

Infrared Camera Calibration for Dense Depth Map Construction

Michael Gschwandtner, Roland Kwitt, Andreas Uhl
Department of Computer Sciences
University of Salzburg, Austria
{mgschwan, rk Witt, uhl}@cosy.sbg.ac.at

Wolfgang Pree
Software & Systems Research Center (SRC)
University of Salzburg, Austria
wolfgang.pree@cs.uni-salzburg.at

Abstract—In this paper, we introduce a novel and cost effective approach to calibrate the geometric properties of a far-infrared (IR) sensor. We further demonstrate that fully automatic sensor-to-sensor calibration is feasible in a setup involving a laser range scanner, IR cameras as well as conventional cameras. The calibration result then serves as a basis for upsampling range measurements to the resolution of the IR or visible-light camera images. Since our approach allows to rely on IR information instead of visible-light information for upsampling, bad light conditions or even no visible light at all are no limitation. From a practical point of view, we only require one calibration board of relatively small size which facilitates application in outdoor environments and further allows seamless integration of the IR camera in an existing multi-sensor platform. Our experimental results demonstrate that IR images are particularly useful to obtain reasonable depth information for living objects, when visible-light cameras are either blind or require impractical exposure times. In fact, our approach provides a convenient solution to IR camera calibration and integration, an issue which is particularly important in scenarios where sensors are not permanently mounted on vehicles and consequently require on-site adjustment and calibration.

I. MOTIVATION

With the ongoing improvement in vision sensor technology, autonomous driving vehicles have become a hot topic in recent years. Milestones in this area have been set in the DARPA Grand Challenge (2005) and the DARPA Urban Challenge (2007). In both competitions, the competitors made excessive use of sensor equipment, including laser range scanners, visible-light cameras or even radar sensors.

In the context of situations with bad light conditions, IR cameras are a useful asset in a multi-sensor system. IR cameras do not rely on visible light but capture IR radiation instead. Given that the ultimate objective of any autonomously driving vehicle is to avoid harming any living objects (humans or animals), IR cameras are of potential interest, since living objects will show up very decisively.

In order to fuse information from multiple sensors, e.g. cameras or laser range scanners, it is essential that the data is represented in a reference coordinate frame common to all sensors. This requires knowledge of intrinsic [1] and extrinsic parameters of every sensor in the system. Most previous works on calibrating a multi-sensor system focus on the case of estimating the rigid transformation between the reference frame of a visible-light camera and the reference frame of a 2-D laser range scanner [2], [3] (assuming that the intrinsic camera parameters are already known). With the emergence

of 3-D laser range scanners, including these sensors in multi-sensor platforms has become popular. In the DARPA Urban Challenge for instance, many vehicles were equipped with a 3-D laser range scanner [4]. In [5], the authors consider the case of calibrating a visible-light camera (in this special case an already calibrated omnidirectional camera) to a 3-D laser range scanner by using hand-selected point correspondences from just one natural scene. In [6], a MATLAB calibration toolbox is presented to determine the external calibration parameters of a 3-D laser range scanner and a visible-light camera. The approach is capable of estimating the camera intrinsics as well, however requires manual selection of the laser returns off a planar calibration board. Furthermore, it is required that the range measurements are not very noisy, a requirement which can probably not be met by a fast rotating 3-D laser range scanner. In situations where the calibration process has to be performed outside the laboratory, we might have to cope with a substantial amount of outliers as well (due to reflections).

For IR camera calibration we cannot directly exploit existing approaches, since the widely-used black/white checkerboard pattern is not visible to an IR camera. However, using separate calibration boards for the intrinsic calibration of the IR and visible-light camera is an impractical solution and complicates extrinsic calibration between the cameras. To facilitate seamless integration of the IR camera, we present an fully automatic approach to determine the intrinsic parameters of the IR and visible-light cameras, as well as the extrinsic parameters of a multi-sensor system using just one planar calibration board. Relying on the automatic calibration parameters, we then demonstrate that IR images provide useful information to upsample range measurements in situations where visible-light cameras cease to provide useful data.

II. SENSOR CALIBRATION

The general approach towards estimation of intrinsic and extrinsic parameters of any visible-light camera is to use some sort of calibration pattern with known dimensions to establish object to pixel correspondences [1]. Usually, this is accomplished using a black/white checkerboard pattern mounted on a planar calibration board. The pattern is then shown to the camera at varying positions. Detecting the corners of each of the squares then allows to estimate the intrinsic parameters as well as the geometric distortion introduced by the lens. Given the intrinsic parameters are

available, it is then possible to determine the extrinsic parameters, defining the position and orientation of the camera reference frame with respect to a known world coordinate frame.

Calibrating multiple sensors to each other, requires that we use a calibration pattern visible to all sensors, though. In our setup, we thus need a pattern visible to a 3-D laser range scanner, a forward-looking IR (FLIR) camera and a visible-light camera. Due to extensive existing work [7], [8], [5] on the calibration of visible-light cameras, our strategy is to keep the black/white checkerboard pattern mounted on a planar surface and to augment it by a set of low-cost electrical components to facilitate IR camera calibration.

The goal is to provide an *easy-to-use* calibration mechanism that can be used outdoors, minimizing the constraints on the scenery. It is required that on-site changes to the test vehicle can be dealt with immediately rather than having to calibrate the system in a laboratory environment.

A. Lens Distortion

To calculate accurate and robust radial (and tangential) distortion parameters, we would have to have calibration points throughout the whole image (especially points close to the image borders). This could be achieved by i) moving the board close to the camera or ii) using a large calibration board. Unfortunately, moving the board close to the camera leads to unusable laser returns since a 3-D range scanner often only provides reliable range measurements for objects at a distance of at least six feet (see Section III). Consequently, it would not be possible to use the same set of calibration images for intrinsic and extrinsic calibration. We would inevitably need two separate sets: one, which is solely used for the calibration of intrinsic camera parameters and one, which is used for the calibration of extrinsic camera parameters.

Using a large calibration board as a potential solution to this problem is quite impractical, though, especially in outdoor environments or in situations where the sensor equipment is not permanently mounted on the vehicle and requires frequent re-calibration.

In our setup, visual assessment of the camera distortion has shown that there is hardly any distortion in the area of interest, which is at least six feet in front of each camera and relatively narrow as well. For that reason, we decided to omit estimation of the lens distortion, however, we emphasize that this step might have to be incorporated depending on the camera type. In that case, we favor the solution of using two separate image sets, due to lower practical impairment.

B. FLIR Camera Calibration

As we already mentioned, we choose a black/white checkerboard pattern mounted on a planar (metal) surface as a basis for calibration. This enables easy estimation of the intrinsic parameters of the visible-light camera by using the method of Zhang [7]. In order to use the same approach to determine the intrinsic IR camera parameters, we need to




Fig. 1: Physical configuration of the calibration board.

augment the calibration board by a pattern visible in the IR images.

Our approach to tackle this problem is to install a pattern of electrical elements emitting IR radiation. Using LEDs to create a detectable pattern in the visible and IR spectrum does not work because the thermal difference to the calibration board is too low, even though LEDs have an efficiency level of ≈ 20 percent leaving enough energy for thermal radiation. As an alternative strategy, we decided to use a set of resistors mounted in the centroid of each square (see Fig. 1a and Fig. 1c). This has the advantage that the geometric corners of the black/white squares are not physically distorted by mounting the resistors. Consequently, the corner points in the image of the visible-light camera can still be automatically detected (e.g. using the algorithm of Vezhnevets, implemented in the OpenCV library [9]) and calibration is not negatively affected.

Regarding the physical configuration of the calibration board, we install an evenly spaced array of 35 (7×5) resistors connected to a 12 volt DC power supply. Each resistor is a 330 ohm metal film resistor. The circuit diagram (see Fig. 1b) shows that there are seven rows attached in parallel to the power supply. Each row consists of 5 resistors connected serially. After adding all wires, the whole circuit draws ≈ 1.12 ampere and thus emits ≈ 13.44 watts of thermal radiation.

By looking at the two exemplary IR images showing the calibration board in Fig. 2a, we notice that simple gray-value thresholding to detect the resistors will not work reliably for the following reason: the person holding the calibration board emits much more IR radiation than the resistors which eventually leads to a broader range of intensity values. It is practically impossible to detect the resistor spots by means of gray-value thresholding alone. One could argue that this problem could be solved by using a (long) mounting bracket, however, such a solution is usually constrained to a laboratory environment.

The basic idea of our approach is to exploit the regularity in the spacing of the resistors (in both directions) to reliably detect the pattern. Our algorithm consists of two steps which are executed iteratively and one finalization step in which we pick the most suitable result and identify any remaining resistors. In one iteration step, we perform gray-value thresholding to identify resistor candidates and then measure the regularity of the candidate resistor pattern which




Fig. 2: Exemplary IR images showing calibration board and the 7×5 resistor pattern.

most likely corresponds to the true resistors arrangement.

a) Candidate Search: In the first step of one iteration, we perform gray-value thresholding with a threshold chosen from a predefined range R . Since the intensity in IR images can differ significantly depending on the environment, the idea is to increase the gray-value threshold $t \in [0, 1]$ in each iteration in order to find a setting where the resistors can at least be segmented from the metal calibration board. In that case, we expect a high regularity in the identified candidate resistors which will be measured in the second step. The range of thresholds $R := [a, b]$, $a := t_O - \epsilon$, $b := t_O + \epsilon$ is determined by Otsu's [10] intensity threshold t_O and a heuristically determined value ϵ^1 . After the IR image is binarized, all connected regions with an area larger than 50 pixel are removed and the region centroids are calculated. Since the calibration board has to be positioned at least six feet away from the sensor, any region larger than 50 pixel most likely corresponds to something other than a resistor. The centroids, denoted by the vertices $V_i \in \mathbb{R}^2$ represent the resistor candidates and are the input to the next step.

b) Measuring Regularity: In the second step of one iteration, we compute the Delaunay triangulation T of the set of vertices V_i . The Delaunay triangulation has the property that the minimum interior angle among all triangles is the greatest possible among all triangulations. Fig. 3 shows the part of the Delaunay triangulation which contains the $x \times y$ resistor pattern (illustrated as black dots). Let us consider the $2(x-1)(y-1)$ triangles which triangulate this pattern. We define that a triangle is a *neighbor* of another triangle if they share a coincident edge. We observe that the $L = 2(x-3)(y-3)$ inner triangles (shaded gray) share a common property: all three immediate neighbors of each of the L triangles as well as their immediate neighbors and the center triangle itself have approximately the same *perimeter* (Fig. 3). We explicitly focus on the L inner triangles, since considering all $2(x-1)(y-1)$ triangles would not allow to establish a rigorous neighborhood criterion. In case of the border triangles of the resistor pattern for instance, the criterion does not hold. Obviously, the inner triangles share other properties as well, but the perimeter led to the most robust detection results in our experiments.

The computational steps to identify the L inner triangles can be summarized as follows: For each of the triangles Δ_j in the triangulation T , we compute the maximum perimeter difference m_j to all its nine neighbors. Formally, given that

$p(\Delta_j)$ denotes the perimeter of triangle Δ_j and $\Delta_{jn}, n = 1, \dots, 9$ denote the nine neighbors, we calculate

$$m_j = \max_{n=1, \dots, 9} \{|p(\Delta_j) - p(\Delta_{jn})|\}. \quad (1)$$

In case a triangle has less than nine neighbors the maximum perimeter difference is set to infinity. Next, we determine the variance σ^2 of the L smallest values of m_1, \dots, m_M , where M denotes the total number of triangles in the triangulation T . This variance is a measure of regularity among the triangles with the smallest maximum perimeter difference to its nine neighbors. Let this set of L triangles be denoted as $\{\Delta_1, \dots, \Delta_L\}$. Now, given that $v(\Delta_j)$ returns the set of vertices that span Δ_j , we can require that the cardinality of the set

$$S = \bigcup_{j=1}^L v(\Delta_j) \quad (2)$$

is $|S| = (x-2)(y-2)$. Otherwise, the detection result cannot be valid, since we only have $(x-2)(y-2)$ (inner) resistors in our pattern.

c) Resistor Identification: Depending on the threshold stepsize s , we perform $\lceil (b-a)/s \rceil$ iterations of the *candidate search* and *regularity measurement* step. Letting the subscript i denote the iteration index, we eventually pick the configuration (t_i, S_i) where the variance σ_i^2 is minimal. The final steps to identify the remaining $K := xy - (x-2)(y-2)$ resistors are straightforward. We pick out the boundary vertices of the $(x-2)(y-2)$ identified vertices and determine all triangles attached to these vertices. The union of the triangle vertices then gives at least $K-4$ of the K remaining resistors. Depending on the triangulation T , we might however miss (at most) the four corner vertices. This can happen because the triangles containing a corner vertex might not be attached to any of the $(x-2)(y-2)$ vertices found by means of the regularity criterion. One of these triangles is illustrated in Fig. 3. The final four resistors are found by performing a nearest neighbor search around the intersection points of the edges determined by the outermost vertices of the detected pattern. Fig. 2b shows two final detection results.

Considering the fact that there is a boundary of metal around the resistors, the next resistor candidate outside the $x \times y$ pattern does not affect the robustness of our approach. Potential disturbances can however occur in case all resistor centroids cannot be identified by one threshold setting. This is obviously possible due to thermal differences of the resistors which can show up as considerable intensity variations.

C. The 3-D Laser Range Scanner

To establish a correspondence between the measurements of the 3-D laser range scanner and both camera systems, we need to automatically detect the 3-D coordinates of the calibration board corners. This is a non-trivial task, due to the inherent noise in the range measurements [11] (we use a Velodyne HDL-64E S2 scanner, see Section III). The problem can be split into two parts: First, we have to identify the point cluster corresponding to the range measurements

¹A value of $\epsilon = 0.25$ has shown to work well in practice.




Fig. 3: Illustration of an exemplary Delaunay triangulation as well as the neighborhood relations.

returned from the calibration board. Second, we have to fit a rectangle with the dimensions of the calibration board to the identified point cluster.

a) Point Cluster Identification: In order to reliably identify the 3-D point cluster corresponding to the range measurements from the calibration board, we physically simplify the task by using a (short) metal stick to hold the calibration board (fixed by means of a vacuum cup). This allows simple distance thresholding to separate the range measurements off the board from range measurements off the person holding the board. Otherwise, it would be significantly harder to automatically separate the point cluster of interest. We further note that the position where we can present the board to the laser scanner is constrained by the smallest field of view angle of one of the cameras, in both vertical and horizontal direction. Since we know the length of the stick, these constraints allow automatic identification of the point cluster. An exemplary segmentation result is shown in Fig. 5b, where the point cluster corresponding to the calibration board is shown in red.


b) Detecting the Corners: Regarding the second part of our problem, the identification of the corner coordinates, we have to cope with the aforementioned sensor noise and the fact that the board might be tilted as well. Since we want to allow outdoor calibration, we need a strategy that can handle possibly occurring outliers caused by reflections as well. In order to cope with these impairments, we rely on the well-known RANSAC algorithm [12] to fit a plane to the point cluster. The points are then projected onto the fitted plane along the viewing direction of the laser scanner. Next, we rotate the projected points to the $x-y$ plane and determine the minimum bounding rectangle (i.e. the bounding box of the convex hull of the points). Rotating the corner points of the bounding rectangle back to our original coordinate system gives the 3-D coordinates of the calibration boards' corners. All three steps are illustrated in Fig. 5a.

We emphasize that we can exploit our knowledge about the dimensions of the calibration board in the corner detection step of our algorithm to validate the quality of the plane fit. Allowing only a certain deviation from the true rectangle area² immediately excludes cases where the RANSAC algorithm followed by the projection and rotation step led to a wrong minimum bounding rectangle. In Fig. 5b, we


²In our experiments, we allow a deviation of 3.9 square inch.



Fig. 4: Boundary of the calibration board projected into the IR images based on the automatically calibrated intrinsic and extrinsic IR camera parameters.



(a) The three steps



(b) Detection/fitting result

Fig. 5: Illustration of the (three) steps to find the calibration board in potentially noisy laser range measurements.

show an exemplary result of fitting the calibration board to the laser range measurements. The data shown in this figure corresponds to the rightmost IR image in Fig. 4. Due to a missing ground truth, we have to rely on visual validation of the result. Considering the fact that we have to deal with noisy measurements, the orientation and dimensions of the board are sufficiently accurate.

D. Extrinsic Calibration

For each calibration image, we estimate the extrinsic parameters of the cameras relative to the calibration pattern (Section II). The distance from the calibration points to the corners of the calibration board is known. This allows to calculate the image coordinates of the corners of the calibration board by projecting the four points into the image, illustrated in Fig. 4, using the previously calculated extrinsic parameters. Consequently, we have knowledge about the 2-D image coordinates and the 3-D coordinates of the corners of the calibration board in all images. The extrinsic camera parameters relative to the laser are then estimated by minimizing the re-projection error (using the OpenCV library). The position and rotation of the cameras relative to the laser coordinates is fixed. This means that we can use the 2D-to-3D point correspondences of every calibration image together, as if we would be using a much bigger and much more complex calibration object. This helps to make the calibration process robust against sensor noise (mainly introduced by the laser scanner).

III. EXPERIMENTS

The objective of the first part of the experiments (i.e. Section III-A) is to evaluate the quality of the IR calibration method against state-of-the-art calibration methods for visible-light cameras. Those experiments are conducted in daylight, otherwise the visible-light cameras would fail to provide any useful data. The second part of the experiments (i.e. Section III-B) provides a visual comparison of depth map estimates at conditions similar to a *driving-by-night* scenario.

Our multi-sensor platform consists of a (visible-light) Bumblebee XB3 camera (15 FPS at a resolution of 1280×960 pixel), a PathfindIR IR camera (8 FPS at a resolution of 360×288 pixel) and Velodyne HDL-64E S2 3-D laser range scanner (operating at 10Hz). The dimensions of our calibration board are $20'' \times 10''$. Image and laser scan acquisition is performed using three equivalent standard PCs, synchronized by a software implementation³ of the IEEE 1584 Precision Time Protocol (PTP). Since calibration is done using static scenes, accurate time synchronization guarantees that the sensors observe the same pose of the calibration board. Obviously, this is an essential requirement for sensor to sensor calibration. For the Bumblebee calibration we rely on the automatic checkerboard corner detection and calibration algorithm implemented in the OpenCV library [9], but any other calibration tool could be used instead as well.

A. Camera Calibration

First, we evaluate the robustness of the IR resistor detection approach on a number of IR images captured under varying environmental conditions and different views. The first test set consists of 15 images captured indoors (i.e. office) with an environmental temperature of $\approx 72^\circ\text{F}$. The second test set consists of 20 images captured outdoors with an environmental temperature of $\approx 50^\circ\text{F}$. Table I lists the fraction of all images where the resistors are successfully detected. The table further lists the re-projection error using the automatically calibrated intrinsic and extrinsic parameters. We compare the re-projection error to the automatic checkerboard corner detection and calibration algorithm implemented in the OpenCV library. The Bumblebee XB3 images are selected to show the same outdoor scene and calibration board pose we used to estimate the IR camera parameters. In order to obtain a fair comparison, we have to crop 40 pixel of the left and right part of each Bumblebee XB3 image (to obtain the same aspect ratio as the IR camera images) and scale the resulting image down to 360×288 pixel. For comparison, the checkerboard is detected two times in every picture: The first time, we detect the checkerboard in the full resolution image (i.e. 1280×960 pixel) and downscale the points afterwards. The second time, we first downscale the images and then detect the checkerboard in the low resolution image. The detection in the full resolution image represents our ground truth and the detection in the

downscaled version of the images represents the scenario which is comparable to resistor detection in the IR images.

With the OpenCV checkerboard detector, the calibration points are found in every image of the *Outdoor* and the *Office* scene using the full resolution Bumblebee XB3 images (see Table I). Detection in the downscaled images, however, is only possible in ≈ 60 percent of all cases. The proposed IR resistor detection achieves considerably better detection rates, ranging between 75 and 100 percent. We further observe that the re-projection error of the IR calibration approach is slightly higher than the re-projection error of the OpenCV calibration result on the Bumblebee XB3 images (≈ 0.3 pixel) but well below one pixel. The re-projection error obtained in the office environment is higher for all three tests because the camera distortion has more impact on objects closer to the camera (and thus close to the image borders). These results clearly demonstrate that the detection of the resistors and the calibration of the IR camera provides comparable results to the calibration of a visible-light camera using a *classic* checkerboard pattern.

B. Depth Map Estimation

As an application of our automatic sensor calibration, we choose a recently proposed range data upsampling approach by Dolson et al. [13] to build dense depth maps from IR images and 3-D laser range measurements. The algorithm is basically a modification of the d -dimensional filtering approach of Adams et al. [14], adapted to enable d -dimensional interpolation.

To obtain accurate depth maps, it is imperative that the sensors are properly calibrated to each other, since we project the laser range measurements into the IR images. According to [13], calibration was performed by hand in the original work. Given that sensors are not permanently mounted on a vehicle or robot, this can be a quite time consuming and error prone task. Our objective is to demonstrate that by using IR images as an input to the algorithm, we obtain reliable depth information for living objects even in environments with no light or bad light conditions. To simulate such a scenario, we choose an underground parking garage where the only source of illumination is emergency light. In order to use the Bumblebee XB3 camera in this scenario, we would have to adjust exposure time to a maximum, consequently introducing an unacceptable amount of motion blur.

Regarding the parameter setting of the upsampling algorithm, we build position vectors p_i as (g, u, v) for each IR image pixel where g represents the intensity value of the pixel (hence encodes the IR information) and u, v represent the pixel location in 2-D. In contrast to the original work, we do not have color information available and further omit the time t in the composition of the position vectors. The extension to a dynamic environment is out of the scope of this work, but is straightforward in the framework of [13]. Obviously, range values outside the IR image plane are discarded. Since the algorithm performs an instance of Gaussian filtering, we have to set the standard deviations of each dimension. As in [13], we perform a grid search to

³<http://ptpd.sourceforge.net/>

Environment	IR (our algorithm)		Visible Light, Full (OpenCV)		Visible Light, Small (OpenCV)	
	Detection	Error [%]	Detection	Error [%]	Detection	Error [%]
Outdoor	15/20	0.2677	20/20	0.0893	12/20	0.0822
Office	15/15	0.3007	23/23	0.2084	15/23	0.1534

TABLE I: Detection rate and re-projection error of resistors in IR images with respect to different capture environments.

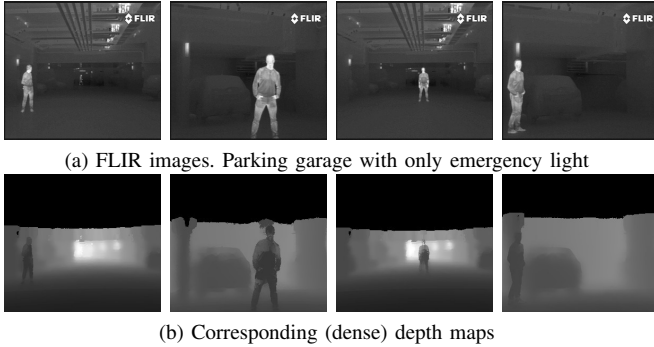


Fig. 6: FLIR images and corresponding depth maps.




Fig. 7: Depth map comparison of the parking garage scene, where the only source of illumination is emergency light.

determine the optimal parameters. Fig. 6 shows a set of IR images (top row) captured in the parking garage, as well as the corresponding (dense) depth maps (bottom row).

To illustrate the difference in estimated depth when relying on images from the (visible-light) Bumblebee XB3 camera a low light scenario is shown in Fig. 6. The images were acquired in the same parking garage with only emergency light. In Fig. 7b, we can see that the camera picture is almost completely black except for the emergency light which illuminates a part of the car. In such a scenario, the upsampling code degenerates to a mere smoothing of the laser range measurements. Consequently, the person standing in front of the wall (see Fig. 7b) is almost invisible. In contrast, the upsampled data based on the IR camera image (see Fig. 7a) preserves the objects contours and thus the person clearly stands out from the wall. This system is especially useful for the detection of humans (and living objects in general), due to the fact that they emit a substantial amount of IR radiation.

IV. CONCLUSION

Avoiding collisions with living objects, such as humans or animals is presumably the primary objective in any situation involving autonomously moving objects (e.g. robots, cars, etc.). However, at night or under severe illumination impairment, vision systems relying on visible-light cameras cease to provide reasonable information. In this work, we

demonstrated that an IR camera is a valuable asset in such cases and we presented a novel approach to allow automatic estimation of the intrinsic IR camera parameters. We further showed that fully automatic estimation of extrinsic multi-sensor system parameters including a 3-D laser range scanner, a visible-light camera and an IR camera is feasible using just one (portable & low-cost) calibration board. Eventually, we exemplified our approach by showing that we can still estimate dense depth maps by fusing IR information and laser range measurements in situations where information from visible-light cameras is useless.

REFERENCES

- [1] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2004.
- [2] S. Wasielewski and O. Strauss, "Calibration of a multi-sensor system laser rangefinder / camera," in *Proceedings of the IEEE Intelligent Vehicles Symposium (IV '95)*, Detroit, MI, United States, Sept. 1995, pp. 472 – 477.
- [3] Q. Zhang and R. Pless, "Extrinsic calibration of a camera and laser range finder (improves camera calibration)," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '04)*, St. Louis, MO, United States, Sept. 2004, pp. 2301 – 2306.
- [4] M. Buehler, K. Iagnemma, and S. Singh, *The DARPA Urban Challenge: Autonomous Vehicles in City Traffic*, 1st ed. Springer Publishing, 2009.
- [5] D. Scaramuzza, A. Martinelli, and R. Siegwart, "A toolbox for easily calibrating omnidirectional cameras," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and System (IROS '06)*, Beijing, China, Oct. 2006, pp. 5695–5701.
- [6] R. Unnikrishnan and M. Hebert, "Fast extrinsic calibration of a laser rangefinder to a camera," Carnegie Mellon Robotics Institute, Pittsburgh, PA, United States, Tech. Rep. CMU-RI-TR-05-09, July 2005.
- [7] Z. Zhang, "A flexible new technique for camera calibration," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 11, pp. 1330–1334, Nov. 2000.
- [8] R. Tsai, "An efficient and accurate camera calibration technique for 3D machine vision," in *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR '86)*, Miami Beach, FL, United States, June 1986, pp. 364–374.
- [9] G. Bradski, "The OpenCV Library," *Dr. Dobbs's Journal*, pp. 120–126, Nov. 2000.
- [10] N. Otsu, "A threshold selection method from gray-level histogram," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 9, no. 1, pp. 62–66, Jan. 1979.
- [11] C. Glennie and D. Lichti, "Static calibration and analysis of the Velodyne HDL-64E S2 for high accuracy mobile scanning," *Remote Sensing*, vol. 2, no. 6, pp. 1610–1624, June 2010.
- [12] M. Fischler and R. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, June 1981.
- [13] J. Dolson, J. Baek, C. Plagemann, and S. Thrun, "Upsampling range data in dynamic environments," in *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR '10)*, San Francisco, CA, United States, June 2010, pp. 1141–1148.
- [14] A. Adams, N. Gelfand, J. Dolson, and M. Levoy, "Gaussian KD-trees for fast high-dimensional filtering," *ACM Transactions on Graphics*, vol. 28, no. 3, pp. 21:1–21:12, July 2009.