

# Intrinsic Calibration of Depth Cameras for Mobile Robots using a Radial Laser Scanner

David Zuñiga-Noël, Jose-Raul Ruiz-Sarmiento, and Javier Gonzalez-Jimenez

Machine Perception and Intelligent Robotics group (MAPIR). Dept. of System Engineering and Automation. Instituto de Investigación Biomédica de Malaga (IBIMA). University of Malaga. Spain.  
{dzuniga,jotaraul,javiergonzalez}@uma.es

**Abstract.** Depth cameras, typically in RGB-D configurations, are common devices in mobile robotic platforms given their appealing features: high frequency and resolution, low price and power requirements, among others. These sensors may come with significant, non-linear errors in the depth measurements that jeopardize robot tasks, like free-space detection, environment reconstruction or visual robot-human interaction. This paper presents a method to calibrate such systematic errors with the help of a second, more precise range sensor, in our case a radial laser scanner. In contrast to what it may seem at first, this does not mean a serious limitation in practice since these two sensors are often mounted jointly in many mobile robotic platforms, as they complement well each other. Moreover, the laser scanner can be used just for the calibration process and get rid of it after that. The main contributions of the paper are: i) the calibration is formulated from a probabilistic perspective through a Maximum Likelihood Estimation problem, and ii) the proposed method can be easily executed automatically by mobile robotic platforms. To validate the proposed approach we evaluated for both, local distortion of 3D planar reconstructions and global shifts in the measurements, obtaining considerably more accurate results. A C++ open-source implementation of the presented method has been released for the benefit of the community.

**Keywords:** Sensor Calibration · Depth Cameras · Mobile Robotics.

## 1 Introduction

Nowadays, many mobile robots get awareness of their workspaces using RGB-D cameras [1,2]. These compact and affordable sensors provide per-pixel depth measurements along with colour information at high frame rates, simplifying a variety of robotic tasks that would be more involved if using a regular camera only, such as 3D object detection and localization [3,4], safe autonomous navigation [5], or map building/scene reconstruction [6,7,8], among others. Alternative sensors providing 3D depth information are LiDAR [9] or Time-of-Flight cameras [10], but they are not as widely spread as structured-light depth sensors, mainly due to their higher price [11].

Unfortunately, affordability of structured-light depth cameras comes at a cost: depth estimates are affected by significant distortion, not always well modeled by factory calibration parameters [12,13]. These errors can be unacceptable for some common robotic applications, and thus require a further calibration by the user. For example, we empirically observed that an obstacle-free path through an open door can be narrowed by intrinsic depth errors up to a point where it appears to the robot as a colliding path. We also experienced the negative effect of inaccurate measurements in algorithms for plane segmentation, scene reconstruction, and human pose estimation [2,14,15].

With the massive deployment of robotic platforms [16], calibration methods suitable to be executed automatically by robots are desirable, seeking to prevent the manual calibration of each depth sensor prior to deployment. However, existing intrinsic calibration methods for structured-light depth cameras cannot be easily automated. For example, the method described in [17] aims to correct depth measurements via visual SLAM and, therefore, has the underlying requirement of a well illuminated, textured enough environment. Authors in [18] argue that their method could be executed automatically, however, it is applicable only for sensors mounted with a near zero pitch angle. Recently, authors in [19] proposed another calibration approach based on the observation of a known checkerboard pattern with a regular RGB camera. In order to enable automatic calibration, their approach requires manipulating the environment to include the visual pattern which, in turn, hampers the deployment process.

In this paper, we first empirically analyze the behaviour of structured-light depth cameras and then present a method to compensate for systematic errors in the measurements, which can be easily executed automatically by mobile robotic platforms. More precisely, the proposed method requires observing, at different distances, a vertical planar surface (*e.g.* a wall) from both the depth camera and another extrinsically calibrated sensor (*e.g.* a 2D laser scanner, device commonly found in robotic platforms) not suffering from those errors. In this way, the second sensor is used to obtain depth references for calibration. Note that planar surfaces are ubiquitous in human-made environments and specific visual calibration patterns are not required. Bias functions for intrinsic depth errors are then calibrated in a Maximum Likelihood Estimation framework. The output of the calibration method are per-pixel quadratic bias functions from which systematic errors in depth measurements can be corrected in an online fashion.

To demonstrate the suitability of our proposal, we collected data from two RGB-D cameras and a 2D laser scanner mounted on a mobile robot (the robotic platform Giraff [20]) when approaching a vertical, planar surface, and carried out an experimental evaluation showing both quantitative and qualitative performance results. A C++, ROS integrated open-source implementation of the presented method is available at: [https://github.com/dzunigan/depth\\_calibration](https://github.com/dzunigan/depth_calibration)

## 2 Related Work

Early works in depth error calibration aimed to calibrate distortions along with the extrinsic parameters with respect to an RGB camera. For example, the authors in [21] considered the calibration of an RGB-D camera pair resorting to a linear depth distortion function, while Herrera *et al.* [22] tackled the calibration of two colour cameras and a depth one. In the latter case the disparity distortion was modelled as a per-pixel offset with exponential decay governed by two global parameters. Both approaches employ planar surfaces for depth compensation, tendency that still holds in recent works. An example of this is the work by Basso *et al.* [19], which proposed a calibration method based on the observation of a planar pattern with a regular camera, while the extrinsic calibration is more a “side effect”.

All above-mentioned works require a visual pattern (typically a checkerboard) in order to compute reference depth measurements, and must be included in the robot workspace to enable automatic calibration. There are works not requiring this, like [18], where the authors proposed a non-parametric calibration approach and they get rid of the visual pattern requirement. However, to perform the calibration on a mobile robot, another sensor (*e.g.* laser scanner) is required in order to provide reference values, and only sensors mounted with a near zero pitch angle can be calibrated.

Another way to get rid of known visual patterns is by using a visual SLAM pipeline to provide the depth references. To the best of our knowledge, depth correction via SLAM was first introduced by Teichman *et al.* [17]. Their method makes the strong assumption that the errors at close ranges (below 2 m) are negligible, and thus are used as reference within the SLAM pipeline. Depth correction factors are then estimated for each pixel and at a number of fixed distances. Another work based on a similar idea was presented in [23], where the authors assume known extrinsic calibration between the RGB and the depth cameras. Their method projects features from a sparse map (generated from the RGB camera) into the depth camera poses in order to estimate the correction factors. They use the thin plate spline as a tool for approximating a dense representation of the sparse correction factors. The main drawback of these methods is that they have the underlying requirements of well illuminated and textured enough environments in order to provide reliable estimates.

The calibration method presented in this work does not require any visual pattern and thus can be easily executed automatically by mobile robots. As in [18], another sensor is needed to provide reference measurements, concretely a radial laser scanner. Notice that this assumption is not very restrictive, since these sensors are commonly used in mobile robotic platforms. Moreover, we can use the laser scanner temporally just for the calibration process and get rid of it after that. In contrast to [18], we argue that the systematic depth errors can be well modeled from a more compact parametric representation. Additionally, our method does not assume a specific orientation of the depth camera to carry out the calibration.

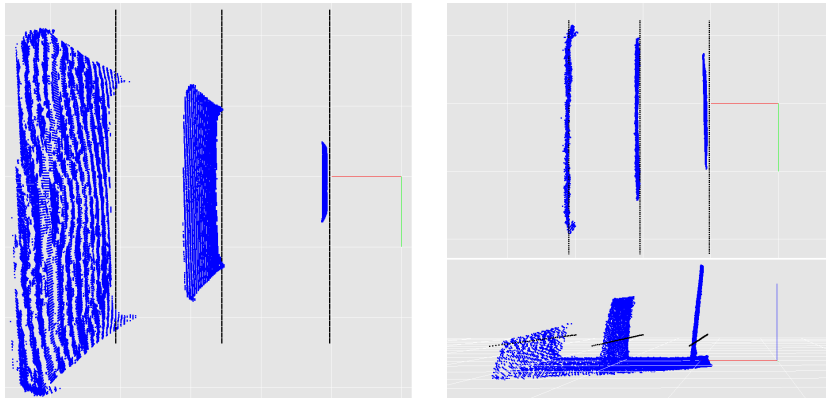


Fig. 1: Illustration of the errors and their variation with distance. Left, a depth camera observing a perpendicular wall at 1-4 m. Right, another camera with a 60 deg pitch observing the same wall, at 1–3 m. Note that the reconstructed ground is parallel to the  $x$ - $y$  plane, while the wall has a noticeable inclination.

### 3 Depth Error Model

In this work, as in [19], we consider both the “local distortion” and “global” errors as the main source of systematic errors. The *local distortion* has the characteristic effect of deforming the resulting point cloud, while the *global errors* shift the average observed depth. Illustrative examples of these errors are shown in Figure 1. We argue that both sources of error can be explained by a depth bias  $\beta_{u,v}$ :

$$z_{u,v} = z_{u,v}^* + \beta_{u,v}, \quad (1)$$

for each pixel  $(u, v) \in \Omega$  in the image domain independently, where  $z_{u,v}^*$  and  $z_{u,v}$  represent the true and the measured depths, respectively. We consider the bias to be normally distributed:

$$\beta_{u,v} \sim \mathcal{N}(\mu_{u,v}(z_{u,v}), \sigma^2(z_{u,v})), \quad (2)$$

where  $\mu_{u,v}$  is a per-pixel mean function and  $\sigma$  is a global standard deviation function modeling the uncertainty in the measurements.

The bias, computed as the difference between the measured depth and the real one, are plotted in Figure 2a as a function of the measured depth, for different pixels. The lines in that figure represent fitted quadratic models (see Section 4.2). It becomes clear that each pixel is affected by a different bias, but the evolution of the biases with respect to depth are well explained by quadratic functions.

Regarding the uncertainty in the measurements, previous research [24] found that it follows a quadratic evolution with respect to depth. We verified this behaviour empirically by analyzing the standard deviation of the measurements in a similar setting as for the biases. The standard deviation plotted against

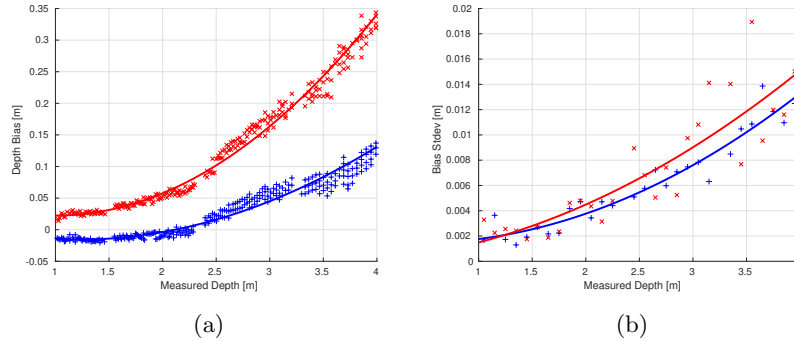


Fig. 2: The observed bias (2a) and bias noise (2b) as a function of the measured depth, along with quadratic curve fits, for two different pixels.

the measured depth are reported in Figure 2b. Notice that, unlike the bias, the uncertainty of the measurements is similar for different pixels. This phenomenon has also been considered in our framework by modeling a single variance function for all pixels (see Section 4.2).

Finally, the systematic depth errors can be compensated by subtracting the bias mean:

$$\bar{z}_{u,v} = z_{u,v} - \mu_{u,v}(z_{u,v}) = z_{u,v}^* + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma^2(z_{u,v})), \quad (3)$$

yielding unbiased depth measurements.

## 4 Calibration Approach

In this section we describe the proposed calibration approach. First, Section 4.1 discusses the process of computing depth reference measurements from the observation of a planar surface by the sensors. Then, the formulation of the calibration problem in a Maximum Likelihood framework and its solution are described in Section 4.2.

### 4.1 Computation of the Depth References

The input of the calibration method are observations of a vertical, planar surface from both a depth camera and another sensor not suffering from the same errors. For the former, observations are in the form of depth images, while for the latter they are in the form of geometric parameters of the observed plane. These parameters are  $(\mathbf{n}, d) \in \mathbb{R}^3 \times \mathbb{R}^+$  such that:

$$\mathbf{n} \cdot \mathbf{x} - d = 0, \quad (4)$$

for any point  $\mathbf{x} \in \mathbb{R}^3$  lying on the plane. Here,  $\mathbf{n}$  represents the unit normal vector (from the origin to the plane) and  $d \geq 0$  the perpendicular distance to the origin (Hessian normal form).

The extrinsic calibration  $(\mathbf{R}, \mathbf{t}) \in \text{SE}(3)$  between the two sensors allow us to express the plane parameters observed by the second sensor into the coordinate system of the depth camera. Clearly, the new normal vector  $\mathbf{n}'$  is affected only by the rotation  $\mathbf{R}$ , while the new distance  $d'$  can be computed as:

$$\mathbf{n}' = \mathbf{R}\mathbf{n}, \quad d' = -\mathbf{n}' \cdot \mathbf{t} - d, \quad (5)$$

which is the distance of the new origin  $-\mathbf{t}$  from the rotated coordinates (before translation).

Depth cameras allow to reconstruct 3D points via back-projection, using the associated depth measurements and the intrinsic camera parameters (provided by the manufacturer). We parameterize the 3D line representing an incoming ray with respect to depth  $z \in \mathbb{R}$  as:

$$\mathbf{l}_{u,v}(z) = z \left( \frac{u - c_x}{f_x}, \frac{v - c_y}{f_y}, 1 \right)^\top, \quad (6)$$

where  $(c_x, c_y) \in \mathbb{R}^2$  refers to the camera center, and  $f_x, f_y \in \mathbb{R}$  to the focal lengths (in each axis). In this way, the reconstructed 3D point can be computed as  $\mathbf{l}_{u,v}(z_{u,v})$ . Therefore, we define the reference depth measure  $z_{u,v}^*$  such that:

$$\mathbf{n}' \cdot \mathbf{l}_{u,v}(z_{u,v}^*) - d' = 0, \quad (7)$$

*i.e.* enforcing the plane constraint in Eq. (4) on the reconstructed 3D point. Finally, since Eq. (7) is linear with respect to  $z_{u,v}^*$ , the solution can be computed as:

$$z_{u,v}^* = \frac{d'}{\mathbf{n}' \cdot \mathbf{l}_{u,v}(1)}, \quad (8)$$

yielding pairs  $(z_{u,v}, z_{u,v}^*)$  that relate measured and reference depth values.

## 4.2 Maximum Likelihood Estimation of the Bias Functions

Once having computed the depth measurement-reference pairs, the estimation of the bias functions is divided into two main steps. First, we fit a quadratic function to the observed deviations, which is common for all pixels (recall Figure 2b). Next, for each pixel independently, we solve for the actual bias parameters (recall Figure 2a).

In first place, we want to estimate the parameters of a quadratic function that best represents the evolution of the bias noise. In a Least Squares sense, this is:

$$\arg \min_{a,b,c} \sum_{k \in \Pi} \|\sigma_k - \sigma(k)\|^2, \quad \sigma(k) = ak^2 + bk + c, \quad (9)$$

given the discrete deviation samples  $\sigma_k$  over the discrete sampling interval  $\Pi$ . In order to compute observed standard deviations, we divide the observed bias into discrete bins for each pixel independently:

$$S_{u,v}^k = \{z - z^* \mid t > |z - k|, \forall (z, z^*) \in M_{u,v}\}, \quad (10)$$

where  $t \in \mathbb{R}$  is a discretization threshold and  $M_{u,v}$  is the set of depth pairs for a pixel  $(u, v) \in \Omega$ . Then, for each set of observations  $S_k$  with  $k \in \Pi$ , we compute the deviation  $\sigma_k$  as:

$$\sigma_k^2 = \frac{1}{\sum_{(u,v) \in \Omega} |S_{u,v}^k|} \sum_{(u,v) \in \Omega} \left( \sum_{z \in S_{u,v}^k} (z - \bar{S}_{u,v}^k)^2 \right), \quad (11)$$

where  $|S|$  represents the set's cardinality and  $\bar{S}$ , the mean. Equation 11 aims to compute the variance of a range of depth measurements, where each pixel can have a different bias mean. This way, we obtain the discrete samples  $\sigma_k$  used to fit the function modeling the bias noise in Eq. (9).

Having an estimation of the uncertainty, we proceed to solve for the parameters of a quadratic approximation to the bias function for each pixel independently. We formulate the calibration problem in a Maximum Likelihood Estimation fashion as:<sup>1</sup>

$$\arg \max_{a,b,c} \prod_{i=1}^N p(z_i | z_i^*, \mu(z_i), \sigma(z_i)), \quad \mu(z) = az^2 + bz + c, \quad (12)$$

for a likelihood function  $p$  and  $N$  independent observations. Under the assumption of normality, the likelihood function becomes:

$$p(z | z^*, \mu, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(z - z^* - \mu)^2}{2\sigma^2}\right). \quad (13)$$

Taking the negative logarithm of Eq. (12) yields an equivalent Least Squares problem:

$$\arg \min_{a,b,c} \sum_{i=1}^N \frac{1}{\sigma^2(z_i)} \|z_i - z_i^* - \mu(z_i)\|^2, \quad (14)$$

which has a closed form solution since the residual expression is linear with respect to the optimization parameters. In this way, solving Eq. (14) we approximate a bias function that can be used to compensate for the measured depths in a per-pixel fashion.

## 5 Experimental Evaluation

The goal of the experimental evaluation is to validate our approach in a real setting. In this respect, we provide a quantitative evaluation of how well the error model presented in Section 3 can handle both the local distortions (Section 5.1) and the global errors (Section 5.2). We also show qualitative improvements in 3D reconstructions after calibration (Section 5.3).

To carry out these evaluations, we recorded two independent sequences with two RGB-D sensors (Orbbec Astra) and a 2D laser scanner (Hokuyo URG-04LX-UG01). The sensors were mounted on a Giraff [20] mobile robot and we

<sup>1</sup> Hereafter, we drop the  $u, v$  subscript to improve readability.

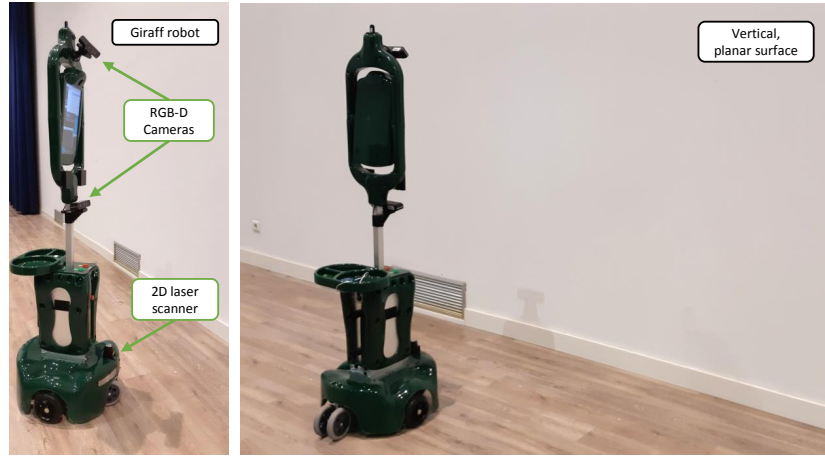


Fig. 3: Left, Giraff robot with annotations of the sensors involved in the calibration process. Right, the robot facing a planar surface during data collection.

recorded the sequences while moving it towards and away a wall. The sensor setup and a snapshot of the collection procedure are depicted in Figure 3. Note that the upper RGB-D camera has a non-negligible pitch, while the other camera and the laser are mounted horizontally, *i.e.* with near zero pitch. The extrinsic calibration parameters between the sensors were estimated using the automatic multi-sensor method proposed in [25]. From the two recorded sequences, one was used to perform the depth calibration described in Section 4, while the other one was used for evaluation purposes. For the sake of reproducibility, the collected data is available at: <https://doi.org/10.5281/zenodo.2636878>.

### 5.1 Local Distortion Evaluation

In order to evaluate the undistortion performance, we follow a similar approach as described in [19]. Since the local distortion errors deform the reconstructed 3D structure, the evaluation method consists of fitting a plane to the point cloud acquired while observing a wall, and then computing the Root Mean Square (RMS) perpendicular distance to the extracted plane for each point belonging to the planar surface.

This is, for a plane  $\pi$ , we have:

$$e_{\perp}(\pi) = \sqrt{\frac{1}{N} \sum_{i=1}^N \|\mathbf{n}_{\pi} \cdot \mathbf{x}_i - d_{\pi}\|^2}, \quad (15)$$

for each 3D point  $\mathbf{x}_i \in \mathbb{R}^3$  of the planar surface.

The evaluation results for the lower and upper cameras, in terms of the RMS perpendicular error, are shown in Figure 4a and 4b, respectively. We can see that



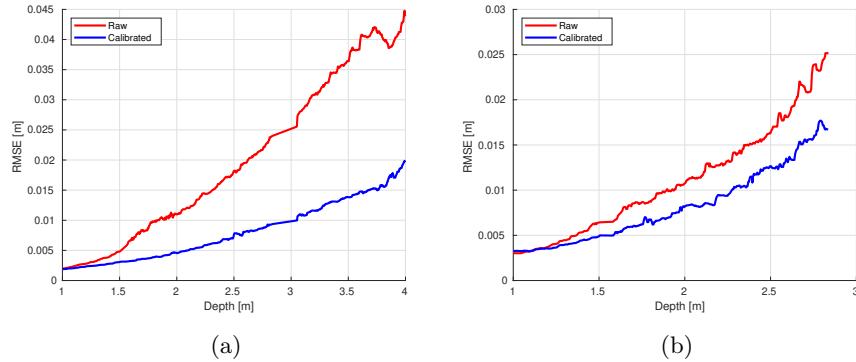


Fig. 4: Local distortion performance evaluation for the two RGB-D cameras (4a lower camera; 4b upper camera). In both cases, calibrated depth measurements show better performance.

calibrated depth measurements achieve better performance in both cases when compared to the original ones. For example, the calibration improves  $\sim 2.5$  cm the RMSE at 4 m for the lower camera, and  $\sim 1$  cm near 3 m for the upper one. It is also noticeable the difficulties that calibrated measurements have in reaching error-free measurements, and how the error grows with respect to depth. This behaviour can be explained by the quantization error of the sensor, as argued in [19]. In the case of the upper camera, errors are even larger. This phenomenon is caused by the nonzero pitch angle of the camera, as noise in the measurements increases when observing surfaces away from the perpendicular orientation, as shown in [26].

## 5.2 Global Error Evaluation

In order to evaluate the global error, we follow a similar approach as before, but in this case we computed the perpendicular error with respect to a reference plane. Recall that the global error shifts the measurements away from their true value. Thus, we can evaluate the error function in Eq. (15) with respect to the plane as observed by the laser scanner.

The RMSE of the compensated and original depth measurements are reported in Figures 5a and 5b for the lower and upper cameras, respectively. Here, the calibration improves up to 4 cm the RMSE for the lower camera (at 4 m) and up to 2.5 cm for the upper one (at 1.5 m). Errors also tend to grow with depth, for the same reasons as before. Additionally, global errors after calibration are higher than the local distortion ones. This is mainly due to other external sources of errors affecting the evaluation, as *e.g.* errors in the laser measurements, extrinsic calibration errors or time delays between the laser and the cameras. Despite of this, in both cases, the use of calibrated depth measurements improves the accuracy of the measurements.

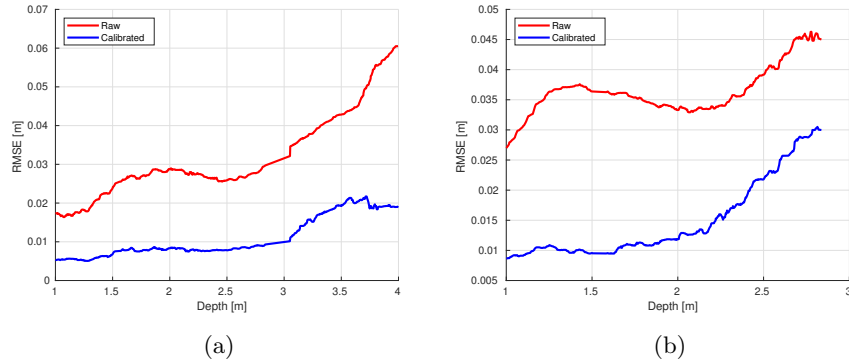


Fig. 5: Global distortion performance evaluation for the two RGB-D cameras (5a lower camera; 5b upper camera). A significantly lower error is shown when using calibrated measurements.

### 5.3 Qualitative Evaluation

In this section, we provide a qualitative evaluation of the obtained, reconstructed 3D point clouds when using compensated depth measurements compared to the original ones. For that purpose, we compare the reconstruction of vertical walls to ground truth measurements before and after calibration. For space reasons, only the results for the lower camera are shown.

The reconstructed point clouds using the raw, original depth measurements are shown in Figure 6-left, while Figure 6-right reports the corrected point clouds after calibration. At closer distances, the distortion is negligible, while small offsets are noticeable in the raw measurements. At higher distances, distortions are clearly visible. It can be observed that, after calibration, both the small offsets in the measurements as well as distortions are significantly corrected.

## 6 Conclusion

In this work, we presented a method to calibrate systematic errors arising from depth cameras that can be easily executed by mobile robot platforms. First, we analyzed and characterized these errors, and then proposed a calibration method based on Maximum Likelihood Estimation. This method requires to observe a planar surface with both the depth camera and another sensor (*e.g.* a radial laser scanner), used to compute reference depth measurements. The output of the calibration are per-pixel parametric bias functions that can be used to compensate for these systematic depth errors. We evaluated the proposed method in a real robotic platform equipped with two RGB-D cameras and a 2D laser scanner, and showed that the proposed model can handle both local distortions and global errors, producing considerably more accurate measurements. We also provided a qualitative evaluation of the method performance, reporting noticeable error

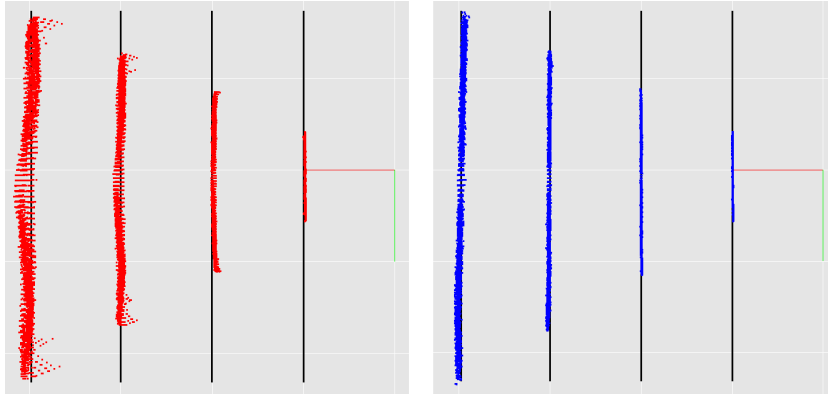


Fig. 6: Reconstructed point clouds form raw (left) and calibrated measurements (right), with reference measurements shown as black lines at 1–4 m.

corrections. In the future we plan to incorporate a robust plane detection mechanism in order to enhance the method performance in cluttered environments.

**Acknowledgments.** This work has been supported by the research projects *WISER* (DPI2017-84827-R), funded by the Spanish Government and the European Regional Development’s Funds (FEDER), *MoveCare* (ICT-26-2016b-GA-732158), funded by the European H2020 program, the European Social Found through the Youth Employment Initiative for the promotion of young researchers, and a postdoc contract from the I-PPIT program of the University of Malaga.

## References

1. Jing, C., Potgieter, J., Noble, F., Wang, R.: A comparison and analysis of RGB-D cameras’ depth performance for robotics application. In: International Conference on Mechatronics and Machine Vision in Practice (M2VIP) (2017)
2. Ruiz-Sarmiento, J.R., Galindo, C., Gonzalez-Jimenez, J.: Robot@Home, a robotic dataset for semantic mapping of home environments. *The International Journal of Robotics Research* **36**(2) (2017)
3. Schwarz, M., Milan, A., Periyasamy, A.S., Behnke, S.: RGB-D object detection and semantic segmentation for autonomous manipulation in clutter. *The International Journal of Robotics Research* **37**(4-5) (2018)
4. Ruiz-Sarmiento, J.R., Galindo, C., Gonzalez-Jimenez, J.: A survey on learning approaches for undirected graphical models. application to scene object recognition. *International Journal of Approximate Reasoning* **83** (2017)
5. Jaimez, M., Blanco, J.L., Gonzalez-Jimenez, J.: Efficient reactive navigation with exact collision determination for 3D robot shapes. *International Journal of Advanced Robotic Systems* **12**(5) (2015)
6. Kähler, O., Prisacariu, V.A., Murray, D.W.: Real-time large-scale dense 3D reconstruction with loop closure. In: Computer Vision – ECCV 2016 (2016)

7. Jamiruddin, R., Sari, A.O., Shabbir, J., Anwer, T.: RGB-depth SLAM review. arXiv preprint [arXiv:1805.07696](https://arxiv.org/abs/1805.07696) (2018)
8. Ruiz-Sarmiento, J.R., Galindo, C., Gonzalez-Jimenez, J.: Building multiversal semantic maps for mobile robot operation. *Knowledge-Based Systems* **119** (2017)
9. Zhang, J., Singh, S.: LOAM: Lidar odometry and mapping in real-time. In: *Robotics: Science and Systems*. vol. 2 (2014)
10. Foix, S., Alenyà, G., Torras, C.: Lock-in time-of-flight (ToF) cameras: A survey. *IEEE Sensors Journal* **11**(9) (2011)
11. Rusu, R.B., Cousins, S.: 3D is here: Point cloud library (PCL). In: *International Conference on Robotics and Automation* (2011)
12. Song, X., Zheng, J., Zhong, F., Qin, X.: Modeling deviations of RGB-D cameras for accurate depth map and color image registration. *Multimedia Tools and Applications* **77**(12) (2018)
13. Fiedler, D., Müller, H.: Impact of thermal and environmental conditions on the Kinect sensor. In: *Advances in Depth Image Analysis and Applications* (2013)
14. Ruiz-Sarmiento, J.R., Galindo, C., Gonzalez-Jimenez, J.: Experimental study of the performance of the Kinect range camera for mobile robotics. Tech. rep., University of Malaga, Dpt. of System Engineering and Automation (2013)
15. Fernandez-Moral, E., Gonzalez-Jimenez, J., Rives, P., Arevalo, V.: Extrinsic calibration of a set of range cameras in 5 seconds without pattern. In: *International Conference on Intelligent Robots and Systems* (2014)
16. Della Corte, B., Andreasson, H., Stoyanov, T., Grisetti, G.: Unified motion-based calibration of mobile multi-sensor platforms with time delay estimation. *IEEE Robotics and Automation Letters* **4**(2) (2019)
17. Teichman, A., Miller, S., Thrun, S.: Unsupervised intrinsic calibration of depth sensors via SLAM. In: *Robotics: Science and Systems*. vol. 248 (2013)
18. Di Cicco, M., Iocchi, L., Grisetti, G.: Non-parametric calibration for depth sensors. *Robotics and Autonomous Systems* **74** (2015)
19. Basso, F., Menegatti, E., Pretto, A.: Robust intrinsic and extrinsic calibration of RGB-D cameras. *IEEE Transactions on Robotics* **34**(5) (2018)
20. González-Jiménez, J., Galindo, C., Ruiz-Sarmiento, J.R.: Technical improvements of the Giraff telepresence robot based on users' evaluation. In: *2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication* (2012)
21. Zhang, C., Zhang, Z.: Calibration between depth and color sensors for commodity depth cameras. In: *International Conference on Multimedia and Expo* (2011)
22. Herrera C., D., Kannala, J., Heikkilä, J.: Joint depth and color camera calibration with distortion correction. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **34**(10) (2012)
23. Quenzel, J., Rosu, R.A., Houben, S., Behnke, S.: Online depth calibration for RGB-D cameras using visual SLAM. In: *International Conference on Intelligent Robots and Systems* (2017)
24. Smisek, J., Jancosek, M., Pajdla, T.: 3D with Kinect. In: *International Conference on Computer Vision Workshops* (2011)
25. Zuñiga-Noël, D., Ruiz-Sarmiento, J.R., Gomez-Ojeda, R., Gonzalez-Jimenez, J.: Automatic multi-sensor extrinsic calibration for mobile robots. arXiv preprint [arXiv:1906.04670](https://arxiv.org/abs/1906.04670) (2019)
26. Nguyen, C.V., Izadi, S., Lovell, D.: Modeling Kinect sensor noise for improved 3D reconstruction and tracking. In: *International Conference on 3D Imaging, Modeling, Processing, Visualization Transmission* (2012)