

Extrinsic Self Calibration of a Camera and a 3D Laser Range Finder from Natural Scenes

Davide Scaramuzza, Ahad Harati, and Roland Siegwart

Abstract—In this paper, we describe a new approach for the extrinsic calibration of a camera with a 3D laser range finder, that can be done on the fly. This approach does not require any calibration object. Only few point correspondences are used, which are manually selected by the user from a scene viewed by the two sensors. The proposed method relies on a novel technique to visualize the range information obtained from a 3D laser scanner. This technique converts the visually ambiguous 3D range information into a 2D map where natural features of a scene are highlighted. We show that by enhancing the features the user can easily find the corresponding points of the camera image points. Therefore, visually identifying laser-camera correspondences becomes as easy as image pairing. Once point correspondences are given, extrinsic calibration is done using the well-known PnP algorithm followed by a non-linear refinement process. We show the performance of our approach through experimental results. In these experiments, we will use an omnidirectional camera. The implication of this method is important because it brings 3D computer vision systems out of the laboratory and into practical use.

I. INTRODUCTION

One of the basic issues of mobile robotics is the automatic mapping of the environments. Autonomous mobile robots equipped with 3D laser range finders are well suited for this task. Recently, several techniques for acquiring three-dimensional data with 2D range scanners installed on a mobile robot have been developed (see [1], [2], and [3]). However, to create realistic virtual models, visually-perceived information from the environment has to be acquired and it has to be precisely mapped onto the range information. To accomplish this task, camera and 3D laser range finder must be extrinsically calibrated, that is, the rigid transformation between the two reference systems must be estimated. Most previous works on extrinsic laser-camera calibration concern calibration of perspective cameras to 2D laser scanners (see [4], [5], and [6]). In contrast to previous works, in this paper we consider the extrinsic calibration of a general camera with a three-dimensional laser range finder. Because of the recent development of 3D laser scanners, only little work about extrinsic calibration of camera and 3D scanners

This work was conducted within the EU Integrated Projects COGNIRON (“The Cognitive Robot Companion”) and BACS (“Bayesian Approach to Cognitive Systems”). It was funded by the European Commission Division FP6-IST Future and Emerging Technologies under the contracts FP6-IST-002020 and FP6-IST-027140 respectively.

D. Scaramuzza, IEEE member, is a PhD student at the Autonomous Systems Laboratory (ASL) at the Swiss Federal Institute of Technology Zurich (ETH), Switzerland. davide.scaramuzza@ieee.org

A. Harati, is a PhD student at the ASL-ETH, Zurich, Switzerland. ahad.harati@mavt.ethz.ch

R. Siegwart, IEEE member, is Full Professor at the ETH Zurich and director of the ASL. r.siegwart@ieee.org

exists. Furthermore, the process of external calibration is often poorly documented. This process usually requires some modification of the scene by introducing landmarks that are visible by both the camera and the laser.

Well-documented work about extrinsic calibration of camera and 3D scanners can be found in [7] and [8]. However, in [7], the authors deal with the case of visible laser traces. Conversely, for the case of the invisible laser in [8], the authors propose a method for fast extrinsic calibration of a camera and a 3D scanner which makes use of a checkerboard calibration target, like the one commonly used for the internal calibration of a camera. Furthermore, they provide a useful laser-camera calibration toolbox for Matlab that implements the proposed calibration procedure [9]. Their method requires the user to collect a few laser-camera acquisitions where the calibration grid is shown at different positions and orientations. This technique however needs several camera-laser acquisitions of the grid for a sufficiently accurate external calibration of the system.

The work described in this paper also focuses on the extrinsic calibration of a camera and a 3D laser range finder but the primary difference is that we do not use any calibration pattern. We use only the point correspondences that the user hand selects from a single laser-camera acquisition of a natural scene. As we use no calibration target, we name our technique self-calibration. This work was motivated while working in the First European Land-Robot Trial [10]. In that contest, we presented an autonomous Smart car equipped with several 3D laser range finders and cameras (both omnidirectional and perspective cameras). The goal was to produce 3D maps of the environment along with textures [11]. Especially when working in outdoor environments, doing several laser-camera acquisitions of a calibration pattern can be a laborious task. For each acquisition, the pattern has to be moved to another position and this process usually takes time. Furthermore, weather conditions (e.g. wind, fog, low visibility) can sometimes perturb or even alter the calibration settings. Hence, the calibration must be done quickly. Because of this, we developed the procedure presented in this paper. The advantages are that now we need only a single laser-camera acquisition and that the calibration input are point correspondences manually selected from the laser-camera acquisition of a natural scene. Once point correspondences are given, the extrinsic calibration problem becomes a camera pose-estimation problem which is well known in computer vision and can be solved using standard methods. The difficulty resides in visually identifying the point correspondences because range images in general lack




Fig. 1. (a) Pin-hole model used for perspective cameras. (b) Image formation model for central catadioptric cameras. (c) Unified spherical projection model for central omnidirectional cameras: every pixel in the sensed image measures the irradiance of the light passing through the effective viewpoint in one particular direction. The vectors are normalized to 1.

in point features. To bypass this problem, we process the range data so that we can highlight discontinuities and orientation changes along specific directions. This processing transforms the range image into a new image that we call Bearing Angle image (BA). Using BA images, we will show that visually identifying point correspondences between the laser and the camera outputs becomes as easy as image pairing. To show the generality of the methodology, in our experiments we will use an omnidirectional camera. The BA images and the application of the method to an omnidirectional camera are the two main contributions of this paper.

This document is organized as follows. Section II describes the projection model of the camera and 3D laser scanner. Section III defines the concept of BA images and explains how to compute them. Section IV describes the calibration procedure. Finally, section V presents some calibration results.

II. LASER-CAMERA PROJECTION MODEL

A. Camera Model

In this work, we deal with central cameras either perspective or omnidirectional. Central cameras satisfy the single effective viewpoint property, that is, they have a single centre of projection (see Fig. 1). For the case of catadioptric omnidirectional cameras, such property can be achieved using hyperbolic, parabolic, or elliptical mirrors [12]. In the last years, central omnidirectional cameras using fisheye lenses have also been built [13].

We assume that the camera has already been calibrated. Some Matlab toolbox to quickly calibrate these sensors can be found in [15], [16], and [17].

Assuming the camera is already calibrated, given a pixel point (u, v) on the camera image plane, we can recover the orientation of the vector \mathbf{X} emanating from the effective viewpoint to the corresponding 3D point (1). Conversely, given a 3D point λX , we can reproject it onto the camera image plane (u, v) (2):

$$\lambda X = \lambda[x, y, z]^T = F(u, v) \quad (1)$$

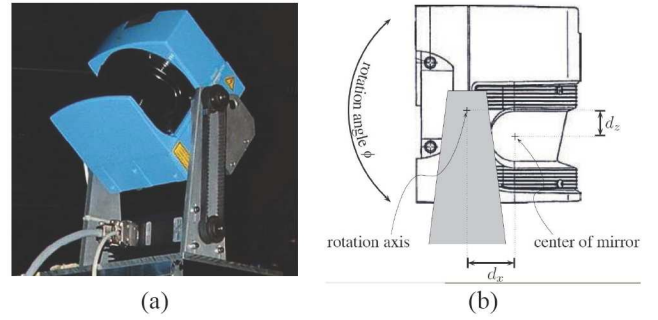


Fig. 2. (a) Our custom-built 3D scanner is composed of a SICK LMS 200 laser range finder mounted on a rotating support. (b) Schematic of the sensor used for calibration.

$$[u, v]^T = F^{-1}(X) \quad (2)$$

where λ is the depth factor and $\|\mathbf{X}\| = 1$. Function F depends on the camera used. Some proposed formulations for F can be found in [13], [14], and [18]. We assume that the origin of the camera coordinate system coincides with the single effective viewpoint. This corresponds to the optical center for perspective cameras and the internal focus of the mirror, in the catadioptric case. The x-y plane is orthogonal to the mirror axis for the catadioptric case or to the camera optical axis for the perspective case (see more in [5], [14], and [18]). Without loss of generality, in this paper, we consider omnidirectional cameras, but the same considerations also apply to perspective cameras. According to what we have mentioned so far, in the remainder of the document we assume that for every sensed pixel we know the orientation of the correspondent vector \mathbf{X} on the unit sphere centered in the mirror frame (Fig. 1.c)

B. Laser Model

3D laser range finders are usually built by nodding or rotating a 2D scanner in a stepwise or continuous manner around its lateral or radial axis. Combining the rotation of the mirror inside the 2D scanner with the external rotation of the scanner itself, spherical coordinates of the measured points are obtained. However, since in reality it is impossible to adjust the two centers of rotation exactly on the same point, the measured parameters are not spherical coordinates and offset values exist. These offset values have to be estimated by calibrating the 3D sensor by considering its observation model. The approach presented in this paper for the extrinsic calibration of a 3D laser scanner with a camera is general and does not depend on the sensor model. Therefore, we assume the laser is already calibrated. Nevertheless, we explain here the scanner model used in our experiments, but a different sensor setup can also be used along with its corresponding observation model.

The 3D range sensor used in this work is a custom-built 3D scanner (Fig. 2). It is composed of a two-dimensional SICK laser scanner mounted on a rotating support, which is driven by a Nanotec stepping motor. The sensor model can be written as:

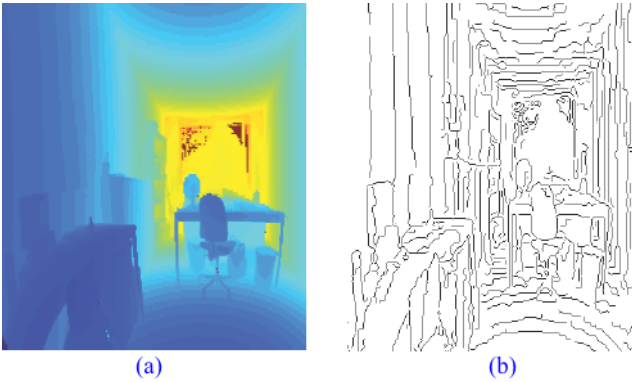


Fig. 3. (a) Depth image of a scene. Jet colormap has been used. The color shade (from blue to red) is proportional to the depth. (b) Sobel based edge map of the depth image.

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} c_i c_j & -c_i s_j & s_i & c_i d_x + s_i d_z \\ s_j & c_j & 0 & 0 \\ -s_i c_j & s_i s_j & c_i & -s_i d_x + c_i d_z \end{bmatrix} \begin{bmatrix} \rho_{ij} \\ 0 \\ 0 \\ 1 \end{bmatrix}$$

$$c_i = \cos(\varphi), c_j = \cos(\theta), s_i = \sin(\varphi), s_j = \sin(\theta) \quad (3)$$

where ρ_{ij} is the j -th measured distance with corresponding orientation θ_j in the i -th scan plane, which makes the angle φ_j with the horizontal plane (Fig. 2.b). The offset of the external rotation axis from the center of the mirror in the laser frame has components d_x and d_z (as observed in Fig. 2.b). $[x, y, z]^T$ are the coordinates of each measured point relative to the global frame (with its origin at the center of the rotation axis, the x -axis pointing forward and the z -axis toward the top). The sensor is calibrated as discussed in [23] based on a known ground truth.

III. BEARING ANGLE IMAGES

In this subsection, we will describe how to highlight depth discontinuities and direction changes in the range image so that the user can easily find the corresponding points of the camera image points. Such features in the range image are called image details. Fig. 3.a shows the range image of an office-like environment extracted by our 3D scanner. In such an environment, we would like emphasizing key points like corners arising from the plane intersections of walls, tables, chairs, and other similar discontinuities. Fig. 3.b shows the result of directly applying a Sobel edge detector on the range image. As observed, edge detection does not directly help for our task, since edges are zones of the range image where the depth between two adjacent points significantly changes. In fact, many details in the range image do not create a big jump in the measured distance. Such edges are called "roof edges" and correspond to sharp direction changes (e.g. tetrahedron shaped corners). Therefore, a measure of direction should be used to highlight all the desired details in the scene. As representative of the surface direction, its corresponding normal vector is usually used (see [19] and [20]). Surface normal vectors are estimated based on the




Fig. 4. (a) Bearing Angles computed along a given scan plan. (b) Plot of a horizontal BA signal.

neighborhood of each point. However, for our application, we avoid the use of surface normals as representatives of the direction. The reason is that we want to highlight the details of the surface along some specific directions (e.g vertical, horizontal, and diagonal). We will show that treating each dimension separately leads to enhanced estimation of the image details. Let the range data coming from the 3D scanner be arranged in the form of a 2D matrix where its entries are ordered according to the direction of the laser beam. This matrix will be referred to as a depth matrix. We compute the surface orientation along four separate directions of the depth matrix, namely the horizontal, vertical, and diagonal one (the latter having $+45^\circ$ and -45° orientation).

We define Bearing Angle (BA) the angle between the laser beam and the segment joining two consecutive measurement points (see Fig. 4.a). This angle is calculated for each point in the depth matrix along the four defined directions (that we call also "traces"). More formally:

$$BA_i = \arccos \frac{\rho_i - \rho_{i-1} \cos d\varphi}{\sqrt{\rho_i^2 - \rho_{i-1}^2 - \rho_i \rho_{i-1} \cos d\varphi}} \quad (4)$$

where ρ_i is the i -th depth value in the selected trace of the depth matrix and $d\varphi$ is the corresponding angle increment (laser beam angular step in the direction of the trace). Performing this calculation for all points in the depth matrix will lead to an image which is referred to as a BA image. BA images can be calculated from the depth matrix along any direction, to highlight the details of the scene in the selected direction. In our application, horizontal, vertical, and diagonal traces suffice for a successful enhancement of the details of the scene (Fig. 5). However, any other direction could be also considered depending on the application. As observed in Fig. 5, these angular measures show the geometry of the scene by highlighting many details that were not distinguishable in the range image (Fig. 3.a). Hence, these will be used in the next section for extracting corresponding features.

IV. EXTRINSIC LASER-CAMERA CALIBRATION

A. Data Collection

Our calibration technique needs a single acquisition of both the laser scanner and the omnidirectional camera. The acquisition target can be any natural scene with a sufficient number of distinguishable key points, i.e. roof edges or depth discontinuities. Our calibration procedure consists of three stages: firstly, we compute the BA images of the




Fig. 5. BA images for a real scan (top left: vertical, top right: horizontal, and bottom: two diagonal directions). Observe that, in BA images, the scene details are very highlighted (e.g. even the corners of the picture hanging on the left wall are now well distinguishable; in the range image 3.a, they were not). In this pictures, jet colormap has been used. The color shade (from blue to red) is proportional to the BA value.

acquired range image. Secondly, we manually select several point correspondences (at least four) between the BA image and intensity image. Finally, extrinsic calibration is done using a camera pose estimation algorithm followed by a non-linear refinement process. Observe that usually not all four BA images are needed. Depending on the scene and the orientation of the laser scanner to the scene, only the horizontal BA image could suffice. However, the remainder BA images can be used anyway, to check if there are further details that would be worth exploiting. At the end of the visual correspondence pairing, we have n laser points in the laser frame and their correspondent points on the camera image plane. We rewrite these points in the following way:

$$\begin{aligned} \theta_C &= [\theta_{C,1}, \theta_{C,2}, \dots, \theta_{C,n}], \theta_L = [\theta_{L,1}, \theta_{L,2}, \dots, \theta_{L,n}], \\ d_L &= [d_{L,1}, d_{L,2}, \dots, d_{L,n}] \end{aligned} \quad (5)$$

where θ_C and θ_L are the unit norm orientation vectors of the camera and laser points in their respective reference frames, and d_L are the point distances in the laser frame.

B. Extrinsic Calibration

Extrinsic calibration of a camera and a 3D laser range finder consists in finding the rotation \mathbf{R} and translation \mathbf{T} between the laser frame and the camera frame that minimizes a certain error function. In photogrammetry, the function to minimize is usually the reprojection error:

$$\min_{\mathbf{R}, \mathbf{T}} \frac{1}{2} \sum_{i=1}^n \|m_i - \hat{m}(R, T, p_i)\|^2 \quad (6)$$

where $\hat{m}(R, T, p_i)$ is the reprojection onto the image plane of the laser point p_i according to equation (2). However, the reprojection error is not theoretically optimal in our




Fig. 6. Estimation of the translation (meters) versus the number of selected points.

application because the resolution of the camera is not uniform. A better error function uses the Riemann metric associated to a sphere as it takes into account the spatial distribution (see [5] and [8]). This metric minimizes the difference of the bearing angles of the camera points and the bearing angles of the laser points after reprojection onto the image, that is:

$$\min_{\mathbf{R}, \mathbf{T}} \frac{1}{2} \sum_{i=1}^n \|\arccos(\theta_{C,i}^T \cdot \theta_{CL,i})\|^2 \quad (7)$$

where θ_{CL} is the unit norm orientation vector of $\hat{m}(R, T, p_i)$. According to equation (2), each correspondence pair contributes two equations. In total, there are $2 \times n$ equations in 6 unknowns. Hence, at least 3 point associations are needed to solve for \mathbf{R} and \mathbf{T} . However, 3 point associations yield up to four solutions and thus a fourth correspondence is needed to remove the ambiguity. This problem has already been theoretically investigated for a long time and is well known in the computer vision community as PnP problem (Perspective from n Points). Some solutions to this problem can be found in [21] and [22]. We implemented the PnP algorithm described in [21] to solve the calibration problem. The output of the PnP algorithm are the depth factors of the camera points in the camera reference frame, that is $d_C = [d_{C,1}, d_{C,2}, \dots, d_{C,n}]$. Then, to recover the rigid transformation between the two point sets, namely \mathbf{R} and \mathbf{T} , we used the motion estimation algorithm proposed by Zhang [24].

C. Non-Linear Optimization

The drawback of using the PnP algorithm is that the solution is quite sensitive to the position of the input points, which were manually selected, and also to noisy range information. Furthermore, we have to take into account that the two sensors can have different resolution and that the rigid transformation is recovered by linear least-square estimation. Thus, the solution given in section B is suboptimal. To refine the solution, we minimized (7) as a non-linear optimization problem by using the Levenberg-Marquardt algorithm. This requires an initial guess of \mathbf{R} and \mathbf{T} , which is obtained using the method described in subsection B.




Fig. 7. Estimation of the rotation (roll, pitch, and yaw angles) versus the number of selected points (the x-axis ranges from 4 to 10).

TABLE I
PARAMETER ESTIMATION

$T(m)$	σ	$R(deg)$	σ	$Pixel\ error$	σ
0.207	0.05	0.64	0.21		
0.042	0.017	-1.24	0.85	1.6	1.2
0.139	0.005	166.95	1.08		

V. RESULTS

The proposed method has been tested on the custom-built rotating scanner described in section II.B and on an omnidirectional camera consisting of a KAIDAN 360 One VR hyperbolic mirror and a SONY XCD-SX910-CR digital camera. The camera resolution was set as 640×480 pixels. The rotating scanner provided 360° field of view range measurements with a vertical angular resolution of 1° and a horizontal resolution of 0.5° . The calibration of the omnidirectional camera was done using a toolbox available on the Internet [16]. The calibration of the rotating scanner was done using a known ground truth as explained in [23]. We evaluated the robustness of the proposed approach with respect to the number of manually selected points. We varied the number of laser-camera correspondences from 4 to 10 and for each combination we did ten calibration trials using different input points. The results shown in Fig. 6 and 7 are the average. Observe that after selecting more than 5 points, the values of the estimated \mathbf{R} and \mathbf{T} become rather stable. This stability occurs when points are chosen uniformly from the entire scene viewed by the sensors. Conversely, when points are selected within local regions of the scene, the estimated extrinsic parameters are biased by the position of this region. We also tried to use more than 10 points but the estimated parameters did not deviate from the average values that had been already estimated. The estimated \mathbf{R} and \mathbf{T} were in agreement with the hand-measured values. Furthermore, the estimated parameters were stable against the position of the input points when these input points were picked uniformly from all around the scene. In table I, the mean and the standard deviation of \mathbf{R} (roll, pitch, and yaw angles) and \mathbf{T} (T_x, T_y, T_z) are shown for the case of ten correspondence pairs. The results were averaged among ten different calibration trials. In table I, we also show the reprojection error (in pixel). For data fusion, this error is the most important. It measures the distance between the laser




Fig. 8. (a) A detail of the BA image. For visualization, we used a gray scale where the intensity is proportional to the BA value. (b) Result of a Sobel edge detector on the BA image. (c) The edges are reprojected onto the image using the computed \mathbf{R} and \mathbf{T} .

points reprojected onto the image using the estimated \mathbf{R} and \mathbf{T} , and the image points. In our experiments, the average reprojection error was 1.6 pixels and the standard deviation was 1.2 pixels. The reprojection of the laser points onto the image also offers an indirect way to evaluate the quality of the calibration. To do this, we chose not to reproject all the laser points onto the image. Rather, we reprojected only those laser points that represent discontinuities in the range image. To select only depth discontinuities automatically, we applied an edge detector to the BA image. The edge points, as representative of depth discontinuities, were then reprojected onto the image. The reprojection results are shown in Fig. 8. As observed, the laser edge points well reprojected onto the edges of the intensity image. In the end, using the estimated \mathbf{R} and \mathbf{T} , we colored an entire 3D scan by reprojecting the scan onto the corresponding image. The results of this color mapping are shown in Fig. 9.

VI. CONCLUSIONS

In this paper, we presented a new approach for the extrinsic calibration of a camera and a 3D laser range finder, that can be done on the fly. The method uses only a few correspondent points that are manually selected by the user from a single laser-camera acquisition of a natural scene. Our method relies on a novel technique to visualize the range information. This technique converts the visually ambiguous 3D range information into a 2D map (called BA image) where natural features of a scene are highlighted. In this way, finding laser-camera correspondences is facilitated. Once correspondence pairs have been given, calibration is done using the PnP algorithm followed by a non-linear refinement process. Real experiments have been conducted using an




Fig. 9. (a) A panoramic picture of a scene unwrapped into a cylindrical image. The size of the original omnidirectional image was set as 640×480 pixels. After extrinsic calibration, the color information was mapped onto the 3D points extracted from a rotating SICK laser range finder. (b), (c), and (d) show the results of this color mapping. The colors are well reprojected onto the 3D cloud.

omnidirectional camera and a rotating scanner, but the same approach can be also applied to any other type of camera (e.g. perspective) or laser range finder. The results showed that selecting the input points uniformly from the whole scene, robust calibration can be done by using only from eight to ten correspondence pairs. The BA images and the application of the method to an omnidirectional camera were the two main contributions of this paper. The implication of the proposed calibration approach is important because it brings 3D computer vision systems out of the laboratory and into practical use. In fact, the proposed approach requires no special equipment and allows the user to calibrate quickly the system in those cases where special settings are difficult to be arranged.

VII. ACKNOWLEDGMENTS

The authors gratefully acknowledge Jan Weingarten for providing the range images from the 3D laser range finder [23].

REFERENCES

- [1] D. Haehnel, W. Burgard, and S. Thrun. Learning compact 3d models of indoor and outdoor environments with a mobile robot. *Robotics and Autonomous Systems*, 44(1):15-27, (2003).
- [2] Weingarten, J. and Siegwart, R. 3D SLAM using Planar Segments. In *Proceedings of IROS'06, Beijing, October 9-15*, (2006).
- [3] O. Wulf, K-A. Arras, H.I. Christensen, and B. Wagner. 2D mapping of cluttered indoor environments by means of 3D perception. In *Proceedings of ICRA'04*, pages 4204-4209, New Orleans, April, (2004.).
- [4] Q. Zhang and R. Pless. Extrinsic Calibration of a Camera and Laser Range Finder (improves camera intrinsic calibration), in *Proc. of IEEE IROS'04*, (2004).
- [5] C. Mei and P. Rives. Calibration between a Central Catadioptric Camera and a Laser Range Finder for Robotic Applications, In *Proceedings of ICRA'06*, Orlando, May, (2006).
- [6] S. Wasielewski and O. Strauss, Calibration of a multi-sensor system laser rangefinder/camera, *Proc. of the Intelligent Vehicles' Symposium*, pp. 472-477, (1995).
- [7] D. Cobzas, H. Zhang and M. Jaegersand, A comparative analysis of geometric and image-based volumetric and intensity data registration algorithms, In *Proc. of IEEE ICRA'02*, (2002).
- [8] R. Unnikrishnan and M. Hebert. Fast Extrinsic Calibration of a Laser Rangefinder to a Camera. *Technical report, CMU-RI-TR-05-09*, Robotics Institute, Carnegie Mellon University, July, (2005).
- [9] R. Unnikrishnan, Laser-Camera Calibration Toolbox for Matlab: <http://www.cs.cmu.edu/ranjith/lcct.html>
- [10] <http://www.m-elrob.eu/>
- [11] P. Lamon, C. Stachniss, R. Triebel, P. Pfaff, C. Plagemann, G. Grisetti, S. Kolsky, W. Burgard, and R. Siegwart. Mapping with an autonomous car. In *IEEE/RSJ IROS'06 Workshop: Safe Navigation in Open and Dynamic Environments*, Beijing, China, (2006).
- [12] Baker, S. and Nayar, S.K. A theory of catadioptric image formation. In *Proc. of the 6th ICCV*, pp. 35-42, ISBN 81-7319-221-9, India, January 1998, IEEE Computer Society, Bombay, (1998).
- [13] Micusik, B. and Pajdla, T. Estimation of omnidirectional camera model from epipolar geometry. In *Proc. of CVPR*. ISBN 0-7695-1900-8, US, June 2003, IEEE Computer Society, Madison, (2003).
- [14] Scaramuzza, D., Martinelli, A. and Siegwart, R. A Toolbox for Easy calibrating Omnidirectional Cameras. In *Proc. of IROS'06*, pp. 5695-5701, China, October 2006, Beijing, (2006).
- [15] Bouguet, J.Y. Camera Calibration Toolbox for perspective cameras: <http://www.vision.caltech.edu/bouguetj/calib4oc/>
- [16] Scaramuzza, D. Matlab calibration toolbox for omnidirectional cameras. Google for "ocamcalib".
- [17] Mei, C. Matlab calibration toolbox for omnidirectional cameras. <http://www.robots.ox.ac.uk/cmei/Toolbox.html>
- [18] Zhang, Z. A Flexible New Technique for Camera Calibration. *IEEE Transactions on PAMI*. Volume 22, No. 11, November 2000, ISSN 0162-8828, (2000).
- [19] Pulli, K. Vision methods for an autonomous machine based on range imaging. *Master's thesis*. ACTA Universitatis Ouluensis, C 72, (1993).
- [20] Bellon, O.R.P. and L. Silva. New improvements to range image segmentation by edge detection. *Signal Processing Letters, IEEE*, 9(2), 43-45, (2002).
- [21] Quan L. and Lan, Z., Linear N-point Camera Pose Determination, *IEEE Transaction on PAMI*, 21(7), (1999).
- [22] Gao, X.S., Hou, X., Tang, J., and Chen, H. Complete Solution Classification for the Perspective-Three-Point Problem, *IEEE Transactions on PAMI*, 930-943, 25(8), (2003).
- [23] Weingarten, J. Feature-based 3D SLAM. *PhD Thesis, Swiss Federal Institute of Technology Lausanne, EPFL*, no 3601, Dir. Roland Siegwart, (2006).
- [24] Zhang, Z., Iterative Point Matching for Registration of Free-Form Curves, *IRA Rapports de Recherche No 1658 Programme 4 Robotique, Image et Vision*, (1992).