. Blue: calibration after reprojection error optimization; red: with RCS optimization.

Table 6: Monte Carlo analysis results for Continental radar – camera calibration.

	Reprojection Error Optimization	RCS optimization
$^r p_{c,x}$ [m]	$\mathcal{N}(0.039, 2.37 \times 10^{-5})$	
$^r p_{c,y}$ [m]	$\mathcal{N}(-0.148, 1.96 \times 10^{-4})$	
$^r p_{c,z}$ [m]	$\mathcal{N}(-0.051, 1.48 \times 10^{-3})$	$\mathcal{N}(0.043, 7.48 \times 10^{-7})$
$^c_r \theta_z$ [°]	$\mathcal{N}(0.12, 4.44 \times 10^{-2})$	
$^c_r \theta_y$ [°]	$\mathcal{N}(6.21, 2.91)$	$\mathcal{N}(5.60, 1.65 \times 10^{-4})$
$^c_r \theta_x$ [°]	$\mathcal{N}(-2.11, 3.26 \times 10^{-1})$	$\mathcal{N}(-1.57, 4.07 \times 10^{-4})$

Results for the Continental radar showed that there is no significant difference between calibration of LiDAR – radar pair and camera – radar pair; therefore, for brevity, we focus on the results of LiDAR – radar calibration. Reprojection error histograms exhibited similar results to the case of the Continental radar when comparing reprojection error optimization, RCS optimization, and 2D reprojection optimization results; hence, they are not repeated here. However, an interesting effect was noticed when comparing the reprojection error optimization estimating the range bias with the reprojection error optimization omitting the bias. Figure 13 shows that when the bias is included, average reprojection error per correspondence is reduced from 0.056 m to 0.045 m. On the other hand, for the case of the Continental radar, estimation of the bias compromised the results when calibrating the radar with a 3D sensor and was not able to reduce the average reprojection error. It can be concluded that with the Delphi radar, an actual bias is present, most likely due to the target design, and the method is able to converge to a local minimum of a significantly lower cost. Success of the RCS optimization is most evident in Fig. 14, where we can see a significant difference in the RCS distribution after two steps of optimization. The origin of this mismatch is convergence to poor values of the less certain parameters in reprojection error optimization. From the figure, one could conclude that there is no pattern in the data after the reprojection error optimization. However, the RCS optimization is able to
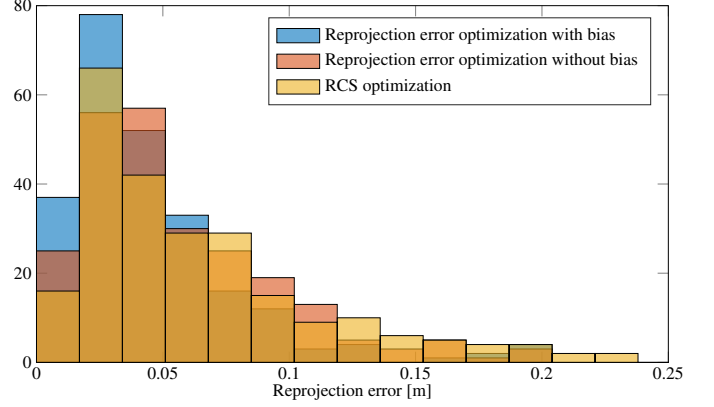

Figure 13: Histogram of reprojection errors for two types of reprojection error optimization and RCS optimization for Delphi radar – LiDAR calibration.
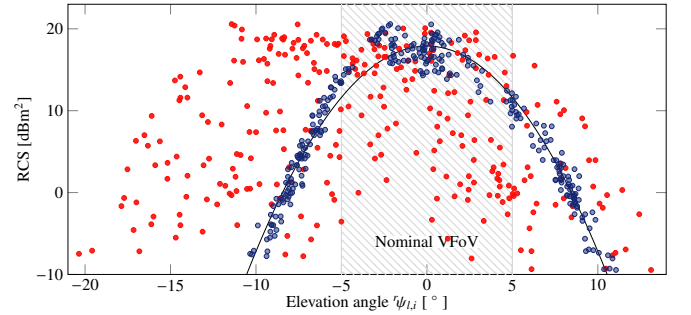

Figure 14: RCS distribution across radar's VFoV for Delphi radar – LiDAR calibration. Red: reprojection error optimization; blue: RCS optimization.

Table 7: Monte Carlo analysis results for Delphi radar – LiDAR calibration.

	Reprojection Error Optimization	RCS optimization
$^r p_{l,x}$ [m]	$\mathcal{N}(-0.064, 8.75 \times 10^{-5})$	
$^r p_{l,y}$ [m]	$\mathcal{N}(0.132, 2.70 \times 10^{-5})$	
$^r p_{l,z}$ [m]	$\mathcal{N}(0.113, 5.79 \times 10^{-4})$	$\mathcal{N}(0.208, 3.08 \times 10^{-6})$
$^l_r \theta_z$ [°]	$\mathcal{N}(-2.93, 5.98 \times 10^{-3})$	
$^l_r \theta_y$ [°]	$\mathcal{N}(4.92, 1.01)$	$\mathcal{N}(2.02, 5.09 \times 10^{-4})$
$^l_r \theta_x$ [°]	$\mathcal{N}(7.50, 2.22 \times 10^{-1})$	$\mathcal{N}(-0.18, 4.85 \times 10^{-5})$
$\Delta^r r_r$ [m]	$\mathcal{N}(0.101, 5.79 \times 10^{-5})$	

find the same quadratic pattern as with the Continental radar without significantly degrading the reprojection error, as seen in Fig. 13.

Finally, Fig. 15 and Table 7 present results for the Monte Carlo analysis. The results are similar to those of the Continental radar, although a slight increase in variance can be seen in the estimation of $^r p_{l,x}$. The cause for the increase could be the performance of the radar or the coupling of the range and bias estimation.

5.4. Radar vertical alignment

In Section 1 we outlined the importance of proper radar vertical alignment, and in this section we present a simple, yet reliable mehod for its assessment. The proposed method requires precise 6DoF extrinsic calibration. Thus, we compared our method, labelled *RCS*, to the other 6DoF calibration method [26], labelled *MAN*, which manually searches for RCS maximums and artificially assigns zero elevation angle to these radar
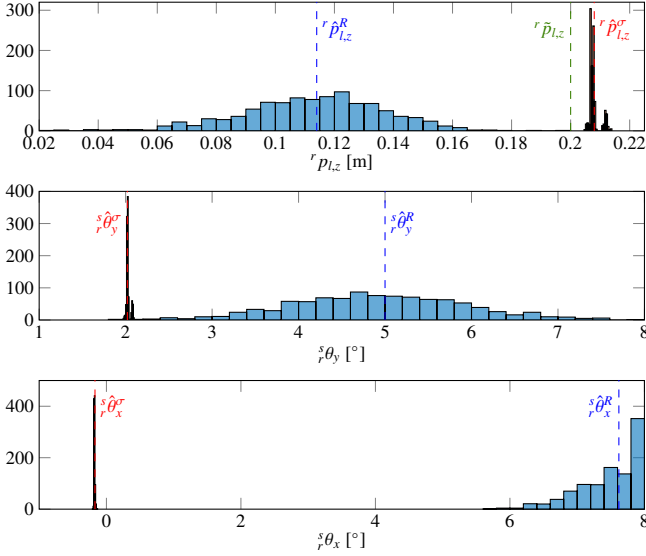
Figure 15: Monte Carlo analysis results for Delphi radar – LiDAR calibration. Blue: calibration after reprojection error optimization; red: with RCS optimization.
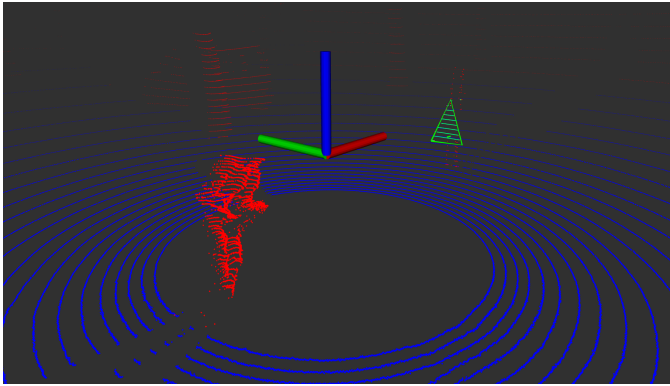

Figure 16: Detected ground plane (blue), target(green) and environment clutter (red) with LiDAR in the purposes of vertical misalignment test.

Table 8: Monte Carlo analysis of vertical misalignment assessment for the Continantal radar using LiDAR-s ground plane estimation with extrinsic calibration results.

	MAN	RCS
$_g^r\theta_y[°]$	$\mathcal{N}(-0.91, 8.86)$	$\mathcal{N}(-4.64, 3.40 \times 10^{-5})$
$_g^r\theta_x[°]$	$\mathcal{N}(6.40, 3.12)$	$\mathcal{N}(0.70, 3.69 \times 10^{-5})$

dinate frame. We set the arbitrary yaw angle around the ground plane normal to $_g^r\theta_z = 0°$ and determined the pitch and roll angles $_g^r\theta_y$ and $_g^r\theta_x$, respectively. The method was tested using the same Monte Carlo analysis described in the Section 5.3 through $N = 1000$ runs.

From the Tab. 8, we can see that *MAN* method produced results with high uncertainty, which is inadequate for vertical misalignment assessment. Namely, the *MAN* method estimated pitch angles in the interval $_g^r\theta_y = [-11.12, 8.19]°$ which surpasses common allowable vertical misalignments, thus providing unreliable misalignment correction guidelines. On the other hand, our method produced stable estimation of ground to radar plane angles, e.g. pitch angles in the interval $_g^r\theta_y = [-4.68, -4.62]°$.

From the results, we can see the *RCS* method estimated vertical misalignment which surpasses allowable tolerance. Namely, Continental SRR20X user manual specifies allowable mounting pitch angle of $\pm 1°$. To elaborate, vertical misalignment causes reduction in range and thus probability of detection. For instance, pitch misalignment of $_g^r\theta_y = 4.5°$ causes a 25% decrease in range for the Delphi SRR2, while such misalignment causes decrease of 80% for the long range radar Delphi ESR. Therefore, radar mounting on a vehicle is a crucial step where our method can provide helpful guidelines. Given that, we conclude that for our sensor setup we should correct the orientation of the Continetal radar according to results because the misalignment would impair the performance.

6. Conclusion

In this paper we have presented a method for extrinsic calibration of a LiDAR – camera – radar sensor system. A special calibration target design was developed to enable all the sensors to detect and accurately localize the target. The extrinsic calibration is based on the proposed two-step optimization procedure which involved: (i) optimization of a reprojection error based on the point-circle constraint which captures radar's lack of elevation angle measurements, and (ii) RCS optimization based on a pattern found in the radar's RCS estimation – again caused by the lack of the elevation angle resolution across substantial FoV thereof. Throughout the identifiability analysis, we have shown that the proposed point-circle geometric constraint requires minimum of 4 non-coplanar points to become identifiable, while the experimentally discovered effect of uneven uncertainty in the extrinsic parameters was confirmed by the FIM analysis. We presented the experimental results for LiDAR and camera sensors in combination with two radars from different manufacturers and have also addressed the radar vertical misalignment problem. In the end, through extensive experimental analysis, we have shown that the proposed method is able to accurately estimate all the six DoF of the extrinsic calibration.

measurements. From Fig. 9, we can see that the measured target reports RCS in the range of $[16, 19]$ dBm^2 at the zero elevation angle. Therefore, for the *MAN* method, we used only the correspondences that surpass the RCS threshold of $\sigma_{th} = 16$ dBm^2, resulting in 32 correspondence groups total. Even though the low number of correspondence groups available from the experiment could affect the accuracy of the *MAN* method, we can see that only a small fraction of measurements could be used in the optimization, thus leading to an expensive calibration data collection.

To find the vertical misalignment we drove the robot and detected the ground plane using a LiDAR, as illustrated by Fig. 16. We used LiDAR for simplicity, but the ground plane can also be found using a single camera as well [44]. The estimated ground plane normals $^l\boldsymbol{n}_{gp}$ from a 2-minute drive were averaged to remove the effects of uneven ground and robot rotation. The averaged normal $^l\boldsymbol{n}_{gp}$ was transformed to the radar $^r\boldsymbol{n}_{gp}$ coordinate frame using the estimated LiDAR to radar extrinsic calibration (for both the *RCS* and *MAN* method). Finally, we expressed the rotation between the ground and radar plane in the radar's coor-

References

[1] R. Y. Tsai, A Versatile Camera Calibration Technique for High-Accuracy 3D Machine Vision Metrology Using Off-the-Shelf TV Cameras and Lenses, IEEE Journal on Robotics and Automation 3 (4) (1987) 323–344.

[2] Z. Zhang, A Flexible New Technique for Camera Calibration (Technical Report), IEEE Transactions on Pattern Analysis and Machine Intelligence 22 (11) (2002) 1330–1334.

[3] D. Scaramuzza, A. Martinelli, R. Siegwart, A toolbox for easily calibrating omnidirectional cameras, in: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2006, pp. 5695–5701.

[4] J. Heikkila, O. Silven, A four-step camera calibration procedure with implicit image correction, IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (1997) 1106–1112

[5] B. Li, L. Heng, K. Koser, M. Pollefeys, A multiple-camera system calibration toolbox using a feature descriptor-based calibration pattern, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (2013) 1301–1307

[6] G. Atanacio-Jiménez, J.-J. . González-Barbosa, J. B. . Hurtado-Ramos, F. . J. Ornelas-Rodríguez, H. . Jiménez-Hernández, T. . García-Ramirez, R. González-Barbosa, LIDAR Velodyne HDL-64E Calibration Using Pattern Planes, International Journal of Advanced Robotic Systems 8 (5) (2011) 70–82.

[7] N. Muhammad, S. Lacroix, Calibration of a rotating multi-beam Lidar, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (2010) 5648–5653

[8] E. Fernández-Moral, J. González-Jiménez, V. Arévalo, Extrinsic calibration of 2D laser rangefinders from perpendicular plane observations, International Journal of Robotics Research 34 (11) (2015) 1401–1417.

[9] Q. Z. Q. Zhang, R. Pless, Extrinsic calibration of a camera and laser range finder (improves camera calibration), in: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2004, pp. 2301–2306.

[10] G. Pandey, J. McBride, S. Savarese, R. Eustice, Extrinsic calibration of a 3D laser scanner and an omnidirectional camera, in: IFAC Symposium on Intelligent Autonomous Vehicles, 2010, pp. 336–341.

[11] L. Zhou, Z. Deng, Extrinsic calibration of a camera and a lidar based on decoupling the rotation from the translation, IEEE Intelligent Vehicles Symposium (IV) (2012) 642–648

[12] A. Geiger, F. Moosmann, O. Car, B. Schuster, Automatic camera and range sensor calibration using a single shot., in: IEEE Conference on Robotics and Automation (ICRA), 2012, pp. 3936–3943.

[13] F. M. Mirzaei, D. G. Kottas, S. I. Roumeliotis, 3D LIDAR-camera intrinsic and extrinsic calibration: Identifiability and analytical least-squares-based initialization, Int. Journal of Robotics Research 31 (4) (2012) 452–467.

[14] E. Olson, AprilTag: A robust and flexible visual fiducial system, International Conference on Robotics and Automation (ICRA) (2011) 3400–3407

[15] J. L. Owens, P. R. Osteen, K. Daniilidis, MSG-cal: Multi-sensor graph-based calibration, in: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2015, pp. 3660–3667.

[16] M. Velas, M. Spanel, Z. Materna, A. Herout, Calibration of RGB Camera With Velodyne LiDAR, WSCG 2014 Communication Papers (2014) 135–144.

[17] K. Kwak, D. F. Huber, H. Badino, T. Kanade, Extrinsic calibration of a single line scanning lidar and a camera, in: IEEE International Conference on Intelligent Robots and Systems (ICRA), 2011, pp. 3283–3289.

[18] X. Wang, L. Xu, H. Sun, J. Xin, N. Zheng, On-Road Vehicle Detection and Tracking Using MMW Radar and Monovision Fusion, IEEE Transactions on Intelligent Transportation Systems 17 (7) (2016) 2075–2084.

[19] J. Ćesić, I. Marković, I. Cvišić, I. Petrović, Radar and stereo vision fusion for multitarget tracking on the special Euclidean group, Robotics and Autonomous Systems 83 (2016) 338–348.

[20] H. Cho, Y.-w. Seo, B. V. K. V. Kumar, R. R. Rajkumar, A Multi-Sensor Fusion System for Moving Object Detection and Tracking in Urban Driving Environments, IEEE International Conference on Robotics and Automation (ICRA) (2014) 1836–1843.

[21] D. Nuss, T. Yuan, G. Krehl, M. Stuebler, S. Reuter, K. Dietmayer, Fusion of laser and radar sensor data with a sequential Monte Carlo Bayesian occupancy filter, in: IEEE Intelligent Vehicles Symposium (IV), 2015, pp. 1074–1081.

[22] R. O. Chavez-Garcia, O. Aycard, Multiple Sensor Fusion and Classification for Moving Object Detection and Tracking, IEEE Transactions on Intelligent Transportation Systems 17 (2) (2016) 525–534.

[23] T. D. Vu, O. Aycard, F. Tango, Object perception for intelligent vehicle applications: A multi-sensor fusion approach, in: IEEE Intelligent Vehicles Symposium (IV), 2014, pp. 774–780.

[24] E. F. Knott, Radar Cross Section Measurements., ITP Van Nostrand Reinhold, 1993.

[25] T. Wang, N. Zheng, J. Xin, Z. Ma, Integrating millimeter wave radar with a monocular vision sensor for on-road obstacle detection applications, Sensors 11 (9) (2011) 8992–9008.

[26] S. Sugimoto, H. Tateda, H. Takahashi, M. Okutomi, Obstacle detection using millimeter-wave radar and its visualization on image sequence, in: International Conference on Pattern Recognition (ICPR), 2004, pp. 342–345.

[27] O. Schwindt, K. Buckner, B. Chakraborty, Systems and Methods for Radar Vertical Misalignment Detection, US 2016/0223649 A1, (2016).

[28] R. Hellinger, O. F. Schwindt, Automotive Radar Alignment, US 2017/0212215 A1, (2017)

[29] B. K. Park, K. K. Im, H. Chang Ahn, Alignment Method And System for Radar of Vehicle, US 9,523,769 B2, (2016)

[30] R. Hermann, A. J. Krener, Nonlinear Controllability and Observability, IEEE Transactions on Automatic Control 22 (5) (1977) 728–740.

[31] F. M. Mirzaei, S. I. Roumeliotis, A Kalman-filter-based algorithm for IMU-camera calibration: observability analysis and performance evaluation., IEEE Transactions on Robotics 24 (5) (2008) 1143–1156.

[32] J. Kelly, G. S. Sukhatme, Visual-Inertial Sensor Fusion: Localization, Mapping and Sensor-to-Sensor Self-calibration, The International Journal of Robotics Research 30 (1) (2011) 56–79.

[33] J. Maye, H. Sommer, G. Agamennoni, R. Siegwart, P. Furgale, Online self-calibration for robotic systems, The International Journal of Robotics Research 35 (4) (2015) 357–380.

[34] J. Peršić, I. Marković, I. Petrović, Extrinsic 6Dof Calibration of 3D Li-DAR and Radar, in: European Conference on Mobile Robotics (ECMR), 2017, pp. 165–170.

[35] Y. Park, S. Yun, C. S. Won, K. Cho, K. Um, S. Sim, Calibration between color camera and 3D LIDAR instruments with a polygonal planar board, Sensors 14 (3) (2014) 5333–5353.

[36] S. Debattisti, L. Mazzei, M. Panciroli, Automated extrinsic laser and camera inter-calibration using triangular targets, in: IEEE Intelligent Vehicles Symposium (IV), 2013, pp. 696–701.

[37] C. G. Stephanis, D. E. Mourmouras, Trihedral rectangular ultrasonic reflector for distance measurements, NDT&E international 28 (2) (1995) 95–96.

[38] P. Furgale, J. Rehder, R. Siegwart, Unified Temporal and Spatial Calibration for Multi-Sensor Systems, in: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2013, pp. 1280–1286.

[39] C. Jauffret, Observability and Fisher information matrix in nonlinear regression, IEEE Transactions on Aerospace and Electronic Systems 43 (2) (2007) 756–759.

[40] J. D. Stigter, J. Molenaar, A fast algorithm to assess local structural identifiability, Automatica 58 (2015) 118–124.

[41] T. J. Rothenberg, Identification in Parametric Models, Econometrica, 39 (3) (1971) 577–591.

[42] K. S. Arun, T. S. Huang, S. D. Blostein, Least-Squares Fitting of Two 3-D Point Sets., IEEE Transactions on Pattern Analysis and Machine Intelligence 9 (5) (1987) 698–700.

[43] G. H. Golub, C. F. Van Loan, Matrix Computations, 2013.

[44] J. Arrospide, L. Salgado, M. Nieto, R. Mohedano, Homography-based ground plane detection using a single on-board camera, IET Intelligent Transport Systems 4 (2) (2010) 149.