

Extrinsic Calibration of a 2D Laser-Rangefinder and a Camera based on Scene Corners

Ruben Gomez-Ojeda, Jesus Briales, Eduardo Fernandez-Moral and Javier Gonzalez-Jimenez

Abstract—Robots are often equipped with 2D laser-rangefinders (LRFs) and cameras since they complement well to each other. In order to correctly combine measurements from both sensors, it is required to know their relative pose, that is, to solve their extrinsic calibration. In this paper we present a new approach to such problem which relies on the observations of orthogonal trihedrons which are profusely found as corners in human-made scenarios. Thus, the method does not require any specific pattern, which turns the calibration process fast and simpler to perform. The estimated relative pose has proven to be also very precise since it uses two different types of constraints, *line-to-plane* and *point-to-plane*, as a result of a richer configuration than previous proposals that relies on plane or V-shaped patterns. Our approach is validated with synthetic and real experiments, showing better performance than the state-of-art methods.

I. INTRODUCTION

The combination of a laser-rangefinder (LRF) and a camera is a common practice in mobile robotics. Some examples are the acquisition of urban models [1] [2], the detection of pedestrians [3], or the construction of semantic maps [4]. In order to effectively exploit measurements from both type of sensors, a precise estimation of their relative pose, that is, their extrinsic calibration, is required. This paper presents a method for such extrinsic calibration which relies on the observation of three perpendicular planes (orthogonal trihedron), which can be found in any structured scene, for instance, buildings. This idea to calibrate the sensors from the elements of the environment was inspired by our previous work for RGB-D cameras [5] and LRFs [6]. In a nutshell, the calibration process is performed by first extracting the three plane normals from the projected junctions of the trihedron, and then imposing co-planarity between the scanned points and those planes.

A. Related Work

The most precise and effective strategy to perform the extrinsic calibration between a camera and a LRF is by establishing some kind of data association between the sensor measurements. For that, the intuitive approach is to detect the laser spot in the image, but this is rarely feasible since most LRFs employ invisible light beams. Then, the common

The authors are with the Mapir Group of Department of System Engineering and Automation, University of Málaga, Spain. +E-mail: {rubengooj|jesusbriales}@gmail.com

This work has been supported by two projects: "GiraffPlus", funded by EU under contract FP7-ICT-#288173, and "TAROTH: New developments toward a robot at home", funded by the Spanish Government and the "European Regional Development Fund ERDF" under contract DPI2011-25483.




Fig. 1. Observation of a trihedron structure, which is defined by three orthogonal planes $\{\Pi_1, \Pi_2, \Pi_3\}$ intersecting at three orthogonal lines $\{L_1, L_2, L_3\}$, by a rig formed by a 2D LRF and a camera.

practice is to establish geometric constraints from the association of different 3D features (e.g. points, lines and planes) observed simultaneously by both sensors. Depending on the nature of the detected features, two methodologies have been considered in the literature. The first one employs *point-to-line* restrictions, establishing correspondences between some identifiable scanned points and line segments in the image. A typical calibration target for this technique is the V-shaped pattern proposed by Wasielewski and Strauss [7]. Their approach computes the intersection point of the LRF scan plane with the V-shaped target, and then minimizes the distance from the projected point to the line segment detected in the image. In general, this procedure requires a large number of observations to have a well-determined problem and to reduce the error introduced by the mismatch between the scan point and the observed line. To overcome this limitation, in [8] the number of *point-to-line* correspondences in each observation is increased to three by introducing virtual end-points, and also the effect of outliers is lowered down with a Huber weighting function.

A different strategy is the one proposed by Zhang and Pless [9], which makes use of a planar checkerboard pattern. Similarly to the camera calibration method in [10], they first estimate the relative pose between the camera and the pattern in each observation. Then, they impose *point-to-plane* restrictions between the laser points and the pattern plane to obtain a linear solution, which is employed as initial value in a final non-linear optimization process. This approach has two problems: the initial value may not be a valid pose, since there is no guarantee that the rotation $\mathbf{R} \in SO(3)$, and it is often a poor guess which leads to a local minimum. These disadvantages are addressed in [11], where the authors reformulate the estimation as a linear Perspective

n Point (PnP) problem, proving that three *line-to-plane* correspondences provide at most eight solutions. They report a minimal solution to the problem in a RANSAC framework and refine it with a non-linear iterative optimization. The main issue of this algorithm is its limited accuracy and its numerical instability. These drawbacks have been recently addressed by Zhou [12], with another improved version of the Zhang and Pless proposal. They report better results in accuracy and numerical stability with a novel minimal solution, employing three *line-to-plane* correspondences as well, estimating the rotation from a polynomial equation system and then calculating the translation with a linear system. However, their method still requires the use of a large checkerboard pattern and suffers from a limited accuracy and robustness due to the weakness of the applied constraints.

B. Contribution

This paper presents the first approach to the calibration of a 2D LRF and a camera without the need of a specific calibration pattern, which makes the method fast and simpler to apply. Concretely, this solution only requires the observation of a scene corner (orthogonal trihedron) commonly found in any human-made environment. Moreover, unlike previous approaches, the proposed method exploits both *line-to-plane* and *point-to-plane* constraints, which yields a more accurate and robust estimation.

The source code of the developed MATLAB library is available online, and will be updated as research progresses. An illustrative video of our calibration system and the source code can be found here: <http://mapir.isa.uma.es>.

II. EXTRACTING CALIBRATION DATA FROM THE TRIHEDRON

Our proposal to calibrate a camera-LRF system relies on the observation of an orthogonal trihedron τ , defined as a structure formed by three perpendicular planes, denoted by $\{\Pi_1, \Pi_2, \Pi_3\}$ (see Figure 1). These planes intersect with each other at three orthogonal lines, $\{L_1, L_2, L_3\}$, so that all planes and lines meet at the vertex \mathbf{C} . The first step of our method consists in detecting these trihedron features from the observations of both sensors, as explained next.

A. Information from the Camera

1) *Trihedron Detection and Tracking*: A trihedron is projected onto the image as three lines intersecting at a vertex (see Figure 2). Hence, it can be defined as a 4-tuple $\xi = (\mathbf{c}, \nu_1, \nu_2, \nu_3)$, where ν_k represents the director vector of the k -th line and \mathbf{c} is the vertex. To gain in accuracy, traditional methods for detecting line segments, such as the LSD [13], are not appropriate since they do not force the three lines to meet at a single point.

Instead, we propose a tracking method based on the fitting of the structure ξ using the image gradient information. The tracking is initialized automatically from the three intersecting lines detected with the LSD detector in a RANSAC framework. After this first step, a rectangular region S_k is defined around each line as depicted in Figure 2 (the




Fig. 2. A single camera observation of a trihedron structure, which is parametrized by its vertex \mathbf{c} and the directions ν_k of the three lines. For its tracking in an image sequence, a region S_k for each line is considered.

part of the line close to the vertex is discarded to avoid unreliable gradients). Within these regions, only the pixels with a gradient magnitude above a threshold \bar{G}_k , and angle almost perpendicular to the line border are considered for the fitting. Let η_k be the normal vector of the k -th line. Hence, the distance from a generic pixel in the k -th region, \mathbf{p}_k^i , to the k -th line is expressed as

$$d_k^i = (\mathbf{p}_k^i - \mathbf{c})^\top \cdot \eta_k \quad (1)$$

Then, the trihedron image parameters ξ are estimated with a Levenberg-Marquardt optimization over the weighted non-linear least squares problem, expressed as

$$\underset{\xi}{\operatorname{argmin}} \sum_{k=1}^3 \sum_{i=1}^{N_k} G_k^i \cdot d_k^i{}^2 \quad (2)$$

where the weights are the gradient magnitude G_k^i of the pixels \mathbf{p}_k^i in the region, and N_k is the number of valid pixels for each region S_k . The initialization for the trihedron detection is only necessary in the first frame. After this, the structure is tracked from the previous frame automatically.

2) *Back-Projected Planes*: The lines of the trihedron in the image, \mathbf{l}_k , are defined by its direction ν_k and the vertex \mathbf{c} , as estimated above. Each line \mathbf{l}_k gives rise to a so-called back-projected plane ${}^\pi\Pi_k$ containing the camera optical center (see Figure 3), which normal vector ${}^\pi\mathbf{n}_k$ is given by the expression [14]

$${}^\pi\mathbf{n}_k = \frac{\mathbf{K}^\top \mathbf{l}_k}{\|\mathbf{K}^\top \mathbf{l}_k\|} \quad (3)$$

where the camera matrix \mathbf{K} is supposed to be known.

3) *Trihedron Plane Normals*: From the computation of ξ and ${}^\pi\mathbf{n}_k$, the relative rotation of the trihedron w.r.t. the camera can be estimated as follows. The back-projected plane and the trihedron plane, defined by their normal directions ${}^\pi\mathbf{n}_k$, are orthogonal to each other (see Figure 3), that is

$${}^\pi\mathbf{n}_k^\top \cdot {}^\pi\mathbf{n}_k = 0 \quad (4)$$

for all three planes $k = \{1, 2, 3\}$ in the trihedron. Notice that the line L_k is perpendicular to its opposite plane Π_k , and it is also co-planar to the back-projected plane ${}^\pi\Pi_k$, as depicted in Figures 1 and 3. Additionally, by imposing the




Fig. 3. Scheme of the back-projected planes.

condition that the normal vectors are unitary, the following six constraints apply:

$$\tau \mathbf{R}^\top \cdot \tau \mathbf{R} = \mathbf{I}_3 \quad (5)$$

with $\tau \mathbf{R} = [\tau \mathbf{n}_1 \ \tau \mathbf{n}_2 \ \tau \mathbf{n}_3]$ being the relative rotation of the trihedron w.r.t. the camera. This states a quadratic equation system which is solved efficiently by employing Gauss-Newton optimization and on-manifold derivatives [15]. For that, we need to re-formulate the system as a vector field $\Phi : SO(3) \rightarrow \mathbb{R}^3$, and to force the field Φ to be zero:

$$\Phi(\tau \mathbf{R}) = \begin{bmatrix} \pi \mathbf{n}_1^\top \cdot \tau \mathbf{R} \cdot \mathbf{e}_1 \\ \pi \mathbf{n}_2^\top \cdot \tau \mathbf{R} \cdot \mathbf{e}_2 \\ \pi \mathbf{n}_3^\top \cdot \tau \mathbf{R} \cdot \mathbf{e}_3 \end{bmatrix} = \mathbf{0} \quad (6)$$

with $\mathbf{e}_1 = [1 \ 0 \ 0]^\top$, $\mathbf{e}_2 = [0 \ 1 \ 0]^\top$, and $\mathbf{e}_3 = [0 \ 0 \ 1]^\top$. Notice that this equation system has 8 possible solutions, one for each quadrant, which might lead the method to a wrong minimum. This can be overcome by initializing the algorithm with a orientation in the actual quadrant, and by comparing the sign of the solution with the projected image directions. This optimization is run at each new observation in the tracking process, thus, it is initialized with the solution from the previous step.

B. LRF Data

The scene corner (trihedron) is sampled by the LRF as three line segments, as shown in Figure 1, which can be extracted in a number of ways [16]. Here, we have implemented a segmentation method based on RANSAC, which searches for the parameters $\{s\nu_x, s\nu_y, \kappa\}$ of a 2D line which maximizes the number of inliers fulfilling the next model

$$s\nu_x x_i + s\nu_y y_i + \kappa \leq \epsilon \quad (7)$$

with the scan points represented by $\{x_i, y_i\}$, and ϵ is a threshold employed to filter the outliers. An interesting advantage of this procedure is that unconnected collinear segments are automatically clustered as the same line, thus the line models obtained are more accurate. For the calibration process, we are interested in the line directions $s\nu = [s\nu_x \ s\nu_y]^\top$ and the corner points \mathbf{q} obtained at the intersection of the scan lines.

III. EXTRINSIC CALIBRATION APPROACH

A. Problem Statement

Let us define a Trihedron Observation TO_i as a set formed by the line segment directions in the scan $s\nu$, the scan intersection points \mathbf{q} , the trihedron plane normals $\tau \mathbf{n}$, and the back-projected plane normals $\pi \mathbf{n}$, expressed as

$$TO_i \equiv \{s\nu_k^i, \mathbf{q}_k^i, \tau \mathbf{n}_k^i, \pi \mathbf{n}_k^i\} \quad (8)$$

where $k = \{1, 2, 3\}$ is the index of the plane correspondence, and $i = 1, \dots, N$ is the index of the observations.

It can be verified that given a minimal sample of three TO correspondences, the problem can be decoupled to obtain the rotation and the translation separately, since both have a closed-form solution. Hence, the optimization algorithm consists of two steps.

First, we obtain the relative rotation between the camera and the LRF, $\mathbf{R} \in SO(3)$, by imposing co-planarity constraints between the lines and planes segmented (*line-to-plane*). This condition is expressed as

$$\tau \mathbf{n}^\top \cdot \mathbf{R} s\nu = 0 \quad (9)$$

where $\tau \mathbf{n} \in S^2$ is the normal vector of the corresponding plane, $s\nu \in \mathbb{R}^3$ is the direction of the scan segment in that plane, and $\mathbf{R} \in SO(3)$ is the relative rotation of the LRF w.r.t. the camera. Notice that the measurements from the scan belongs to the plane defined by \mathbf{x}_s and \mathbf{y}_s (see Figure 1), and they can be transformed to 3D vectors by setting the \mathbf{z}_s coordinate to be zero.

Secondly, we compute the relative translation $\mathbf{t} \in \mathbb{R}^3$ by imposing co-planarity constraints between the scan corner points \mathbf{q} , and the planes back-projected from camera center through trihedron lines (*point-to-plane*). This condition is given by

$$\pi \mathbf{n}^\top \cdot (\mathbf{R} \mathbf{q} + \mathbf{t}) = 0 \quad (10)$$

for each correspondence, with $\pi \mathbf{n} \in S^2$ being the back-projected plane normal.

B. Optimization

1) *Formulation*: Given a set of TO s extracted from different poses of the Camera-LRF rig, the aim is to find the relative pose $[\mathbf{R}|\mathbf{t}]$ that minimizes the errors of the constraints in (9) and (10). For that, we can assume independence between the TO s and that the measurements are affected by unbiased Gaussian noise. Without loss of generality, we take the camera pose as the reference coordinates.

Since the unknowns \mathbf{R} and \mathbf{t} are decoupled, the above problem can be formulated as two independent maximum likelihood estimation (MLE) processes for the given TO s, whose solutions are obtained by solving the weighted non-linear least squares problem expressed as

$$\underset{\mathbf{x}}{\operatorname{argmin}} \sum_{i=1}^N \mathbf{r}_i^\top \mathbf{w}_i \mathbf{r}_i \quad (11)$$

where \mathbf{r}_i is the residual error defined for each problem, and \mathbf{w}_i the weight of the corresponding residual from TO_i .

All the weights employed in this work derive from the propagation of the uncertainty from the sensor measurements ($w_i = \Sigma_i^{-1}$), which has been propagated analytically with a first order approximation to the covariance matrix [17]. Finally, RANSAC is used to discard possible outliers in the TO correspondences, before the estimation of the calibration.

2) *Solving for the Rotation:* The problem in (11) is reformulated using Lie algebra so the optimization steps are performed on the manifold tangent space $\mathfrak{so}(3)$ [15]:

$$\mathbf{R}_{n+1} = e^{\varepsilon^n} \mathbf{R}_n \quad (12)$$

where \mathbf{R}_n refers to the rotation value in the n -th iteration, and $e^{\varepsilon^n} \in SO(3)$ is the infinitesimal pose increment. As expressed in (9), the rotation is estimated imposing that the LRF line segments belongs to the trihedron plane normals. Hence, the residual error is given by

$$r_k^i = \boldsymbol{\pi} \mathbf{n}_k^i \top \cdot \mathbf{R}^s \boldsymbol{\nu}_k^i \quad (13)$$

$$\mathbf{r}^i = [r_1^i \ r_2^i \ r_3^i]^\top \quad (14)$$

3) *Solving for the Translation:* Once the rotation is estimated, the translation is recovered by imposing three coplanarity conditions in the 3D space. Concretely, the scan corner point \mathbf{q}_k (in the LRF reference frame) must lie on the k -th back-projected plane defined by the normal $\boldsymbol{\pi} \mathbf{n}_k$, as expressed in (10) (see Figure 3). Hence, the optimal translation is obtained by solving the optimization problem in (11) with Levenberg-Marquardt defining the following residual error

$$r_k^i = \boldsymbol{\pi} \mathbf{n}_k^i \top \cdot (\mathbf{R} \mathbf{q}_k^i + \mathbf{t}) \quad (15)$$

for each observation TO_i , and the correspondent weights obtained by propagating the rotation uncertainty from previous steps.

IV. EXPERIMENTAL RESULTS

In this section we validate the proposed approach with a number of experiments with both synthetic and real data. We compare our results with several methods, for which we have implemented the paper versions of [7] [8], and also with the method in [11].

A. Simulation

The first set of experiments has been designed to prove the correctness of the proposed method and to compare our results with the state-of-art methods. In our simulation environment, a rig formed by a LRF and a camera is randomly generated at different relative poses in the range of $\pm 45^\circ$ and ± 50 cm for the rotation and translation respectively. The simulated sensors are characterized as the Sick LMS 200 and the stereo camera Point Grey Bumblebee[®]2 employed in the real experiments. In order to compare with the techniques previously presented in literature, we have also simulated two calibration targets besides the trihedron employed in our algorithm. First, we have employed the traditional checkerboard with 8×8 squares, for those algorithms which exploits *point-to-plane* constraints (see [11]). Also, we have created a V-shaped target to compare with the *point-to-line* algorithms

[7] [8]. For the sake of fairness, all the targets used here have similar dimensions, i.e. a $L \times L$ checkerboard, a corner with a height of $2L$ and sides of L , and a trihedron with sides of L , with $L = 1.5$ m. Notice that whereas the size of the two targets has been set with realistic values, the trihedron structures present in most structured scenes would have bigger dimensions, which would produce better results than those shown here.

The observations have been generated with variable levels of unbiased and uncorrelated Gaussian noise, σ_s and σ_c for both scan and camera observations, respectively. Also, we have tested the effect of the number of correspondences in the different methods by performing the same experiments with variable amounts of correspondences. The number of Monte Carlo simulations has been set to 500 configurations for each calibration method. In order to estimate the accuracy of the calibration results, the following error metrics are employed

$$e_{\mathbf{R}} = 2 \arcsin \left(\frac{\|\mathbf{R} - \hat{\mathbf{R}}\|_{\mathfrak{F}}}{2\sqrt{2}} \right) \quad (16)$$

for the rotation error and

$$e_{\mathbf{t}} = \|\mathbf{t} - \hat{\mathbf{t}}\| \quad (17)$$

for translation, where $[\mathbf{R}|\mathbf{t}]$ is the estimated pose and $[\hat{\mathbf{R}}|\hat{\mathbf{t}}]$ is the ground truth. The value of $e_{\mathbf{R}}$ represents the minimal angular distance between two rotations in $SO(3)$ in radians, while $e_{\mathbf{t}}$ is the Euclidean distance in meters between two \mathbb{R}^3 vectors.

In order to combine the noise information of both sensors in one plot, the noises have been modeled with a variable factor of proportionality k_σ which multiplied by the standard deviation of each sensor, i.e. $\sigma_c = 1$ pixel and $\sigma_s = 0.03$ m for the camera and the LRF respectively, provides the noise standard deviation for each sensor

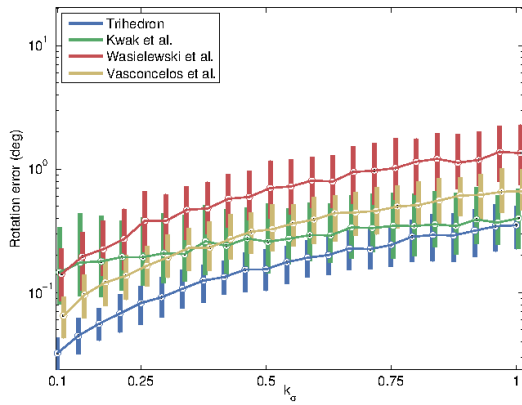
$$\hat{\sigma}_j = k_\sigma \sigma_j \quad (18)$$

with $j = \{c, s\}$ for the camera and the scan, respectively. The average errors of the relative rotation and translation are shown in Figure 4 in degrees and meters, respectively. The experiment has been performed with a fixed number of 20 correspondences. The graphic shows the better results of our method in accuracy and precision.


The effects of the number of correspondences are represented in Figure 5, where the standard deviation of the LRF is set to $\sigma_s = 0.03$ m, and that of the camera is set to $\sigma_c = 1$ pixel. We observe that errors decrease asymptotically with the number of correspondences. Also, the results provided by our method are superior to the state-of-art methods, specially with low number of correspondences.

B. Real Data

The rig employed in our experiments is composed of a Sick LMS 200 and a stereo camera Point Grey Bumblebee[®]2 (see Figure 6). The experiments consist in estimating the calibration of the laser w.r.t. the left and right camera, so



(a) Rotation error



(b) Translation error

Fig. 4. Simulated calibration results comparing our method with [11] [7] and [8]. The experiments have been performed for 20 correspondences, and are plotted for increasing levels of gaussian noise (proportional to k_σ).

that we can compute the relative pose of the stereo set and compare it with the calibration provided by the manufacturer, which is considered here as ground truth (a similar procedure was used in [18]). The results are showed in Table I for a variable number of correspondences. We observe that the average errors decrease when raising the number of trihedron observations. We have also tested the calibration results

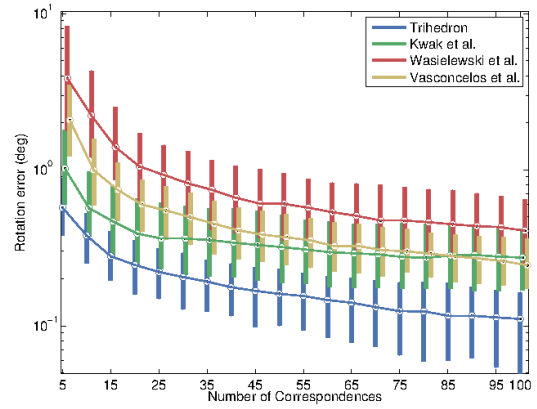
TABLE I

AVERAGE ERRORS OF REAL CALIBRATIONS WITH DIFFERENT NUMBER OF CORRESPONDENCES.


TOs	5	10	20	50	100
e_R (deg)	0.7912	0.6307	0.4218	0.3911	0.2927
e_t (cm)	1.1664	0.7809	0.6776	0.5787	0.4269

obtained with our method by projecting the laser points into the image. This experiment is showed in Figure 8, where the projected points from the LRF fit to the objects in the scene. The line segments have been detected with the method employed in Section II-B and have been represented with a different color for each object. The experiment validates the performance of our method as the line segments match with the planar surfaces from the scene.

Also, we have performed an experiment which allows us to observe the real LRF points on different surfaces. For



(a) Rotation error



(b) Translation error

Fig. 5. Simulated calibration results comparing our method with [11] [7] and [8]. The experiments have been performed for a fixed noise level of $\sigma_s = 0.03$ m and $\sigma_c = 1$ pixel, increasing the number of correspondences.




Fig. 6. Rig employed in the real simulations formed by a Point Grey Bumblebee@2 and a Sick LMS 200.


that, we have employed the sensitivity of common cameras to the infrared beams of some LRFs [19]. To avoid the use of special infrared filters, we have recorded a scene in dark conditions with long exposure times. An example of such images is depicted in Figure 7, where the exposure time of the stereo camera was set to 2 min. This experiment serves for visual evaluation of the accuracy of our method, which is evident as the lines segmented from the laser scans coincide with the faces of the different boxes in the scene.

V. CONCLUSIONS

A new methodology for calibrating the extrinsic parameters of a rig formed by a 2D LRF and a camera has been presented in this paper. It relies on the observations of an orthogonal trihedron typically encountered as scene corners in buildings and human-made structures. The method



(a) Ground truth image



(b) Projected LRF points in the illuminated scene

Fig. 7. (a) The real LRF beam of the Sick LMS 200 can be observed in the objects of dark scenes when the image is acquired with a high exposure time. (b) The correctness of our estimation can be checked visually by comparing the real LRF beam trace in (a) with the projection of the LRF points in the same illuminated scene with the estimated calibration.



Fig. 8. Projection of the LRF points that belongs to planar surfaces onto the image with the estimated calibration. Our estimation is visually validated as the projection of points fits with the planar objects in the image.

establishes correspondences between the plane observations from the camera and the line segmented from the laser scans data. The problem is solved in a probabilistic framework which takes into account the uncertainty in the measurements from the sensors to weight the optimization processes. Our calibration technique has the advantage of being accurate and easy to perform for the user, avoiding the need of specific calibration patterns. The method has been extensively tested in simulation and in real case experiments, yielding results that validate the claimed features. Future work will focus on reducing the errors caused by the assumption that the selected corner is ideal by considering the orientation of the plane normals as additional variables in the optimization process.

ACKNOWLEDGMENTS

We would like to thank to Francisco Vasconcelos for providing us the source code of his calibration method.

REFERENCES

- [1] C. Früh and A. Zakhor, "An automated method for large-scale, ground-based city model acquisition," *International Journal of Computer Vision*, vol. 60, no. 1, pp. 5–24, 2004.
- [2] J.-L. Blanco, F.-A. Moreno, and J. Gonzalez, "A collection of outdoor robotic datasets with centimeter-accuracy ground truth," *Autonomous Robots*, vol. 27, no. 4, pp. 327–351, 2009.
- [3] C. Premebida and U. J. C. Nunes, "Fusing lidar, camera and semantic information: A context-based approach for pedestrian detection," *The International Journal of Robotics Research*, vol. 32, no. 3, pp. 371–384, 2013.
- [4] B. Douillard, D. Fox, F. Ramos, and H. Durrant-Whyte, "Classification and semantic mapping of urban environments," *The international journal of robotics research*, vol. 30, no. 1, pp. 5–32, 2011.
- [5] E. Fernández-Moral, J. González-Jiménez, P. Rives, and V. Arévalo, "Extrinsic calibration of a set of range cameras in 5 seconds without pattern," in *Intelligent Robots and Systems (IROS), in 2014 IEEE International Conference on*, IEEE, 2014.
- [6] E. Fernández-Moral, V. Arévalo, and J. González-Jiménez, "Extrinsic calibration of a set of 2d laser rangefinders," in *International Conference on Robotics and Automation (ICRA 2015)*, IEEE.
- [7] S. Wasielewski and O. Strauss, "Calibration of a multi-sensor system laser rangefinder/camera," in *Intelligent Vehicles' 95 Symposium, Proceedings of the*, pp. 472–477, IEEE, 1995.
- [8] K. Kwak, D. F. Huber, H. Badino, and T. Kanade, "Extrinsic calibration of a single line scanning lidar and a camera," in *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, pp. 3283–3289, IEEE, 2011.
- [9] Q. Zhang and R. Pless, "Extrinsic calibration of a camera and laser range finder (improves camera calibration)," in *Intelligent Robots and Systems, 2004.(IROS 2004). Proceedings. 2004 IEEE/RSJ International Conference on*, vol. 3, pp. 2301–2306, IEEE, 2004.
- [10] Z. Zhang, "Flexible camera calibration by viewing a plane from unknown orientations," in *Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on*, vol. 1, pp. 666–673, IEEE, 1999.
- [11] F. Vasconcelos, J. P. Barreto, and U. Nunes, "A minimal solution for the extrinsic calibration of a camera and a laser-rangefinder," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 34, no. 11, pp. 2097–2107, 2012.
- [12] L. Zhou, "A new minimal solution for the extrinsic calibration of a 2d lidar and a camera using three plane-line correspondences," *IEEE Sensors Journal*, vol. 14, no. 2, pp. 442–454, 2014.
- [13] R. G. Von Gioi, J. Jakubowicz, J.-M. Morel, and G. Randall, "Lsd: A fast line segment detector with a false detection control," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 4, pp. 722–732, 2010.
- [14] R. Hartley and A. Zisserman, *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- [15] J.-L. Blanco, "A tutorial on se(3) transformation parameterizations and on-manifold optimization," tech. rep., University of Malaga, Sept. 2010.
- [16] V. Nguyen, A. Martinelli, N. Tomatis, and R. Siegwart, "A comparison of line extraction algorithms using 2d laser rangefinder for indoor mobile robotics," in *Intelligent Robots and Systems, 2005.(IROS 2005). 2005 IEEE/RSJ International Conference on*, pp. 1929–1934, IEEE, 2005.
- [17] J. C. Clarke, "Modelling uncertainty: A primer," *Tutorial of Department of Eng. Science*, pp. 1–21, 1998.
- [18] O. Naroditsky, A. Patterson, and K. Daniilidis, "Automatic alignment of a camera with a line scan lidar system," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, pp. 3429–3434, IEEE, 2011.
- [19] H. Yang, X. Liu, and I. Patras, "A simple and effective extrinsic calibration method of a camera and a single line scanning lidar," in *Pattern Recognition (ICPR), 2012 21st International Conference on*, pp. 1439–1442, IEEE, 2012.