

Extrinsic calibration of 2D laser rangefinders from perpendicular plane observations

Eduardo Fernández-Moral*, Javier González-Jiménez and Vicente Arévalo

Universidad de Málaga. MAPIR Group. E.T.S. de Ingeniería Informática, Campus de Teatinos s/n, 29071, Málaga, Spain.

Abstract

Many applications in the fields of mobile robotics and autonomous vehicles employ two or more 2D laser rangefinders (LRFs) for different purposes: navigation, obstacle detection, 3D mapping or SLAM. The extrinsic calibration between such sensors (i.e. finding their relative poses) is required to exploit effectively all the sensor measurements and to perform data fusion. In the literature, most works employing several LRFs obtain their extrinsic calibration from manual measurements or from *ad-hoc* solutions. In this paper we present a new method to obtain such calibration easily and robustly by scanning perpendicular planes (typically corners encountered in structured scenes), from which geometric constraints are inferred. This technique can be applied to a rig with any number of LRFs in almost any geometric configuration (a minimum of two LRFs whose scanning planes are not parallel is required). Experimental results are presented with synthetic and real data to validate our proposal. A C++ implementation of this method and a dataset are also provided.

Keywords: Extrinsic calibration, sensor fusion, 2D laser rangefinder

1 INTRODUCTION

The extrinsic calibration of several 2D laser rangefinders or LIDAR is of very practical interest for autonomous vehicles and for mobile robotics. Such calibration is required to put all the measurements in a common reference frame, which is a prerequisite to perform other tasks as localization, mapping or naviga-




Figure 1: Observation of a corner structure by a rig with two LRFs.

tion. Combinations of LRFs have been employed for 3D mapping in outdoor (Borrmann et al., 2008; Barber et al., 2008; Haala et al., 2008) and indoor environments (Thrun et al., 2000), and also for safe navigation (Victorino et al., 2003). This calibration problem becomes more relevant with the advent of autonomous cars (Thrun et al., 2006; Campbell et al., 2010; Bohren et al., 2008; Petrovskaya and Thrun, 2009; Miller et al., 2011; Leonard et al., 2008), where the information provided by such sensors is essential to avoid possible collisions.

This paper presents a novel solution for the general problem of extrinsic calibration of 2D LRFs, which is based on the observation of perpendicular planes. These planes may be already present in the environment (something common in structured scenes), as it is the case in our experiments, or they may constitute a calibration pattern built for this purpose. The calibration is constrained by imposing co-planarity and perpendicularity constraints on the line segments extracted

*Corresponding author: eduardofernandez@uma.es

by the different laser scanners. A rough approximation of the sensor relative poses must be provided, which can be guessed from simple visual inspection of the rig. This method can be used to calibrate any set of rigidly jointed LRFs where there are at least two sensors with non-parallel scanning planes.

In summary, the necessary conditions for our method to work are:

- the LRF rig observes at least two perpendicular planes from two different viewing directions
- not all the LRFs scan parallel planes, in other words, the measurement planes of at least two sensors must intersect

1.1 Related works

Among the robotic systems found in the literature that employ a combination of LRFs, only a few of them report a calibration technique (Huang et al., 2010; Blanco et al., 2009; Gao and Spletzer, 2010). Many other works like (Thrun et al., 2006; Miller et al., 2011; Campbell et al., 2010; Bohren et al., 2008; Petrovskaya and Thrun, 2009), do not report any calibration process, so, it is reasonable to suppose that they obtain the sensor's relative poses from manual measurements on their setups, like in (Blanco-Claraco et al., 2014). Such procedures are prone to errors that may severely affect the performance of mapping and navigation methods, especially when the laser scanners have a long working range, so that small rotation errors can produce significant distortions in the map (Miller et al., 2011; Leonard et al., 2008). Apart from the limitations in accuracy, measuring the sensors relative poses by hand is also tedious and time consuming.

Generally, the preferred strategy to calibrate exteroceptive sensors is to use their own measurements to establish some kind of data association between their observations. The extrinsic calibration of 2D range scanners in arbitrary poses proves to be more difficult than for RGB or depth cameras, since distinctive features are significantly more scarce in the first. This problem has been tackled in some particular scenarios, but the solutions reported share one or more of the next limitations: they need supervised data association in controlled conditions, they need external information (extra sensors, a pattern or landmarks placed manually in the environment), or they are specific for a particular configuration of the sensor rig. For example, vertical posts of traffic signs are segmented and matched in a supervised way in

(Huang et al., 2010). In (Gao and Spletzer, 2010), a solution is presented based on the matching of reflecting landmarks that are manually placed in the environment. Without using particular targets, calibration is achieved in (Blanco et al., 2009) by making use of the vehicle's odometry to maximize the fitting of the 3D point clouds built from the different LRFs, what requires extra sensors (cameras or precision GPS) to improve the accuracy of the vehicle's odometry.

Another approach consists of matching the trajectories of dynamic objects (or people) in the scene (Schenk et al., 2012), (Glas et al., 2010). For that, the trajectory of one, or several objects, is tracked independently by each LRF, and these trajectories are registered to constraint the sensor's relative poses. This solution is suitable for static systems where all the LRFs scan a common space (nearly in the same plane), but it is not valid to calibrate LRFs in arbitrary poses. Our previous approach (Fernández-Moral et al., 2015) makes use of planar surfaces from the environment, but it requires at least three LRFs observing simultaneously the same plane, and presents further restrictions in its applicability with respect to the method presented here. Also, a broader description of several methods to calibrate the extrinsic parameters of different kinds of range sensors can be found in (Fernández-Moral, 2014).

Ego-motion has been used in combination with wheel odometry to determine the intrinsic parameters of the odometry together with the relative pose of a laser sensor with respect to the robot's frame (Censi et al., 2013). However, this strategy is only applicable when the laser scanner moves in its own plane of measurement (typically, planar movement of a vehicle with an horizontal LRF), otherwise the laser ego-motion cannot be estimated. This last work addresses a different, but complementary problem to the one tackled here. Therefore, a combination of this technique with the one proposed here would be interesting in many problems to calibrate several LRFs with respect to the vehicle's frame, what is required in the field of autonomous cars.

1.2 Contribution

This paper presents, to the best of our knowledge, the first general approach to calibrate a set of LRFs in arbitrary positions (with a minimum of 2 LRFs scanning non-parallel planes). This solution only requires observing perpendicular planes from any structured scene. The observability of the problem, formulated as a maximum likelihood estimation, is analysed from the rank

of the Fisher Information Matrix, and a minimal solution using a single observation of the rig is presented. The convergence of the method is also studied. A C++ implementation of this method is also provided together with a testing dataset¹².

In comparison to previous approaches, our method is applicable to almost any geometric configuration of the sensors; it is accurate and fast, indeed, the calibration can be achieved from a single observation; and finally, our method provides an estimation of the calibration uncertainty.

In the following section we describe the calibration method using constraint equations derived from the coplanarity and perpendicularity of the observed planes (section 2). Section 3 addresses the optimization problem stated upon these constraints within a probabilistic framework that takes into account the precision of the sensor measurements. We analyse the observability conditions (section 4) and study the convergence region for the solution (section 5). Experimental results are presented to validate our approach with both simulated and real data (section 6). Finally, a discussion on our method’s applicability is presented (section 7) and the conclusions are outlined (section 8).

2 CALIBRATION APPROACH

Our proposal for finding the extrinsic calibration (i.e. relative poses) between two or more LRFs relies on establishing geometric constraints from the simultaneous observation of pairs of perpendicular planes (see figure 1). For readability, let us define a corner as a pair of perpendicular planar surfaces, not necessarily intersecting³. Then, the geometric constraints are inferred from: 1) the co-planarity of the observed line segments lying on each face of the corner, and 2) the perpendicularity of both planar surfaces. An overview of the calibration procedure is depicted in figure 2 showing its different stages, which are detailed below.

¹<https://github.com/EduFdez/mrpt/tree/LRF-calib/apps/LRF-calib/>

²https://github.com/EduFdez/mrpt/blob/LRF-calib/share/mrpt/datasets/3LRFs_dataset_demo.rawlog

³This procedure of calibration can make use of pairs of oblique (non-parallel) planes, not requiring perpendicularity. However, by considering only perpendicular corners we facilitate both the readability of this paper and its implementation and application since any structured scene contains perpendicular planes.




Figure 2: Calibration procedure.

2.1 Gathering observations

The first process is to collect data in an environment which contains physical corners. A minimum of two corner observations from different orientations are required to calibrate a pair of LRFs as it is described in section 4. In order to take such observations, either the sensor rig can be rotated with respect to the scene, or a calibration pattern made with 2 perpendicular planes can be rotated in front of the sensor rig. This last may be useful when it is difficult (or impossible) to move the sensor rig. In this paper we follow the first option to carry out our experiments, as we see it is the most convenient in most practical situations. It implies to segment lines from the raw data to generate corner candidates, which are evaluated later on according to some constraint equations to obtain the required corner observations (*COs*) as explained below.

We would like to remark that the rig can be calibrated from a single observation when three perpendicular planes are visible by the sensors, as it is shown in figure 4. However, it is preferable to take more observations to compensate for the noise in the measurements, thus increasing the accuracy of the calibration.

2.2 Line representation and segmentation

The planar structures of the environment are sampled by the LRFs as line segments. These lines can be extracted from the scans provided by each LRF in a number of ways (Nguyen et al., 2005). Here we have implemented a segmentation method based on RANSAC (Fischler and Bolles, 1981), although other approaches like those based on region growing (Borges and Aldon, 2000) or on Hough transform (Forsberg et al., 1995) may also be applied. The RANSAC procedure searches

for the parameters $\{A, B, C\}$ of a 2D line which maximize the number of points $\mathbf{p}_i = (x_i, y_i)^T$ supporting the model

$$|Ax_i + By_i + C| \leq \varepsilon \quad (1)$$

being ε a threshold used to differentiate between inliers and outliers. An advantage of using RANSAC is that unconnected collinear segments are automatically clustered as the same line, simplifying the subsequent optimization process.

The segmented lines are represented in 2D in the LRF's reference system by the normalized direction vector $l = (l_x, l_y)^T$ and an arbitrary point of the line (see figure 3). For such a point we have chosen the centroid of the line's inliers $c = (c_x, c_y)^T$ since it has less uncertainty to belong to the line than any measured point. These parameters and their covariances are estimated assuming unbiased, identically distributed (i.i.d.) Gaussian noise in the LRF measurements. In the literature, it is often assumed a model where the noise only affects the range measurements, with exact bearing directions (Arras and Siegwart, 1998; Diosi and Kleeman, 2003). However, such model introduces linearization errors that produces biased estimates of the line parameters. To avoid this, a common approach (see (Arras and Siegwart, 1998; Diosi and Kleeman, 2003)) that we follow here is to approximate the covariance of each point \mathbf{p}_i in Euclidean coordinates as $\Sigma_{p_i} = \sigma_i^2 I$. Then, the maximum likelihood estimation of the centroid \mathbf{c} is calculated as

$$\mathbf{c} = \frac{1}{N} \sum_{i=1}^N \mathbf{p}_i \quad (2)$$

and its covariance Σ_c is given by

$$\Sigma_c = \begin{bmatrix} \frac{\sigma^2}{N} & 0 \\ 0 & \frac{\sigma^2}{N} \end{bmatrix} \quad (3)$$

The line direction vector \mathbf{l} is obtained as the eigenvector corresponding to the largest eigenvalue of the point dispersion matrix M

$$M = \sum_{i=1}^N (\mathbf{p}_i - \mathbf{c})(\mathbf{p}_i - \mathbf{c})^T \quad (4)$$

and its covariance $\Sigma_l = H^+$ is calculated following (Pathak et al., 2010) as the Moore-Penrose generalized inverse of

$$H = \frac{1}{\sigma^2} \sum_{i=1}^N \begin{bmatrix} y_i^2 & -y_i x_i \\ -x_i y_i & x_i^2 \end{bmatrix} \quad (5)$$




Figure 3: A rig of 3 LRFs in different orientations (i.e. the z axis of each LRF are linearly independent).

In the rest of the paper the lines are represented in 3D in the LRF reference system (see figure 3) by setting the z component of the line parameters and their covariances to zero.

2.3 Corner constraints

Let S_1 and S_2 be two rigidly jointed LRFs, each one observing two line segments $\{L_1^a, L_1^b\}$ and $\{L_2^a, L_2^b\}$ from two perpendicular planes Π_a and Π_b . Let $[R_1|\mathbf{t}_1], [R_2|\mathbf{t}_2] \in \mathbb{SE}(3)$ be the LRFs poses with respect to a common reference frame, with the rotations $R \in \mathbb{SO}(3)$ represented as 3×3 matrices and the translations $\mathbf{t} \in \mathbb{R}^3$.

2.3.1 Co-planarity constraint

a co-planarity constraint is inferred for the lines segmented by two LRFs that observe the same plane, for instance, the lines $\{L_1^a, L_2^a\}$ on the plane Π_a in figure 1. Both line direction vectors and the vector joining any two points \mathbf{p}_1^a and \mathbf{p}_2^a of the observed lines, all referred to the same coordinate system, form a matrix of deficient rank as they are all parallel to the plane Π_a . Therefore, the determinant of that matrix, which is equal to the mixed triple product of the 3 vectors, must be zero. This is expressed as

$$(\mathbf{R}_1 \mathbf{l}_1^a \times \mathbf{R}_2 \mathbf{l}_2^a) \cdot (\mathbf{R}_1 \mathbf{p}_1^a + \mathbf{t}_1 - \mathbf{R}_2 \mathbf{p}_2^a - \mathbf{t}_2) = 0 \quad (6)$$

This condition can also be interpreted as a statement for the perpendicularity between the plane's normal vector

$\mathbf{n}^a = \mathbf{R}_1 \mathbf{l}_1^a \times \mathbf{R}_2 \mathbf{l}_2^a$ and any vector joining a pair of points from the scanned lines. For such points, we make use of the centroids \mathbf{c}_1^a and \mathbf{c}_2^a , simplifying the optimization procedure presented later on.

2.3.2 Orthogonality constraint

another constraint can be stated for the relative rotation of a pair of LRFs that observe a corner defined by the perpendicular planes Π_a and Π_b , so that

$$\mathbf{n}^a \cdot \mathbf{n}^b = (\mathbf{R}_1 \mathbf{l}_1^a \times \mathbf{R}_2 \mathbf{l}_2^a) \cdot (\mathbf{R}_1 \mathbf{l}_1^b \times \mathbf{R}_2 \mathbf{l}_2^b) = 0 \quad (7)$$

Notice that the first constraint affects only the observation of a plane, while the second one implies the observation of a corner. Regarding the calibration problem, the former involves the relative rotation and translation of the sensors, while the latter affects only the rotation.

2.4 From corner candidates to corner observations (COs)

Let us call *CO* the corner observed by a pair of rigidly jointed LRFs. Specifically, a *CO* is described by the set: $CO = \{j, L_j^a, L_j^b, k, L_k^a, L_k^b\}$, where j and k are the indices of the LRFs, a and b refer to the respective corner faces (orthogonal planes), and the lines are represented by $L_j^a = \{\mathbf{c}_j^a, \mathbf{l}_j^a, \Sigma_{c_j}^a, \Sigma_{l_j}^a\}$.

Detecting a *CO* is not trivial when the relative poses of the sensors are not known, as it is the case here. Indeed, since the information provided by the LRF measurements is purely geometric, we can only know that two pairs of lines observed by two LRFs come from a corner once we know the sensor's calibration. Assuming that corners are visible by the LRFs, the resulting problem implies detecting the *COs* and estimating the calibration simultaneously. This can be tackled in a hypothesize-and-test framework, where many corner candidates (potential *COs*) are generated from the rig observations by grouping sets of two pairs of lines seen by a pair of LRFs.

After a number of observations are taken from different poses of the sensor rig, inconsistent corner candidates are ruled out robustly using RANSAC (Fischler and Bolles, 1981) taking into account the restrictions in equations (6-7). For that, a candidate extrinsic calibration is calculated from a minimum set of randomly selected corner candidates (as explained in the next section). Then, the number of corner candidates consistent

with such calibration is evaluated. This process is repeated iteratively searching for the maximum consensus of corner candidates. The result of this process is the largest set of consistent candidates (inliers), which are the *COs* from which the calibration is computed. Empirically, we have verified that the correct calibration always corresponds to the largest number of inliers, even when the number of outliers (wrong corner candidates) is considerably larger than the number of inliers ($n_{inliers}/n_{outliers} \sim 0.1$). This situation where the number of outliers is much larger than that of inliers results in a slow calibration process, however, it is not critical since this can be done offline.

It is worth mentioning that if we have a rough knowledge of the sensors relative poses, which is very common in practice, some constraints can be set for the selection of *COs* so that less outliers are selected at a first instance. Such restrictions are not applied here for the sake of generality.

The process to obtain the *COs* can perform automatically from the streaming data of the sensors (see for example Extension 1 or the video at <http://youtu.be/v6Ls9NZWOZM>). Errors derived from the motion of the sensor rig are neglected here since the rig's velocity is small with respect to the acquisition time. In order to decide when should we stop gathering *COs*, a convergence condition for the maximum uncertainty (covariance) of the resulting calibration is set to stop this process automatically (Fernández-Moral et al., 2014). Finally, once the *COs* are obtained, the calibration is computed (last step in figure 2) as detailed in the following section.

3 PROBLEM FORMULATION

Let us consider a set of *COs* gathered with a rig of m LRFs $\{S_1, \dots, S_j, \dots, S_m\}$. Without loss of generality, the sensor S_1 is chosen as the reference coordinate frame, so that each LRF S_j is located with a relative transformation $[R_j | \mathbf{t}_j] \in \mathbb{SE}(3)$ with respect to S_1 . Then, we want to estimate the optimal $\{R, \mathbf{t}\} = \{[R_2 | \mathbf{t}_2], \dots, [R_j | \mathbf{t}_j], \dots, [R_m | \mathbf{t}_m]\}$ that minimize the errors of the constraints in eqs. 6-7, assuming independence between the *COs* and that the measurements are affected by unbiased Gaussian noise as modelled in section 2.2.

The above problem is formulated as the maximum likelihood estimation (MLE) of the relative poses $\{R, \mathbf{t}\}$ for the given *COs*, which is calculated from the maxi-

mization of the log-likelihood

$$\operatorname{argmax}_{\{\mathbf{R}, \mathbf{t}\}} \left(\ln \prod_{i=1}^N p(CO_i | \{\mathbf{R}, \mathbf{t}\}) \right) \quad (8)$$

where the likelihood of $\{\mathbf{R}, \mathbf{t}\}$ for the i -th CO is calculated from the constraints presented above (eqs. 6-7). It is expressed as the multiplication of the likelihood for each constraint, two co-planarity constraints from eq. 6 (one per plane in the corner) that affect the estimate of both rotation and translation and one constraint from eq. 7 that affects only the estimate of the relative rotation. Thus, for a given CO_i observed by the LRFs j and k , we have:

$$p(CO_i | \{\mathbf{R}, \mathbf{t}\}) = p(CO_i^a | \mathbf{R}_j, \mathbf{t}_j, \mathbf{R}_k, \mathbf{t}_k) \cdot p(CO_i^b | \mathbf{R}_j, \mathbf{t}_j, \mathbf{R}_k, \mathbf{t}_k) \cdot p(CO_i^{ab} | \mathbf{R}_j, \mathbf{R}_k)$$

with the superindices a and b referring to the co-planarity constraints inferred from each plane, and the superindex ab referring to the perpendicularity condition. As verified experimentally through Monte Carlo simulations in appendix A, the above probabilities follow Gaussian distributions. Concretely, the first two elements of the right term in eq. 9 are given by

$$p(CO_i^a | \mathbf{R}_j, \mathbf{t}_j, \mathbf{R}_k, \mathbf{t}_k) = \frac{1}{\sqrt{2\pi}\sigma_i^a} \exp\left(-\frac{1}{2} \frac{e_i^2}{(\sigma_i^a)^2}\right) = \frac{1}{\sqrt{2\pi}\sigma_i^a} \exp\left(-\frac{(\mathbf{n}_{jk}^a \cdot (\mathbf{R}_j \mathbf{c}_j^a + \mathbf{t}_j - \mathbf{R}_k \mathbf{c}_k^a - \mathbf{t}_k))^2}{2(\sigma_i^a)^2}\right) \quad (10)$$

with

$$\mathbf{n}_{jk}^a = \mathbf{R}_j \mathbf{l}_j^a \times \mathbf{R}_k \mathbf{l}_k^a \quad (11)$$

being σ_i^a the standard deviation of the error e_i (eq. 6). The expression of the probability from the co-planarity constraint of the plane Π_b is the same with the exception of the superindices. On the other hand, the probability inferred from eq. 7 is given by

$$p(CO_i^{ab} | \mathbf{R}_j, \mathbf{R}_k) = \frac{1}{\sqrt{2\pi}\sigma_i^{ab}} \exp\left(-\frac{(\mathbf{n}_{jk}^a \cdot \mathbf{n}_{jk}^b)^2}{2(\sigma_i^{ab})^2}\right) \quad (12)$$

being σ_i^{ab} the standard deviation of the error of (eq. 7). Both standard deviations (σ_i^a and σ_i^{ab}) are computed through linearisation from a first order Taylor approximation of the error functions. Their derivation is detailed in the appendix B.

When these standard deviations are constant with respect to the model parameters, the solution of the MLE in (8) coincides with that of the weighted, non-linear least squares problem expressed as

$$\operatorname{argmin}_{\{\mathbf{R}, \mathbf{t}\}} \sum_{i=1}^N \left(\omega_i^a (\mathbf{n}_{jk}^a \cdot (\mathbf{R}_j \mathbf{c}_j^a + \mathbf{t}_j - \mathbf{R}_k \mathbf{c}_k^a - \mathbf{t}_k))^2 + \omega_i^b (\mathbf{n}_{jk}^b \cdot (\mathbf{R}_j \mathbf{c}_j^b + \mathbf{t}_j - \mathbf{R}_k \mathbf{c}_k^b - \mathbf{t}_k))^2 + \omega_i^{ab} (\mathbf{n}_{jk}^a \cdot \mathbf{n}_{jk}^b)^2 \right) \quad (13)$$

where ω_i^x (the superindex x stands for a , b or ab) is the weight of the corresponding residual from CO_i

$$\omega_i^x = \frac{1}{(\sigma_i^x)^2} \quad (14)$$

This problem is reformulated using Lie algebra (Fernández-Madrigo and Claraco, 2013) to represent the poses with a minimal parameterization on a manifold. For that, the rotations are represented as the composition of a guessed rotation and a rotation increment represented with the exponential map ($e^{\mu_j} \mathbf{R}_j$), with the rotation increment $e^{\mu_j} \in \mathbb{SO}(3)$. The translations are also represented as the sum of a guessed translation plus an increment ($\mathbf{t}_j + \Delta \mathbf{t}_j$), both in \mathbb{R}^3 . The resulting non-linear least squares problem is solved iteratively using Levenberg-Marquardt

$$[\mu_2^k \Delta \mathbf{t}_2^k, \dots, \mu_m^k \Delta \mathbf{t}_m^k]^T = -(H + \lambda \operatorname{diag}(H))^{-1} g \quad (15)$$

being λ the Levenberg-Marquardt's damping factor. H is the Hessian (a symmetric matrix of dimension $6(m-1)$) and g is the Gradient (a column vector of dimension $6(m-1)$) of the cost function, which are calculated as

$$H = \sum_{i=1}^N J_i^T \omega_i J_i, \quad g = \sum_{i=1}^N J_i^T \omega_i r_i \quad (16)$$

where N is the number of constraints of this optimization, being r_i the residual and J_i the Jacobian for each constraint (remember that each CO provides three constraints). The Jacobian J_i^a corresponding to the constraint from eq. 6 is calculated as

$$J_i^a = [\dots; \underbrace{(\mathbf{l}_j^{aW} \times (\mathbf{l}_k^{aW} \times (\mathbf{c}_j^{aW} - \mathbf{c}_k^{aW})) + \mathbf{R}_j \mathbf{c}_j^a \times \mathbf{n}_{jk}^a)}_{J_{\mu_j}}; \underbrace{\mathbf{n}_{jk}^a}_{J_{\Delta_j}}; \dots; \dots; \underbrace{-(\mathbf{l}_k^{aW} \times (\mathbf{l}_j^{aW} \times (\mathbf{c}_j^{aW} - \mathbf{c}_k^{aW})) - \mathbf{R}_k \mathbf{c}_k^a \times \mathbf{n}_{jk}^a)}_{J_{\mu_k}}; \underbrace{-\mathbf{n}_{jk}^a}_{J_{\Delta_k}}; \dots]^T \quad (17)$$

where the superindex W refers to the common reference frame, so that⁴

$$l_j^{aW} = \mathbf{R}_j \mathbf{l}_j^a, l_k^{aW} = \mathbf{R}_k \mathbf{l}_k^a$$

$$\mathbf{c}_j^{aW} = \mathbf{R}_j \mathbf{c}_j^a + \mathbf{t}_j, \mathbf{c}_k^{aW} = \mathbf{R}_k \mathbf{c}_k^a + \mathbf{t}_k$$

This Jacobian (a row vector of dimension $6(m-1)$) contains four blocks of 1×3 vectors corresponding to the derivatives of the residual with respect to μ_j , $\Delta \mathbf{t}_j$ and μ_k , $\Delta \mathbf{t}_k$, respectively.

On the other hand, the Jacobian J_i^{ab} of the residual from eq. 7 is given by

$$J_i^{ab} = [\dots; \underbrace{\mathbf{l}_j^{bW} \times (\mathbf{l}_k^{bW} \times \mathbf{n}_{jk}^a) + \mathbf{l}_j^{aW} \times (\mathbf{l}_k^{aW} \times \mathbf{n}_{jk}^b)}_{J_{\mu_j}}; \dots]$$

$$\dots; \underbrace{-\mathbf{l}_k^{bW} \times (\mathbf{l}_j^{bW} \times \mathbf{n}_{jk}^a) - \mathbf{l}_k^{aW} \times (\mathbf{l}_j^{aW} \times \mathbf{n}_{jk}^b)}_{J_{\mu_k}}; \dots]^T \quad (18)$$

The two blocks of 1×3 vectors of this Jacobian correspond to the derivatives of the residual with respect to μ_j and μ_k respectively, being zero the blocks corresponding to the rest of elements in $\{\mathbf{R}, \mathbf{t}\}$.

This optimization is solved iteratively

$$\mathbf{R}_j^{k+1} = e^{\mu_j^k} \mathbf{R}_j^k, \mathbf{t}_j^{k+1} = \Delta \mathbf{t}_j^k + \mathbf{t}_j^k, j \in [2, m] \quad (19)$$

from an initial guess for the sensor relative poses, which may be obtained from a rough measurement of the rig. Once the problem is solved, the covariance of the resulting calibration is calculated as the inverse of the Hessian of the cost function in (13) (Fernández-Madriral and Claraco, 2013).

4 OBSERVABILITY

The problem of estimating the relative poses has $6(m-1)$ DoF, with m being the number of LRFs. From the formulation presented in the previous section, we have seen that each CO leads to three constraints for the relative poses of the corresponding pair of LRFs. Therefore, at least $2(m-1)$ CO s are needed to provide as many equations as unknowns for solving the problem.

We are also interested in knowing how these observations should be taken in order to provide the necessary information to solve the calibration. The analysis of the

⁴For clarity, the CO index i will be omitted in subsequent operations that affect only to the same CO .




Figure 4: Observation of a corner with three perpendicular planes by two LRFs.

observability of calibration problems provides valuable information about the procedure to gather such data (Martinelli, 2011; Censi et al., 2013). Such analysis is carried out here by studying the rank of the Fisher Information Matrix (FIM) of the estimation problem (see appendix C). The key concept here is that when the FIM is singular, the information carried by the data (observations) is not sufficient and the problem is under-constrained ($\{\mathbf{R}|\mathbf{t}\}$ is unobservable). We wish to identify these situations of unobservability in order to avoid them in practice.

From the Monte Carlo simulation carried out in 8.1, the probabilities in eqs. 10 and 12 can be assumed to follow Gaussian distributions. Thus, following appendix C, the FIM can be expressed as

$$FIM = J^T \Omega J \quad (20)$$

where J is a matrix concatenating the Jacobians J_i of the different constraints, and Ω is a diagonal matrix containing the weights ω_i of such constraints. Since Ω is diagonal with all the elements being positive, the rank of FIM is given by the rank of J

$$\text{rank}(FIM) = \text{rank}(J^T J) = \text{rank}(J) \quad (21)$$

Therefore, the problem has a solution when

$$\text{rank}(J) = 6(m-1) \quad (22)$$

By analysing the structure of the different Jacobians J_i for each CO , we can notice that a CO provides three

linearly independent rows for J when a corner is observed in a new orientation (eqs. 17 and 18). To get a deeper insight into this, consider a block of the Jacobian in eq. 18, for instance J_{μ_j} . Each corner observation in a new, linearly independent direction, expressed by $\mathbf{n}_{jk}^a \times \mathbf{n}_{jk}^b$, contributes to constrain the problem for μ_j . For the Jacobian in eq. 17, it can be verified that each plane observation providing a linearly independent \mathbf{n}_{jk} results in a linearly independent J_i which constrains both the relative rotation and the relative translation between the sensors S_j and S_k . Therefore, two observations of a corner from different orientations suffice to solve the calibration of a pair of sensors. Moreover, a single observation of a corner with three perpendicular planes (as shown in figure 4), provides enough information to solve the problem since it contains already 3 independent normal vectors of the plane, and 3 independent orthogonal constraints.

An interesting case of unobservability occurs for planar movement of a sensor rig, when all the visible planes are perpendicular to the plane of movement. In such case, it can be clearly seen that there is a free degree of freedom for the translation since the rank of the matrix concatenating the plane's normal vectors will always be deficient. This situation arises for a vehicle with an horizontal LRF which only observes vertical planes. In order to calibrate such system, the scene should contain oblique planes, or the rig should be tilted in order to take observations from non-vertical planes.

Finally, the ratio $\eta = \mu_{6(m-1)}/\mu_1$ between the smallest and largest eigenvalues of the FIM is also an indicator of how well distributed the measurements are along the different directions (DoF) of the domain ($\eta = \text{condition_number}^{-1}$). So, in the best case $\eta = 1$ which means that all plane observations are equally distributed in the space, while when $\eta \rightarrow 0$, the system becomes ill-conditioned.

5 CONVERGENCE

Considering that the calibration is observable, another important issue is to know if the solution converges to the correct value. This problem is not trivial for a non-linear optimization whose domain is not convex, containing local minima, as it is the case here. The local convexity of the error function around the solution depends on a number of parameters including: the configuration of the sensor rig, the amount of corner observations and their positions, and the noise in the sensor

measurements. Thus, a mathematical condition for the convergence cannot be established in general.

However, given a configuration of the rig and a set of observations (COs), we can sample the optimization domain to conduct a qualitative analysis of its convexity. In figure 5 we display the residual error of eq. (13) for a rig with two LRFs which observe 20 COs from different orientations. The error is shown with respect to the six parameters of the calibration by grouping pairs of DoFs in rotation and translation. The first row of figure 5 shows three sections of the sphere of possible rotations⁵, corresponding to the planes $x - y$, $y - z$ and $z - x$, while the second row shows the translation domain, where each DoF takes values in $[0.5, -0.5]$ meters around their true value which is at the center $\{0,0\}$ of each graph. The resulting residuals are shown with a 2D heat map with the contour lines. In all the graphs, we see clearly a minimum at the correct calibration. We observe that there are local minima in the orientation domain, what implies that the initial values for the relative rotation must be given in a local region around the solution. On the other hand, the problem is convex for the translation as it is inferred from the formulation (the error depends linearly on the translation), therefore, the result does not depend on its initialization.

To give a further insight on this topic, we show the percentage of convergence of the test above with respect to the initial rotation offset measured in degrees. A test is said to converge when the rotation error with respect to the groundtruth is lower than 1° and the translation error is below 1 cm. In this experiment, the initial rotation is given by $\mathbf{R}_{Init} = \mathbf{R}_{Offset} \mathbf{R}_{groundtruth}$ with $\mathbf{R}_{Init}, \mathbf{R}_{Offset}, \mathbf{R}_{groundtruth} \in \mathbb{SO}(3)$. The domain of all possible initializations for the rotation is represented by a sphere $\tau \in \mathfrak{so}(3)$ where $\mathbf{R}_{Offset} = e^\tau$ with radius $r = \text{norm}(\tau)$, $r \in [0, \pi]$. With this experiment we check the "convergence rings" around the true solution. Such convergence rings represent spherical surfaces (fixed radius) in \mathbb{S}^2 . A discretization with 1 degree sampling is employed for $[\theta, \phi] \in \mathbb{S}^2$ in each spherical surface. The results of this experiment are shown in figure 7, where each bar represents the rate of convergence in a surface of possible rotations of a given angular offset. In analogy to the graph of figure 5, where the contour maps represent planar sections of the sphere of possible rotations, here the bars represent different spherical surfaces (different modules of the rotation offset). As we

⁵Note that the domain for the rotation in the Lie algebra $\mathfrak{so}(3)$ is a sphere of radius $r = \pi$

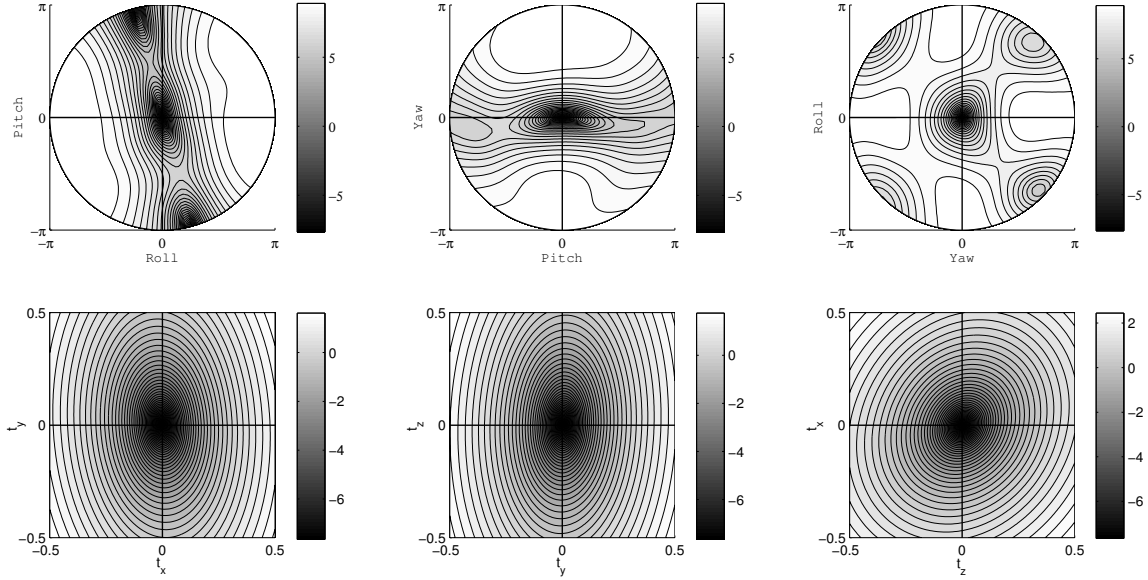


Figure 5: Error maps of the calibration of a pair of LRFs for 6 different combinations of 2 DoFs using 100 COs. The configuration of the rig corresponds to the one depicted in fig. 11 for the sensors S_1 and S_2 . The heat maps show the residual error from eq. 13 in logarithmic scale, with the correct calibration at the center of each graph, lying on a local minimum. The free DoFs in rotation take all possible values in the domain $\tau \in [-\pi, \pi]$, while for the translation the DOFs are given values $\delta \in [-0.5m, 0.5m]$.

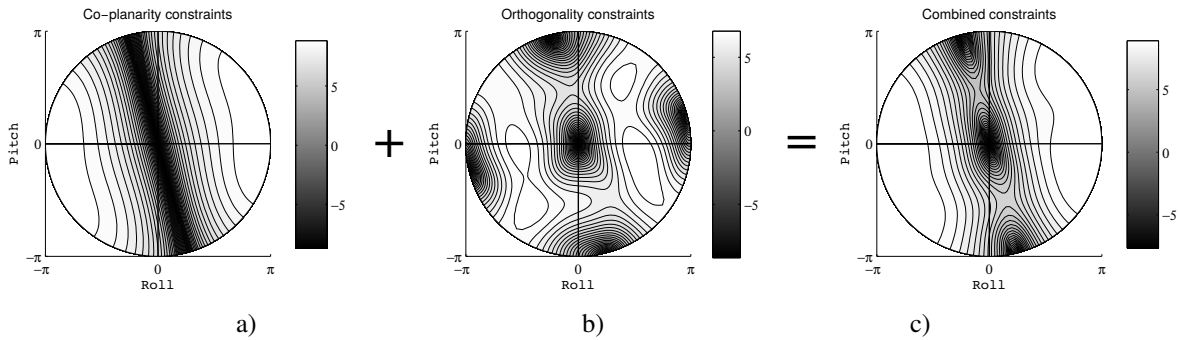


Figure 6: Error maps for different cost functions in logarithmic scale: a) co-planarity constraints, b) orthogonality constraints, c) the combination of these two. These graphs correspond to the previous simulation for two DoFs in the rotation domain (the graph c is the same as the graph shown at the top-left of figure 5). Note that the correct calibration at $(0,0)$ lies on a local minimum only when the orthogonality constraints are applied.

can see, the further the initializations are from the true solution, the more reduced the convergence rate is.

The tests above have been repeated for a number of rig configurations and for different sets of COs, where the results are qualitatively similar. For all such tests, we observe a similar trend for the error surfaces and

the convergence rates, indicating that there is at least one local minimum for the error function at the correct solution, and that there are local and global minima distributed in the domain. A particularly interesting case of wrong convergence occurs when the relative rotation of a pair of LRFs is initialized in a way that their scanning




Figure 7: Calibration convergence rate of the sensors S_1 and S_2 with respect to the angular offset of the initial rotation. Please, see explanation in the text.

planes coincide. Note that there exists a global minimum for such set of parameters, where the error will be zero no matter the COs . Initializing the calibration near such global minima can drive the optimization to wrong calibration.

Another interesting point is to know if the calibration can be performed only from plane observations since, in fact, co-planarity constraints already restrict both the rotation and the translation of the relative poses in the rig. Several tests have indicated that this form of calibration is not possible because the cost function is not locally convex near the solution, and thus the problem cannot be properly constrained. As shown in figure 6, the introduction of the orthogonality constraint (eq. 7) makes the correct solution lie on a local minimum. Figure 6 shows the convergence error maps of different cost functions from: a) co-planarity constraints, b) orthogonality constraints, and c) a combination of both, which is actually the sum of the previous two. Note how these two functions complement well making the correct calibration lie on a local minimum.

In summary, the region of convergence in which we can provide a valid initial rotation depends on the particular configuration of the sensor rig. Therefore, a closed domain to guarantee convergence cannot be defined. From an engineering point of view, we observed that such convergence region is wide enough to provide good initial values for the relative poses from simple visual inspection of the rig.

6 EXPERIMENTS

A number of experiments has been carried out to validate the present approach from both simulated and real

LRF rigs.

6.1 Simulation

In our simulation environment, a rig consisting of two non-parallel LRFs is placed at different distances and orientations with respect to a corner in order to gather measurements from several poses. The sensors are modelled according to the parameters of the Hokuyo UTM-30LX rangefinder, and the observations are generated with unbiased, uncorrelated Gaussian noise with $\sigma = 0.03$ m. The line features and their covariances are extracted from these synthetic observations. The calibration is estimated for the cases of weighted (MLE) and unweighted optimization (standard least squares) for a varying number of COs . The average errors of the calibration with respect to the true poses are obtained from a Monte Carlo simulation with 10^5 trials for every set of COs . For each test, the initial relative pose is uniformly generated around the groundtruth at distance $d \in [0, 1m]$ and at an angle $|\tau| \in [0, \pi/4]$. The average errors of the relative rotation and translation are shown in figure 8 in degrees and millimeters, respectively. We observe that these errors diminish asymptotically with the number of COs . Also, we see how the MLE solution that takes into account the covariance of the measurements is consistently more accurate than the solution which ignores that information. This test was repeated for several configurations of the LRF rig (different relative poses between the sensors) obtaining similar results.

We also study the bias and covariance of our method from the above Monte Carlo simulation by analysing the distribution of the calibration results. The six dimensional errors of the calibrated poses are shown in figure 9, by grouping pairs of DoF for the rotation and the translation, respectively. This figure shows the distribution of the 10^5 samples around the groundtruth (at the intersection of the two lines), the 2σ confidence ellipses of Monte Carlo, and the one corresponding to the estimated covariance of one sample through the Cramér-Rao Bound (see appendix C). Note that the bias (difference between the center of the Monte Carlo ellipses and the groundtruth) and the covariance of the solution are very small as they are in the order of millimetres and 10^{-3} radians for the translation and rotation respectively. Also, the CRB covariance, which is the best covariance we may obtain from the sensor measurements when all the assumptions are correct, is very close to the true covariance given by Monte Carlo,




Figure 8: Calibration error of 2 LRFs with and without covariance weighting, represented by the norms of the rotation (left) and translation (right) error vectors.

which implies that our estimator is highly efficient.

Finally, we examine the problem of calibrating several LRFs, where the corner observations relate pairs of sensors. This problem is solved as the optimization of a graph of constraints, similar to the problem of pose graph SLAM (Grisetti et al., 2010). To give some insight on this, we simulate the observation of a number of corners by a rig with 4 LRFs. The relative poses of these were randomly generated around the reference sensor *LRF1* which is located at the reference coordinates. Figure 10.a) shows the average rotation errors of the lasers *LRF2*, *LRF3* and *LRF4* with respect to the number of corner observations. We compare these errors with those obtained from the calibration of the different pairs (1-2, 1-3 and 1-4) shown in fig. 10.b). As expected, using all the information available (full graph optimization) results in less error than pair-wise calibration.

6.2 Real data

We have also validated the proposed calibration method in real case scenarios employing: 1) a rig with three LRFs and 2) the sensors mounted on two autonomous cars. The characteristics of the calibrated LRFs are shown in table 1.

Test rig

In the first case, the test rig is composed of 3 Hokuyo UTM-30LX (see figure 11). The LRF recordings of

Table 1: Characteristics of the LRFs calibrated in this paper.

	Hokuyo UTM-30LX	Sick LMS 291-S05
Range (m)	[0, 60]	[0, 80]
σ (m)	0.03	0.01
Resolution	0.25°	0.25°
Field of view	270°	180°



Figure 11: Test LRF rig with three Hokuyo UTM-30LX.

this experiment are publicly available at www.mrpt.org/LRF-calib_dataset. The sensors' synchronization effect is neglected in this test since the rig is smoothly waved at a low velocity while the LRFs scan at a frame rate of 40 Hz (see Extension 1 or the video at <http://youtu.be/v6Ls9NZWOZM>). The accuracy of the resulting calibration cannot be estimated di-

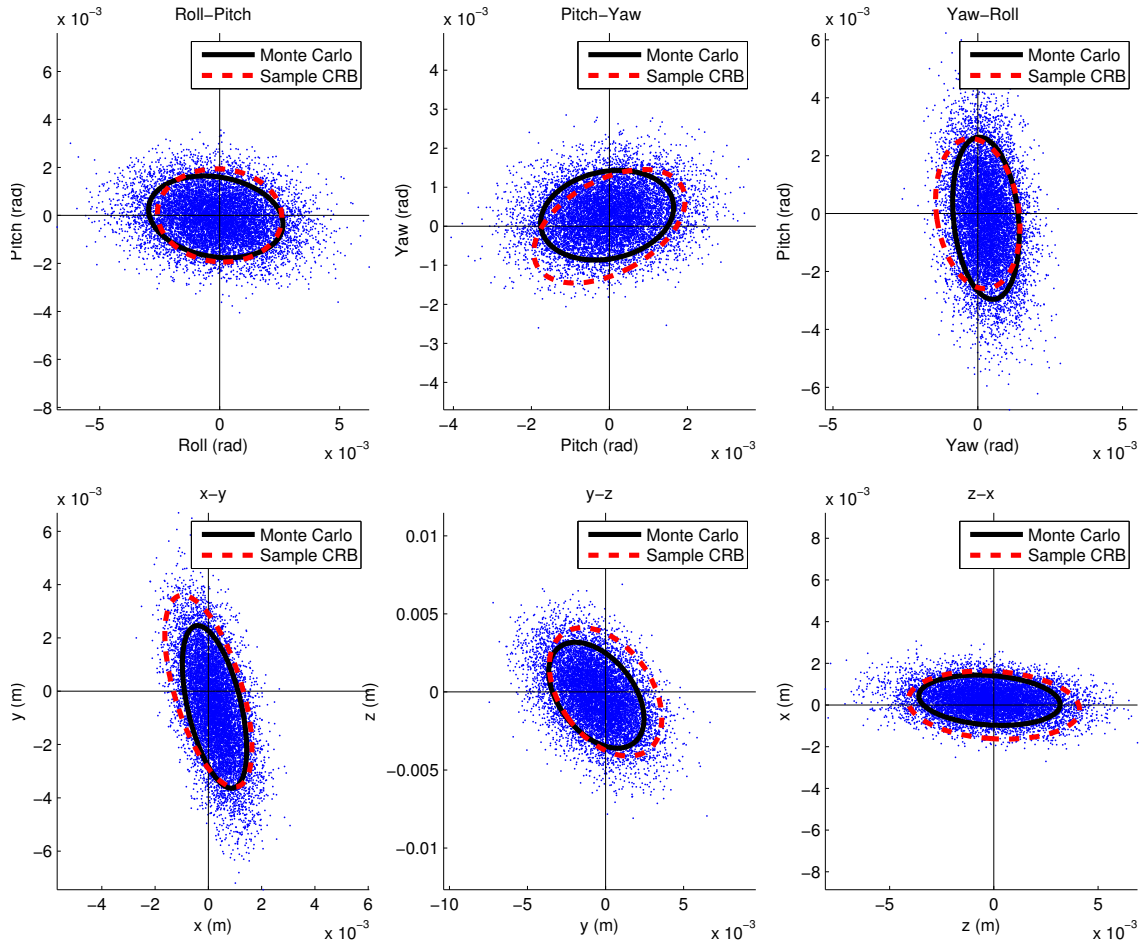


Figure 9: Error distribution and 2σ covariance ellipses of the calibration of two LRFs from a Monte Carlo simulation with 10^5 samples. The 2σ covariance ellipses from the Cramér-Rao Bound (CRB) of one of these samples is also drawn, centred at the groundtruth.

rectly since a groundtruth for the sensors relative poses is not available. Instead, we evaluate the accuracy of the method by checking that the pose composition from calibrating the different pairs closes a loop ($R_{12}R_{23}R_{31} = I$ and $\mathbf{t}_{12} + \mathbf{t}_{23} + \mathbf{t}_{31} = \mathbf{0}$).

In this test we validate our approach using a varying amount of *COs*, which are obtained at different orientations of the rig while it moves as it is shown in the video. The initial poses required by our method are given around a guess obtained from visual check of the rig, in a range of $[0, 40]$ degrees for the rotation and $[0, 1]$ meters for the translation, with respect to the correct calibration (several initializations are tested to check the robustness of our method). Table 2 shows the

results of this test for different numbers of *COs*, from a minimum of 2 *COs* to 100 *COs*. The first three columns show the average residuals of the calibration of each pair of sensors, and the last two columns show the average deviation with respect to the loop closure condition of the three independent calibrations. From this table, we observe that, as expected, the average residuals and the loop closure deviations decrease with the number of *COs*.

In this experiment, we have also calibrated the three LRFs by optimizing the full graph of constraints between them, so that the above loop closure condition is guaranteed. This way, the calibration should be more accurate since it uses all the information available. Ta-

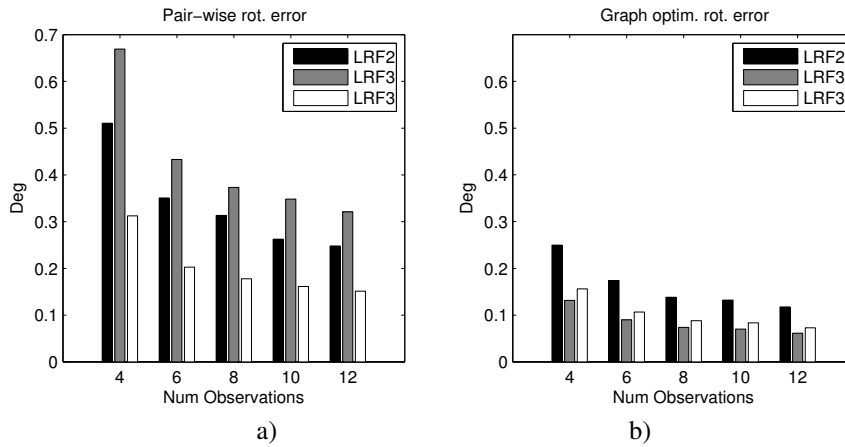


Figure 10: Errors obtained from the calibration of 4 LRFs for a different amount of COs .

Table 2: Average residual and total deviation of the calibration for a varying amount of COs

COs	res_{12}	res_{23}	res_{31}	R dev (deg)	t dev (cm)
2	2.74	3.81	1.41	1.03	5.31
20	1.48	1.70	1.25	0.63	1.72
40	1.46	1.66	1.23	0.51	0.54
60	1.39	1.66	1.22	0.49	0.34
80	1.33	1.62	1.22	0.48	0.29
100	1.32	1.62	1.21	0.47	0.27

Table 3: Deviations between global calibration and the calibration of each pair for a varying amount of COs

COs	Rotation Dev. (deg)			Translation Dev. (cm)		
	r_{12}	r_{23}	r_{31}	t_{12}	t_{23}	t_{31}
2	0.84	0.54	0.65	1.21	1.03	0.83
20	0.41	0.40	0.52	1.10	0.93	0.72
40	0.40	0.36	0.50	0.96	0.91	0.74
60	0.39	0.35	0.43	0.89	0.79	0.65
80	0.33	0.33	0.29	0.88	0.68	0.66
100	0.32	0.34	0.27	0.73	0.67	0.61

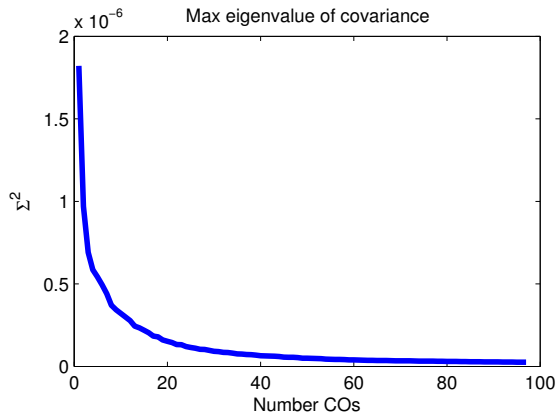


Figure 12: Maximum eigenvalue of the calibration covariance with respect to the number of COs .

Table 3 shows the deviation of the relative pose between each pair of sensors and this global calibration.

The covariance of the resulting calibration depends on the information provided by the COs . In general,

providing more COs contributes to reduce the uncertainty of the solution. This is confirmed in figure 12, which displays the maximum eigenvalue of the calibration covariance with respect to the number of COs for the experiment above. We observe how such value of the variance decreases asymptotically with the number of COs . This feature is relevant since it allows the user to set the maximum uncertainty for the calibration, so that the process of gathering COs stops after such limit is reached.

Autonomous car datasets


We have also validated our method by calibrating the sensors mounted on two different autonomous vehicles, using two publicly available datasets^{6,7}. For the dataset in (Blanco-Claraco et al., 2014), the vehicle has five

⁶<http://www.mrpt.org/MalagaUrbanDataset>

⁷<http://grandchallenge.mit.edu/wiki/index.php?title=PublicData>



(a)



(b)

Figure 13: A semi-autonomous car which incorporates several laser rangefinders at different orientations. b) Scheme of the sensors.

LRFs in total, three Hokuyo UTM-30LX and two Sick LMS 291-S05, whose configuration is shown in figure 13. The Sick sensors scan horizontal planes, and therefore, the calibration cannot be fully constrained unless they observe non-vertical planes (for that, either the rig must be tilted or the scene should contain oblique planes like shop awnings, otherwise there is only one linearly independent *CO* visible by the Sick LRFs). This situation does not occur in the dataset, so, only the Hokuyo sensors are considered. Note that two of these three sensors (labeled as Hokuyo2 and Hokuyo3) scan the same vertical plane, however they can still be calibrated since the sensor Hokuyo1 has a different orientation. For calibration, we have chosen an extract of the dataset⁸ where

⁸<http://www.youtube.com/watch?v=qZMlc5UeUpE>




Figure 14: Snapshot of the corner observed by the lasers 2 and 4 of the MIT's autonomous vehicle Talos during the 2007 Darpa Urban Challenge.

the car travels through some streets with buildings at the sides. The corner observations were selected in a supervised way because the clutter in the scene (other cars, trees, etc.) introduces a huge amount of wrong correspondences that prevents a correct corner detection. Our method could be applied however automatically when there is less clutter, like in the previous experiment.

The calibration was computed from 12 *CO*s taking the extrinsic parameters provided with the dataset for the initialization which, according to the authors, were manually measured from the rig. All the corners selected for this calibration come from the floor and a wall, which are assumed to be perpendicular. The average angle between these planes for the 12 corners after calibration was 89.4 degrees, with a standard deviation of 0.68, while, by using the calibration provided with the dataset the average angle was 85.9 degrees, with a standard deviation of 5.6. Also, the visualization of the calibrated laser scans for both cases shows that our estimation is clearly more accurate since the intersection of the scanning planes produces coincident points, and the straight segments observed by each laser lie on the scene planes. On the contrary, such visual alignment is not so good for the calibration proposed in the dataset.

The second dataset used here (Huang et al., 2010) corresponds to the recordings of a car which participated in the Darpa Challenge, which has 13 LRFs. The dataset was taken in open outdoor spaces which are scarce in corner structures. At some point of the video sequence however, the car passes near a building where a pair of LRFs observe two corners (see figure 14). Such corners are seen at the time 15m03s of the dataset "2007-11-03-log-uce-scrubbed.mission1", by the lasers

2 and 4. From this single observation we perform calibration using different initializations for their relative pose. Such initializations were randomly generated around the calibration proposed in the dataset, in a similar way to the previous experiments. In this case, the lines were segmented using a parameter independent line fitting method (Prasad et al., 2011), which works better given the low angular resolutions of the sensors in the dataset (1/4 of the maximum resolution). The estimated pose differs from the one reported by the authors by $\sim 0.8^\circ$ in the rotation and ~ 20 cm in the translation. It is hard to say which calibration is more accurate in terms of the rotation, while for the case of the translation, the calibration provided with the dataset looks more accurate according to the visualization of the reconstructed laser scans. We attribute this difference to the COs obtained (only 2) and to the curvature of one of the observed "planes" (the floor). Also, the observations are not simultaneous since the scans were taken with no synchronization while the car moves at a considerable speed, contributing to increase the calibration error.

There are several ways to improve the calibration obtained for this practical example. The first thing would be to take observations in a more structured scenario, like in a city, or just in front of a wall (a wall with the floor constitute a corner). In this way, a bigger number of observations could be taken to compensate for different sources of error. Finally, taking synchronized (or still) observations will also help to obtain more accurate results.

Table 4: MIT Dataset calibration from a single observation with different initializations.

Av Rot deviation	0.79°
Rot precision	0.02°
Av Trans deviation	23.2 cm
Trans precision	3.65 cm

7 DISCUSSION

This section presents to the reader some practical information about how our method can be applied. The most convenient case is when we can move freely the LRF rig, so that we can gather the data in a scenario with large, clearly visible corner structures. This case is shown in our experiments with the test rig. In some

cases we might need to tilt manually the robot where the rig is already mounted on, as it is shown in the accompanying video demonstration. This procedure is common when the scanning plane of one sensor is parallel to the floor. Whereas this may be less practical for large vehicles, it could be very useful for robotic wheelchairs or small service robots, which constitute a growing class of applications which require precise sensor calibration for navigating in real-world environments.

For the case when moving the rig is not practical, a good option may be building a corner pattern which can be waved manually in front of the sensors to provide them the corner observations. This may also be the option to calibrate an static rig (impossible to move), unless it observes already a corner with three perpendicular planes like the one shown in figure 4, in which case the rig can be calibrated without the need to build a pattern. Depending on the sensor setup (dimensions and weight of the system), and on the particular environment, gathering the measurements and/or producing a specific calibration pattern (with two perpendicular planes) may be more or less inconvenient. But we reckon that in most scenarios, this is still more practical than taking manual measurements on the sensor set-up or engineering a specific *ad-hoc* solution.

It is also worth to mention that the only actual case where our method cannot be applied is when the measurement planes of all the sensors are parallel. This means that two parallel sensors can still be calibrated if there is a third sensor scanning a plane which is not parallel to them. This case is common for experimental autonomous cars, which besides having some parallel LRFs have also other ones in different non-parallel orientations, what allow to calibrate the system without need of an extra LRF with the only purpose of calibration.

Regarding the effort needed for calibration, we observe that in most cases the biggest effort comes from taking the rig to an amenable place (in many cases the current working environment is fine) and taking a few seconds of measurements to take enough corner observations. The most relevant time cost of applying our method comes from the fact of taking the different measurements (the computation time is in the order of milliseconds for about 100 COs). The optimal amount of COs may be chosen by the user according to the accuracy he requires, the particular noise model of his sensors, and the facility to take observations of corners in different viewing orientations. As a rule of thumb, we

have checked that some tens of *COs* provide a solution very close to the maximum accuracy and may suffice in practice.

8 CONCLUSIONS

We have presented, to the best of our knowledge, the first general solution to calibrate the extrinsic parameters of a rig of 2D range scanners. The method relies on the observation of perpendicular planes to constraint the relative poses of the different LRFs. This problem is solved in a probabilistic framework that takes into account the uncertainty in the measurements of the sensors, and as a result, it also provides the uncertainty of the estimated calibration. The observability and the convergence conditions for the problem are studied, showing that there exists a minimal solution which only requires a single observation from the LRF rig.

The calibration method proposed here presents important advantages with respect to previous approaches, since it is applicable to almost any sensor configuration, it is easy to use and easy to automatize, while being robust and accurate. Also, its probabilistic formulation allows to calibrate different models of sensors, as each error is weighted according to its uncertainty. We have conducted several experiments to validate our approach, both with synthetic and real data, which have demonstrated the claimed features of our proposal.

APPENDIX

8.1 Gaussian error model

It is common to assume Gaussian error models in robotics. This section verifies such assumption for the error models employed here. In the simulation experiments carried out in this work, the depth measurements are generated with additive Gaussian noise of zero mean, and no error in the bearing, although, the covariances of the measured points are assigned as diagonal matrices (eq. 3) in Euclidean coordinates (a common practice in most of the related works (Arras and Siegwart, 1998; Diosi and Kleeman, 2003)). Moreover, several operations are applied to Gaussian random variables in the error model, so that the resulting distribution is not necessarily Gaussian. In this appendix section, we conduct some Monte Carlo simulation with 10^5 samples, to check how the planarity error is distributed

(eq. 6) and the perpendicularity error (eq. 7) of a corner observation for different normal errors in the measurements. These errors are simulated for three different corners at different orientations, and for three rig configurations with two LRFs (nine evaluations are shown). The results of these simulations are shown in figure 15 together with a χ^2 goodness of fit test, confirming that both errors follow normal distributions of zero mean. This test has been repeated for different resolutions of the LRFs, obtaining similar results.

8.2 Propagation of uncertainty

The solution presented in this paper propagates the uncertainty of the sensor measurements to the resulting calibration. For that, the covariance of different functions must be estimated. This is systematically done through linearization, so that the covariance Σ_y of

$$y = f(x) \quad (23)$$

is computed as

$$\Sigma_y \approx J_x \Sigma_x J_x^T \quad (24)$$

being $J_x = \frac{\partial f(x)}{\partial x}$ the Jacobian of $f(x)$. This is an approximation when $f(x)$ is not linear, otherwise the formula is exact.

In order to compute the variance for the error function in eq. 10

$$r_{jk} = \underbrace{R_j \mathbf{l}_j \times R_k \mathbf{l}_k}_{n_{jk}} \cdot \underbrace{(R_j \mathbf{c}_j + \mathbf{t}_j - \hat{R}_k \mathbf{c}_k - \mathbf{t}_k)}_d$$

the linearization above is applied, resulting in

$$\sigma_i^2 \simeq \mathbf{n}_{jk}^T \sigma_d \mathbf{n}_{jk} + \mathbf{d}^T \Sigma_{n_{jk}} \mathbf{d} \quad (25)$$

$$\Sigma_{n_{jk}} = [R_j \mathbf{l}_j]_x R_k \Sigma_{l_k} R_k^T [R_j \mathbf{l}_j]_x^T + [R_k \mathbf{l}_k]_x R_j \Sigma_{l_j} R_j^T [R_k \mathbf{l}_k]_x^T \quad (26)$$

$$\sigma_{d_i}^2 = R_j \Sigma_{c_j} R_j^T + [\mathbf{c}_j]_x \Sigma_{R_j} [\mathbf{c}_j]_x^T + R_k \Sigma_{c_k} R_k^T + [\mathbf{c}_k]_x \Sigma_{R_k} [\mathbf{c}_k]_x^T \quad (27)$$

where $[\cdot]_x$ is the skew-symmetric matrix operator.

For the error distribution in eq. 12,

$$r_i = \mathbf{n}_{jk}^a \cdot \mathbf{n}_{jk}^b \quad (28)$$

the linearization above is applied as

$$\sigma_i^2 = (\mathbf{n}_{jk}^a)^T R_j^T R_k \Sigma_{\mathbf{n}_{jk}^b} R_k^T R_j \mathbf{n}_{jk}^a + (\mathbf{n}_{jk}^b)^T R_k^T R_j \Sigma_{\mathbf{n}_{jk}^a} R_j^T R_k \mathbf{n}_{jk}^b \quad (29)$$

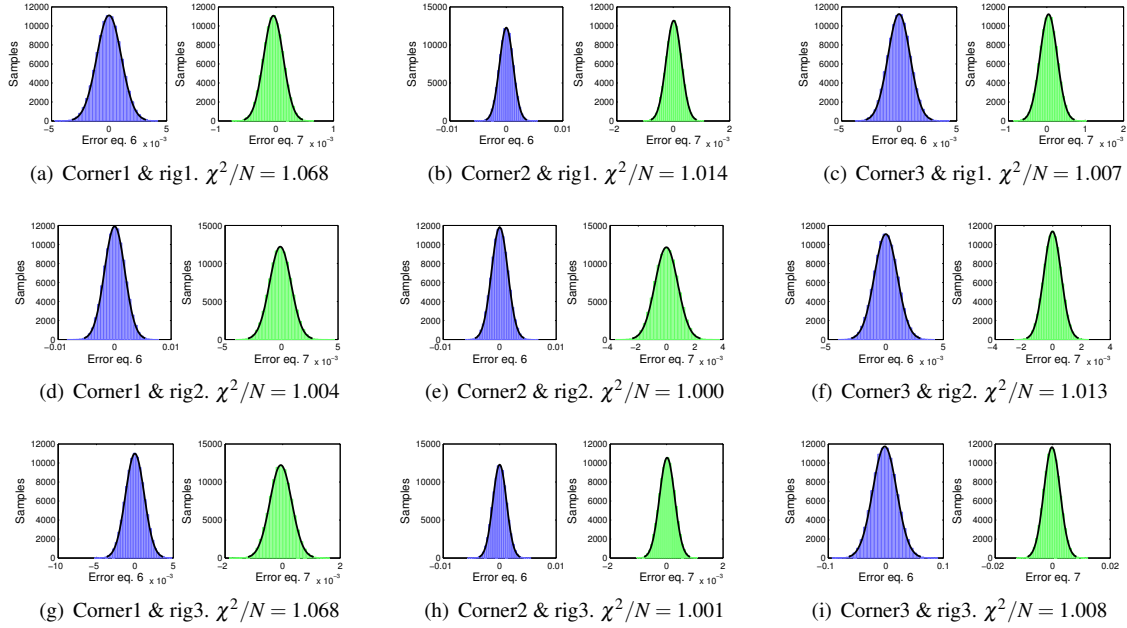


Figure 15: Monte Carlo simulation of a corner observation at three different orientations (a different corner is shown in each column) with three different rigs of two LRFs (a different rig is shown in each row). The distribution of the planarity and perpendicularity errors are shown respectively at the left and right of each graph, together with the Gaussian functions which best fit these distributions. For each one, we compute the χ^2 goodness of fit test, where N represents the number of samples.

These linearizations imply that the product of two Gaussian random variables is approximated as a Gaussian distribution, despite that the result follows a χ^2 distribution. The error of this approximation diminishes asymptotically with the number of samples (observations) (Severo and Zelen, 1960; Zhang, 2005).

8.3 Fisher Information and Cramér-Rao Bound

The Fisher Information Matrix (FIM) is a statistic measure of how much information an observable random variable X carries about the parameters $\theta = [\theta_1, \theta_2, \dots, \theta_N]^T$ upon which the likelihood function $P(X|\theta)$ depends. For an unbiased estimator, the FIM is defined as

$$\mathcal{I}(\theta) = \mathbb{E}[(\nabla_{\theta} \log P(X|\theta))(\nabla_{\theta} \log P(X|\theta))^T | \theta] \quad (30)$$

which is a $N \times N$ positive semi-definite symmetric matrix. Formally, the FIM corresponds to the expected value of the observed information. It has important

implications regarding the observability of estimation problems,

concretely, the problem is observable (it has a solution) if and only if the FIM has full rank, i.e. $\text{rank}(\mathcal{I}) = N$.

Considering an estimation problem in which the observable random variables X follow an unbiased, asymptotically Gaussian distribution, the FIM can be calculated as

$$\mathcal{I}(\theta) = J_E^T \Sigma_X J_E \quad (31)$$

where J_E is the Jacobian of the observation equations (i.e. the cost function of the estimation problem) (Van Trees and Bell, 2007). This formulation has a relevant significance in this paper, implying that the FIM of the maximum likelihood estimation in eq. 8 corresponds to the Hessian of the negative log-likelihood of eq. 9, which is minimized through least squares, where the Hessian is given in eq. 16.

The Cramér-Rao Bound (CRB) defines a lower bound for the variance of estimators, therefore, it is a measure of the performance of estimators (Van Trees

and Bell, 2007). The CRB states that the covariance of an unbiased estimator cannot be lower than the inverse of the Fisher information (Van Trees and Bell, 2007).

$$\text{cov}(\hat{\theta}) \geq \mathcal{I}^{-1}(\theta) \quad (32)$$

where $\hat{\theta}$ is the true value of the estimator. If this limit is achieved, the estimator is called efficient. Thus, the FIM defines a lower bound for the noise of our estimate, which can be used to find the best estimator for the problem.

8.4 Lie algebra and Lie groups

A Lie algebra can be described as a representation of a Lie group in a vector space where infinitesimal transformations can be applied. The special orthogonal group $\mathbb{SO}(3)$ which represents all rotations in 3D Euclidean space \mathbb{R}^3 , is an example of a Lie group. Each rotation in $\mathbb{SO}(3)$ is expressed as a 3×3 orthonormal matrix. The Lie algebra corresponding to the Lie group $\mathbb{SO}(3)$ is expressed as $\mathfrak{so}(3)$, and coincides with \mathbb{R}^3 . It constitutes a minimal parameterization for the rotations which can be arithmetically manipulated as a vector space. The transformation from the Lie algebra $\mathfrak{so}(3)$ to its corresponding Lie group $\mathbb{SO}(3)$ is defined by the exponential map operation

$$\begin{aligned} \exp: \mathfrak{so}(3) &\rightarrow \mathbb{SO}(3) \\ \omega &\rightarrow R \end{aligned}$$

which is given by the Rodrigues' formula

$$R = I + \frac{\sin(\theta)}{\theta} [\omega]_{\times} + \frac{(1 - \cos(\theta))}{\theta^2} [\omega]_{\times}^2 \quad (33)$$

where $\theta = |\omega|$, and $[\cdot]_{\times}$ represents the skew-symmetric matrix operator. The inverse operation is called logarithm map, which in this case is obtained from

$$\theta = \arccos\left(\frac{\text{trace}(R) - 1}{2}\right) \quad (34)$$

$$\omega = \begin{cases} 0 & \text{if } \theta = 0 \\ \frac{\theta}{2\sin(\theta)}(R - R^{\top}) & \text{if } \theta \neq 0 \text{ and } \theta \in (-\pi, \pi) \end{cases} \quad (35)$$

Figure 16 shows an intuitive representation for the smooth manifold of a Lie group (e.g. $\mathbb{SO}(3)$) and the Euclidean tangent space of its Lie algebra, showing the transformations between both. Note that the tangent point is located at the identity of the Lie group.




Figure 16: Lie group and Lie algebra for the space of 3D rotations.

The special Euclidean group $\mathbb{SE}(3)$ represents 3D rigid motions (rotation plus translation), and an element is expressed as a 4×4 matrix

$$T = [R|\mathbf{t}] = \left[\begin{array}{c|c} R & \mathbf{t} \\ \hline 0 & 1 \end{array} \right] \quad (36)$$

where the rotation $R \in \mathbb{SO}(3)$ and the translations $\mathbf{t} \in \mathbb{R}^3$. The optimization algorithm presented in this paper makes use of the above transformations between Lie groups and Lie algebras, so that direct operations to compute errors employ the Lie group formulation, while the optimization problem is derived in the tangent space of the manifold given by the Lie algebra. For a more detailed description of Lie groups and Lie algebras in the context of mobile robotics, the reader is referred to (Blanco, 2010).

9 ACKNOWLEDGMENT

The authors would like to thank our colleges in the group MAPIR, and specially Francisco Melendez, for their help to carry out experiments and to develop supplementary material for this paper. This work has been funded by the Spanish Government under project DPI2011-25483 and the EU FP7 Programme under the GiraffPlus Project.

References

- Arras, K. O. and Siegwart, R. Y.: 1998, Feature extraction and scene interpretation for map-based navigation and map building, *Intelligent Systems & Advanced Manufacturing*, pp. 42–53.
- Barber, D., Mills, J. and Smith-Voysey, S.: 2008, Geometric validation of a ground-based mobile laser scanning system, *ISPRS Journal of Photogrammetry and Remote Sensing* **63**(1), 128–141.

- Blanco-Claraco, J.-L., Moreno-Dueñas, F.-Á. and González-Jiménez, J.: 2014, The Málaga urban dataset: High-rate stereo and lidar in a realistic urban scenario, *The International Journal of Robotics Research* **33**(2), 207–214.
- Blanco, J.-L.: 2010, A tutorial on $se(3)$ transformation parameterizations and on-manifold optimization, *University of Malaga, Tech. Rep.*
- Blanco, J.-L., Moreno, F.-A. and Gonzalez, J.: 2009, A collection of outdoor robotic datasets with centimeter-accuracy ground truth, *Autonomous Robots* **27**(4), 327–351.
- Bohren, J., Foote, T., Keller, J., Kushleyev, A., Lee, D., Stewart, A., Vernaza, P., Derenick, J., Spletzer, J. and Satterfield, B.: 2008, Little ben: The ben franklin racing team’s entry in the 2007 darpa urban challenge, *Journal of Field Robotics* **25**(9), 598–614.
- Borges, G. A. and Aldon, M.-J.: 2000, A split-and-merge segmentation algorithm for line extraction in 2d range images, *IEEE CVPR 2000*, Vol. 1, IEEE, pp. 441–444.
- Borrmann, D., Elseberg, J., Lingemann, K., Nüchter, A. and Hertzberg, J.: 2008, Globally consistent 3d mapping with scan matching, *Robotics and Autonomous Systems* **56**(2), 130–142.
- Campbell, M., Egerstedt, M., How, J. P. and Murray, R. M.: 2010, Autonomous driving in urban environments: approaches, lessons and challenges, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **368**(1928), 4649–4672.
- Censi, A., Franchi, A., Marchionni, L. and Oriolo, G.: 2013, Simultaneous calibration of odometry and sensor parameters for mobile robots, *Robotics, IEEE Transactions on* **29**(2), 475–492.
- Diosi, A. and Kleeman, L.: 2003, Uncertainty of line segments extracted from static sick pls laser scans, *ACRA 2003*.
- Fernández-Madrigal, J.-A. and Claraco, J. L. B.: 2013, *Simultaneous Localization and Mapping for Mobile Robots: Introduction and Methods*, Information Science Reference.
- Fernández-Moral, E.: 2014, *Contributions to metric-topological localization and mapping in mobile robotics*, PhD thesis, Universidad de Málaga.
- Fernández-Moral, E., Arévalo, V. and González-Jiménez, J.: 2015, Extrinsic calibration of a set of 2d laser rangefinders, *International Conference on Robotics and Automation (ICRA)*, IEEE.
- Fernández-Moral, E., González-Jiménez, J., Rives, P. and Arévalo, V.: 2014, Extrinsic calibration of a set of range cameras in 5 seconds without pattern, *IEEE/RSJ IROS 2014*, IEEE.
- Fischler, M. A. and Bolles, R. C.: 1981, Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography, *Communications of the ACM* **24**(6), 381–395.
- Forsberg, J., Larsson, U. and Wernersson, A.: 1995, Mobile robot navigation using the range-weighted hough transform, *Robotics & Automation Magazine, IEEE* **2**(1), 18–26.
- Gao, C. and Spletzer, J. R.: 2010, On-line calibration of multiple lidars on a mobile vehicle platform, *IEEE ICRA 2010*, IEEE, pp. 279–284.
- Glas, D. F., Miyashita, T., Ishiguro, H. and Hagita, N.: 2010, Automatic position calibration and sensor displacement detection for networks of laser range finders for human tracking, *IEEE/RSJ IROS 2010*, IEEE, pp. 2938–2945.
- Grisetti, G., Kummerle, R., Stachniss, C. and Burgard, W.: 2010, A tutorial on graph-based slam, *Intelligent Transportation Systems Magazine, IEEE* **2**(4), 31–43.
- Haala, N., Peter, M., Kremer, J. and Hunter, G.: 2008, Mobile lidar mapping for 3d point cloud collection in urban areas - a performance test, *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* **37**, 1119–1127.
- Huang, A. S., Antone, M., Olson, E., Fletcher, L., Moore, D., Teller, S. and Leonard, J.: 2010, A high-rate, heterogeneous data set from the darpa urban challenge, *The International Journal of Robotics Research* **29**(13), 1595–1601.
- Leonard, J., How, J., Teller, S., Berger, M., Campbell, S., Fiore, G., Fletcher, L., Frazzoli, E., Huang, A.,

- Karaman, S. et al.: 2008, A perception-driven autonomous urban vehicle, *Journal of Field Robotics* **25**(10), 727–774.
- Martinelli, A.: 2011, State estimation based on the concept of continuous symmetry and observability analysis: The case of calibration, *Robotics, IEEE Transactions on* **27**(2), 239–255.
- Miller, I., Campbell, M. and Huttenlocher, D.: 2011, Efficient unbiased tracking of multiple dynamic obstacles under large viewpoint changes, *Robotics, IEEE Transactions on* **27**(1), 29–46.
- Nguyen, V., Martinelli, A., Tomatis, N. and Siegwart, R.: 2005, A comparison of line extraction algorithms using 2d laser rangefinder for indoor mobile robotics, *IEEE/RSJ IROS 2005*, IEEE, pp. 1929–1934.
- Pathak, K., Vaskevicius, N. and Birk, A.: 2010, Uncertainty analysis for optimum plane extraction from noisy 3d range-sensor point-clouds, *Intelligent Service Robotics* **3**(1), 37–48.
- Petrovskaya, A. and Thrun, S.: 2009, Model based vehicle detection and tracking for autonomous urban driving, *Autonomous Robots* **26**(2-3), 123–139.
- Prasad, D., Quek, C., Leung, M. K. H. and Cho, S.-Y.: 2011, A parameter independent line fitting method, *Pattern Recognition (ACPR), 2011 First Asian Conference on*, IEEE, pp. 441–445.
- Schenk, K., Kolarow, A., Eisenbach, M., Debes, K. and Gross, H.-M.: 2012, Automatic calibration of a stationary network of laser range finders by matching movement trajectories, *IEEE/RSJ IROS 2012*, IEEE, pp. 431–437.
- Severo, N. C. and Zelen, M.: 1960, Normal approximation to the chi-square and non-central f probability functions, *Biometrika* **47**(3-4), 411–416.
- Thrun, S., Burgard, W. and Fox, D.: 2000, A real-time algorithm for mobile robot mapping with applications to multi-robot and 3d mapping, *IEEE ICRA 2000*, Vol. 1, IEEE, pp. 321–328.
- Thrun, S., Montemerlo, M., Dahlkamp, H., Stavens, D., Aron, A., Diebel, J., Fong, P., Gale, J., Halpenny, M., Hoffmann, G. et al.: 2006, Stanley: The robot that won the darpa grand challenge, *Journal of field Robotics* **23**(9), 661–692.
- Van Trees, H. L. and Bell, K. L.: 2007, Bayesian bounds for parameter estimation and nonlinear filtering/tracking, *AMC* **10**, 12.
- Victorino, A. C., Rives, P. and Borrelly, J.-J.: 2003, Safe navigation for indoor mobile robots. part i: a sensor-based navigation framework, *The International Journal of Robotics Research* **22**(12), 1005–1118.
- Zhang, J.-T.: 2005, Approximate and asymptotic distributions of chi-squared-type mixtures with applications, *Journal of the American Statistical Association* **100**(469), 273–285.