

On calibration of a low-cost time-of-flight camera

Alina Kuznetsova and Bodo Rosenhahn

Institute fuer Informationsverarbeitung,
Leibniz University Hannover

Abstract. Time-of-flight (ToF) cameras are becoming more and more popular in computer vision. In many applications 3D information delivered by a ToF camera is used, and it is very important to know the camera's extrinsic and intrinsic parameters, as well as precise depth information. A straightforward algorithm to calibrate a ToF camera is to use a standard color camera calibration procedure [12], on the amplitude images. However, depth information delivered by ToF cameras is known to contain complex bias due to several error sources [6]. Additionally, it is desirable in many cases to determine the pose of the ToF camera relative to the other sensors used.

In this work, we propose a method for joint color and ToF camera calibration, that determines extrinsic and intrinsic camera parameters and corrects depth bias. The calibration procedure requires a standard calibration board and around 20-30 images, as in case of a single color camera calibration. We evaluate the calibration quality in several experiments.

Keywords: calibration, time-of-flight cameras

1 Introduction

The availability of three dimensional information is desirable in many areas of computer vision. Such information can be obtained using time-of-flight (ToF) sensors, which are becoming more and more affordable. However, ToF sensor measurements are subject to different sources of noise and systematic errors [6] (see Figure 1(a)), such as:

- Systematic errors due to simplification in distance calculation formulas and numerical imprecision.
- Intensity-related distance errors that depend on the reflectivity of different objects.
- External factors, such as lightning conditions, presence of other infra-red devices, etc.

As a consequence, precise calibration and sensor fusion become impossible without accounting for these errors, while good calibration quality is required in many computer vision problems, e.g. 3D reconstruction, or high-quality ground truth data acquisition.

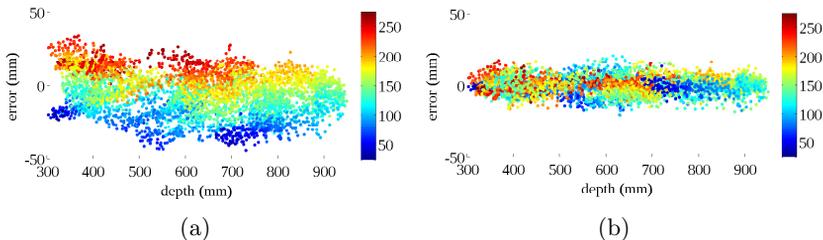


Fig. 1. 1(a) The dependency of the depth error measurement on the distance to the object and the horizontal image coordinate of the pixel is shown (horizontal coordinate is shown by color). The measurements are done using the Intel Creative Gesture Camera. 1(b) Corrected depth error: the clear structure of the error bias disappears and only the noise component is left.

1.1 Contribution

Several approaches for ToF camera calibration were proposed in [7, 8]. However, they require knowing ground truth distance to the calibration object to correct the depth bias. Additionally, no approach was presented so far for simultaneous calibration of ToF and color sensors.

An algorithm for simultaneous calibration of a depth sensor and a color camera was presented in [3] for Microsoft Kinect, which exploits structured light technology to obtain depth image. In our experiments with ToF camera, however, the precision of the resulting calibration was not sufficient, due to noise and depth bias in ToF depth data.

In our work, we propose a novel calibration method for a ToF camera, along with its simultaneous calibration with a color camera or potentially other sensors. Our approach has several advantages:

- Ground truth depth measurement are not required to correct the depth bias.
- Only 20 – 30 images of a checkerboard pattern are needed for reliable calibration, which is the same amount of data as required for a single color camera calibration.

We employ sensor fusion to compensate for different sources of the calibration error, such as low resolution, depth systematic bias and noise (in Fig 1(b) the depth error is shown after running our algorithm; it can be seen, that the systematic component of the error disappeared).

We evaluate our approach using the Intel Creative Gesture Camera [4] in two experiments and show, that our calibration delivers better results than manufacturer-provided calibration for this camera.

2 Related work

There exists an extensive literature on color camera calibration. The basic method is proposed in [12]: several images of a checker pattern are recorded, checker-

board corners are detected and then used to determine the intrinsic and extrinsic parameters through a non-linear optimization procedure.

Initially, ToF cameras were calibrated using amplitude images, delivered by most ToF cameras, as color images, and applying color camera calibration procedure [12]. This allows to obtain the intrinsic parameters of the camera. However, the resolution of ToF cameras is usually small (320×240 for the Intel sensor), therefore the obtained calibration is not very precise. Additionally, depth bias is not corrected in this case. In [7, 8, 5] it is shown that the bias is non-constant and shows high dependency on the distance to the camera. In the first two works it is proposed to correct it using ground truth distance measurements by fitting either a B-spline [7] to the measurement errors or creating a look-up-table [8]. Unfortunately, both methods rely on the ground truth data, which is not easy to obtain without a special setup. In [5], the bias is computed by fitting a 6-degree polynomial to the difference between the depth values, obtained through checkerboard-based calibration, and the measured depth values. However, the fact that the amplitude-based calibration is not itself accurate due to low resolution of the amplitude images is not taken into account.

Joint depth and color calibration method was proposed in [3] for another depth sensor, Microsoft Kinect, that does not provide amplitude images. In this method, 3D planes are used instead of checkerboard corners to find the parameters of the depth sensor, since the checkerboard corners are not visible in the depth image. The depth measurement bias, called depth distortion, is corrected on the basis of several images of a planar surface. However, in our experiments with the Intel Creative Gesture Camera, this method was inapplicable due to the imprecision and greater amount of noise in depth data.

Joint calibration of a system of several ToF and RGB cameras was proposed in [2]. There, projective alignment is used to find the mapping between different sensors. However, this method does not account for non-linear depth error and can only be used in case when a color camera is present.

Finally, in some applications depth bias is compensated after the calibration directly during the actual processing of 3D data, as in [1]. The 3D point cloud is computed initially using a standard corner-based calibration, and the bias is taken into account during 3D object reconstruction. However, this approach adds complexity in the data processing method and cannot be directly transferred to the problems other than 3D reconstruction.

We address the disadvantages of the approaches mentioned above by combining plane-based calibration [3] with the standard calibration procedure for the amplitude image, therefore compensating for the low resolution of the ToF images and depth inaccuracies. We also model depth bias, using a non-parametric kernel regression approach [9] for error estimation. The error is obtained by comparing the depth values measured by the ToF camera and the values obtained by the calibration procedure.

3 Calibration

For the ToF calibration, 20 to 30 images of a calibration board are required. The calibration board should consist of a black-and-white pattern (see Section 3.4 for further explanation) and sufficient space without this pattern (see Figure 3(a)).

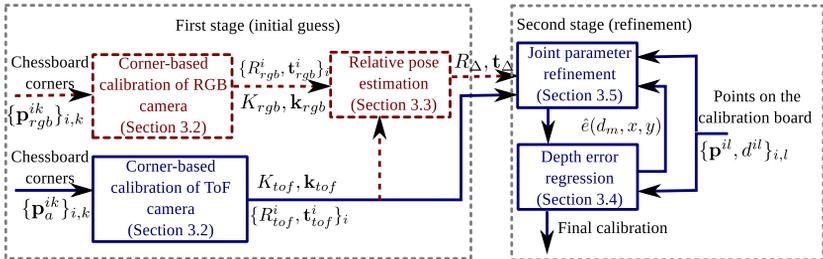


Fig. 2. Overview of the calibration procedure.

We divide our algorithm in two stages (see Figure 2). In the first stage, we obtain the initial guess for the parameters of the ToF camera. In case there is also a color camera present, we initialize its parameters and its pose relative to the ToF camera. An initial guess for the ToF intrinsic parameters is obtained using corner-based calibration on the amplitude image (Section 3.2). Moreover, for each image the relative checkerboard position is obtained. The same procedure is done for the color camera and the corresponding parameters are determined. Then, the relative pose of color and ToF cameras can be obtained.

In the second stage, we iteratively refine the initial guess by subsequent re-estimation of the camera parameters and the depth bias:

1. Re-estimate the parameters of the cameras using joint optimization (Section 3.5).
2. Depth bias is estimated as described in Section 3.4.

3.1 Camera model

We adopt the pinhole camera model to describe lens properties of both ToF and RGB cameras. Let $\mathbf{p} = (x, y)^T$ be image coordinates of a point and $(X, Y, Z)^T$ are its world coordinates. Then, the relation between these coordinates is given

by:

$$s \begin{pmatrix} x_n \\ y_n \\ 1 \end{pmatrix} = (R|t) \begin{pmatrix} X \\ Y \\ Z \\ 1 \end{pmatrix} \quad (1)$$

$$\begin{pmatrix} x \\ y \\ 1 \end{pmatrix} = K \begin{pmatrix} x_n^d \\ y_n^d \\ 1 \end{pmatrix} = \begin{pmatrix} f_x & 0 & c_x \\ 0 & f_y & c_y \end{pmatrix} \begin{pmatrix} x_n^d \\ y_n^d \\ 1 \end{pmatrix}, \quad (2)$$

where K is the projection matrix with unknown intrinsic parameters f_x, f_y, c_x, c_y , R, t is the transformation from world coordinate system to camera coordinate system (extrinsic parameters), s is a scale factor, and $x_n^d = x_n, y_n^d = y_n$ in case no lens distortion is assumed. Otherwise it can be modeled as:

$$r^2 = x_n^2 + y_n^2 \quad (3)$$

$$x_n^d = (1 + k_1 r^2 + k_2 r^4 + k_5 r^6) x_n + 2k_3 x_n y_n + k_4 (r^2 + 2x_n^2) \quad (4)$$

$$y_n^d = (1 + k_1 r^2 + k_2 r^4 + k_5 r^6) y_n + k_3 (r^2 + 2y_n^2) + 2k_4 x_n y_n, \quad (5)$$

where $\mathbf{k} = (k_1, k_2, \dots, k_5)^T$ are distortion coefficients [3].

3.2 Corner-based calibration

To obtain an initial guess for the extrinsic and intrinsic parameters for RGB and ToF cameras, corner-based calibration, as described in [12], is used. As a pre-processing step, checkerboard corners $\mathbf{p}^{ik} = (x^{ik} \ y^{ik})^T, k = 1 \dots K$ are extracted from each image with index $i = 1 \dots I$. Intrinsic parameters are then initialized using image size. The relative position of the checkerboard with respect to the camera $\{R^i, \mathbf{t}^i\}_i$ is initialized using homographies. In the next step, parameter refinement is performed using a non-linear minimization of the following functional:

$$E_0 = \sum_i \sum_k \|\mathbf{p}^{ik} - \hat{\mathbf{p}}^{ik}\|^2 \quad (6)$$

where $\hat{\mathbf{p}}^{ik} = (\hat{x}^{ik}, \hat{y}^{ik})^T$ is the projection of each 3D point of the 3D checkerboard model from the world coordinate system to the image coordinate system, using equations (1)-(5).

After this step, the initial guess for the ToF camera parameters $K_{tof}, \mathbf{k}_{tof}, \{R_{tof}^i, \mathbf{t}_{tof}^i\}_i$ and color camera parameters $K_{rgb}, \mathbf{k}_{rgb}, \{R_{rgb}^i, \mathbf{t}_{rgb}^i\}_i$ is found.

3.3 ToF and RGB relative pose estimation

In case a color camera is available, it can be used to further refine calibration. Furthermore, in many applications, knowing the relative pose between color and ToF camera is essential.

The relative pose is computed from the estimated relative poses of the planes $\{R_{tof}^i, \mathbf{t}_{tof}^i\}_i$ and $\{R_{rgb}^i, \mathbf{t}_{rgb}^i\}_i$ in each image. Firstly, plane parameterization is recomputed in the form of normal and offset from its relative pose to the camera. The parameters of the planes in ToF camera frame are denoted by $\mathbf{n}_{tof}^i, \theta_{tof}^i$ and the parameters in the color camera frame are denoted by $\mathbf{n}_{rgb}^i, \theta_{rgb}^i$. The normal vectors form the matrices $N_{tof} = (\mathbf{n}_{tof}^1, \mathbf{n}_{tof}^2, \dots, \mathbf{n}_{tof}^I)$ and $N_{rgb} = (\mathbf{n}_{rgb}^1, \mathbf{n}_{rgb}^2, \dots, \mathbf{n}_{rgb}^I)$ and the offsets — the row vectors $\mathbf{D}_{tof} = (\theta_{tof}^1, \theta_{tof}^2, \dots, \theta_{tof}^I)$ and $\mathbf{D}_{rgb} = (\theta_{rgb}^1, \theta_{rgb}^2, \dots, \theta_{rgb}^I)$. The relative rotation R_Δ is determined using the solution to the Procrustes problem to find the relative rotation [10]:

$$R_\Delta = VU^T \quad (7)$$

$$USV^T = N_{tof}N_{rgb}^T, \quad (8)$$

where (8) is SVD decomposition of the matrix $N_{tof}N_{rgb}^T$. The translation is found by minimizing the difference between the distances from the camera origin to each plane [11]:

$$\mathbf{t}_\Delta = (N_{rgb}N_{rgb}^T)^{-1}N_{rgb}(\mathbf{D}_{rgb} - \mathbf{D}_{tof}) \quad (9)$$

The relative pose is then used to determine the plane position relative to the color camera as follows:

$$R_{rgb}^i = R_\Delta R_{tof}^i \quad (10)$$

$$\mathbf{t}_{rgb}^i = R_\Delta \mathbf{t}_{tof}^i + \mathbf{t}_\Delta \quad (11)$$

$R_\Delta, \mathbf{t}_\Delta$ are refined during joint optimization, as described in Section 3.5.

3.4 Depth correction

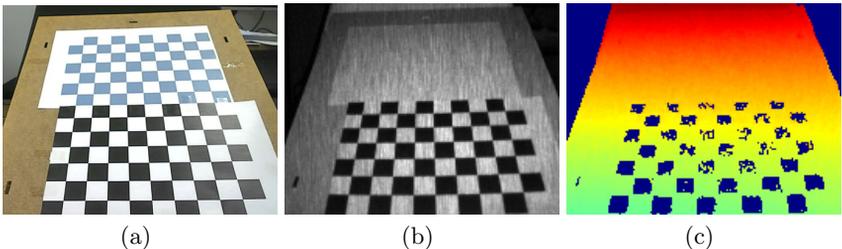


Fig. 3. 3(a) Three-color calibration pattern; 3(b) amplitude image: the black pattern is visible, while the color pattern is not; 3(c) depth image: the areas with black squares have invalid depth values.

We model depth measurements in the following way:

$$d_r = d_m + e + \xi, \quad (12)$$

where d_r corresponds to the real distance to the object, d_m is the depth measured by the camera, e is the bias specific for each camera, and $\xi \sim \mathcal{N}(0, \sigma)$ is the Gaussian noise.

In our experiments, we observed that e depends on three parameters: the real distance to the object d_r at a pixel (x, y) , as well as pixel coordinates on the camera matrix, i.e. $e = e(d_r, x, y)$.

Given that ground truth measurements are available, the dependence can be modeled by fitting a regression function to $e = d_r - d_m$. We use a non-parametric kernel regression with Gaussian kernel [9], as it is a fairly simple and fast approach.

In this case, $\hat{e}(d_r, x, y)$ is represented as follows:

$$\hat{e}(d_r, x, y) = \frac{\sum_j K(\mathbf{q}, \mathbf{q}_j) e_j}{\sum_j K(\mathbf{q}, \mathbf{q}_j)}, \quad \mathbf{q} = \begin{pmatrix} d_r \\ x \\ y \end{pmatrix}, \quad K(\mathbf{q}, \mathbf{q}_j) = e^{-\frac{1}{2h^2} \|\mathbf{q} - \mathbf{q}_j\|^2} \quad (13)$$

Here (\mathbf{q}_j, e_j) are the points in the image, for which the error was measured; h is the bandwidth parameter, that is optimized using grid search. The regression result is shown in Figure 4(a).

However, obtaining ground truth measurements for each pixel is not a trivial task, which cannot be performed outside of a lab setup. We propose to avoid it by firstly calibrating extrinsic and intrinsic parameters of the camera without taking into account the depth bias and then using depth, predicted from the calibration, to determine $\hat{e}(d_r, x, y)$. Since d_r is not available at run-time, we estimate the bias as $\hat{e}(d_m, x, y)$, using $e_j = \hat{d}^{il} - d^{il}$, i.e. the difference between depth measurement d^{il} at pixel $\mathbf{p}^{il} = (x^{il}, y^{il})^T$ and the estimated depth \hat{d}^{il} and $\mathbf{q}_j = (\hat{d}^{il}, x^{il}, y^{il})^T$.

To predict the expected depth value at a given pixel, we use the current estimates of the ToF camera parameters and render a plane onto each image. From the plane pose relative to the camera $R_{tof}^i, \mathbf{t}_{tof}^i$, the plane parameterization with normal \mathbf{n}_{tof}^i and offset θ_{tof}^i can be computed. Then, the predicted depth value at pixel l , \mathbf{p}^{il} , is computed as:

$$\hat{d}^{il} = \frac{\theta_{tof}^i}{n_{tof,1}^i x_n^{il} + n_{tof,2}^i y_n^{il} + n_{tof,3}^i}, \quad (14)$$

where x_n^{il}, y_n^{il} are computed from \mathbf{p}^{il} by first inverting (2) and then undistorting the normalized coordinates.

To be able to estimate the depth error reliably, however, it is required either to use a three-color pattern (see Figure 3(a)), or an additional space on the calibration board without any pattern on it. This is because the black checks,

while visible on the amplitude image, produce unreliable depth data on the depth image, which is not usable for depth correction (see Figure 3(b),3(c)). The area, free of black-and-white squares, can be instead used to compare the measured distance and the distance predicted using estimated parameters of the camera and the plane pose.

The blue-and-white part of the pattern can be used to determine re-projection error from ToF image to color camera, for example, for calibration evaluation.

3.5 Joint optimization

After the initial guess for the parameters is obtained, an iterative process is performed:

1. The parameters of the cameras $K_{tof}, \mathbf{k}_{tof}, \{R_{tof}^i, \mathbf{t}_{tof}^i\}_i$, and $K_{rgb}, \mathbf{k}_{rgb}, \{R_{rgb}^i, \mathbf{t}_{rgb}^i\}_i, R_{\Delta}, \mathbf{t}_{\Delta}$ in case of presence of a color camera, are jointly optimized.
2. Depth bias is estimated and corrected, as described in Section 3.4.

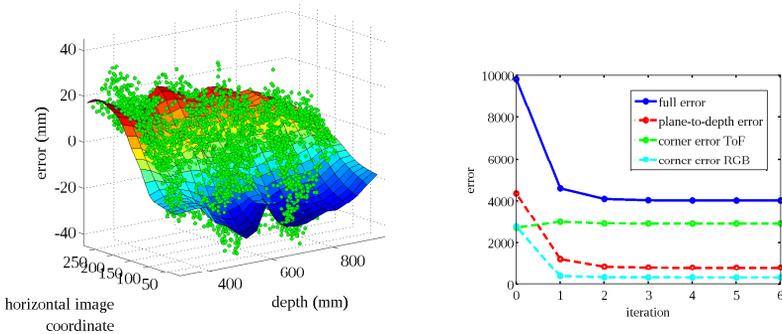


Fig. 4. 4(a)Depth error e measured relative to d_m and the x coordinate of the image (grey dots); fitted regression surface $e(d_m, x)$. 4(b) Decrease of E_{full} depending on the iteration, as well as decrease in each of the sums of the error of E_{full} : depth-to-plane error (red), corner error for the amplitude images (green) and corner error on the rgb images (azure); as can be seen, 5 iterations are enough so that the algorithm converges.

The idea behind joint optimization of the parameters is the following. As mentioned above, corner-based calibration suffers from the low resolution of the ToF images and plane-based calibration from [3] gets confused by the systematic error and noise in depth measurements. We therefore fuse both approaches to compensate for their drawbacks.

Joint optimization of the parameters is performed by minimizing a corner-based error term and a plane-based error term together:

$$E_a = \sum_i \left(\underbrace{\sum_k \frac{\|\mathbf{p}_a^{ik} - \hat{\mathbf{p}}_a^{ik}\|^2}{\sigma_a^2}}_{\text{corner error ToF}} + \underbrace{\sum_l \frac{(d^{il} - \hat{d}^{il})^2}{\sigma_d^2}}_{\text{plane-to-depth error}} \right) \quad (15)$$

Here, \mathbf{p}_a^{ik} is the k -th checkerboard corner detected on image i , and $\hat{\mathbf{p}}_a^{ik}$ is the projection of the model's k -th corner using ToF camera distortion parameters and its projection matrix, \hat{d}^{il} is depth prediction computed using (14) and d^{il} is the depth measurement delivered by the ToF camera. Each error term is normalized by the variance in the corresponding error: σ_a^2 and σ_d^2 , computed directly after the initialization step.

Plane-to-depth error estimation follows the idea described in [3]. Unlike [3], we use the same distortion model for the ToF camera as for the color camera, and not the reverse one.

In case of presence of a color camera, a new term is added to E_a :

$$E_{full} = \sum_i \left(\sum_k \frac{\|\mathbf{p}_a^{ik} - \hat{\mathbf{p}}_a^{ik}\|^2}{\sigma_a^2} + \underbrace{\sum_k \frac{\|\mathbf{p}_{rgb}^{ik} - \hat{\mathbf{p}}_{rgb}^{ik}\|^2}{\sigma_{rgb}^2}}_{\text{corner error RGB}} + \sum_l \frac{(d^{il} - \hat{d}^{il})^2}{\sigma_d^2} \right), \quad (16)$$

where \mathbf{p}_{rgb}^{ik} is the k -th checkerboard corner detected in the corresponding color image i and $\hat{\mathbf{p}}_{rgb}^{ik}$ is the projection of the k -th corner into the color image using the relative checkerboard position from equations (10)-(11).

We show in Figure 4(b), that the error decreases when iterating between joint optimization and error regression, and in our experiments became stable after around 5 iterations. Note the slight increase of the corner error for the amplitude images: we attribute it to the fact that corner detection on the amplitude images is not very precise due to their low resolution.

4 Evaluation

To evaluate the presented calibration approach, we performed a series of experiments with Intel Creative Gesture Camera. The setup consists of both color and ToF sensors, the resolution of the ToF sensor is 320×240 and the admissible distance is from 100mm to 1000mm. In practice, however, the distances below 200mm and above ≈ 950 mm are estimated very unreliably. The resolution of a color camera is 640×480 pixels.

To obtain the calibration, we recorded 25 images of the calibration board. For calibration quality it is important to cover most of the calibrated volume and to ensure that the poses are uniformly distributed in the volume, otherwise the calibration will be biased.

We evaluate the resulting calibration in two experiments. Firstly, we reconstruct 3D planes and compute the relative distance and the angular error between them.

Secondly, we evaluate the alignment between depth and color sensors by re-mapping depth images to RGB images using the obtained relative calibration. Note, that although the SDK [4], available for the camera, provided the capability to map a depth image onto a color image, the actual intrinsic parameters, relative pose and depth bias are not available. That makes it difficult to include additional sensors without completely re-calibrating the whole system.

The camera intrinsic parameters are presented in Table 1 and compared to the values provided by the manufacturer.

Table 1. Intel Creative Gesture Camera intrinsic parameter estimation.

	estimated	manufacturer
(f_x, f_y)	(231.09, 231.16)	(224.50, 230.49)
(c_x, c_y)	(150.87, 118.22)	(160, 120)
\mathbf{k}	(-0.14, -0.03, 0.00, 0.00, 0.23)	-

4.1 Reconstruction quality

In this experiment, we measure the relative distance error and angular error between reconstructed planes to evaluate the quality of the calibration, and then compare it with the manufacturer’s calibration.

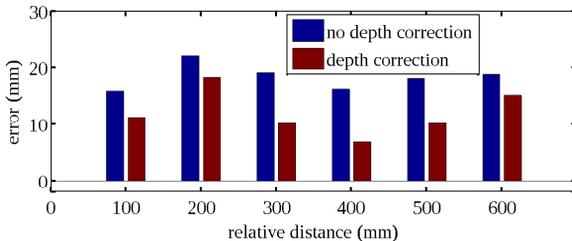


Fig. 5. Error in relative distance measurements with and without correction.

Relative distance error The distance between parallel planes is measured and compared with the ground truth value. In Figure 5, the average distance error is compared for several planes. As reference plane, a plane at the distance of 300mm from the camera was taken. As it can be seen, the correction allows to reduce the measurement error almost by a factor of 2 with respect to depth data without correction.

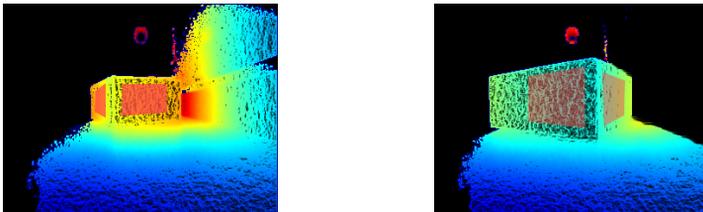


Fig. 6. Images of the two perpendicular planes, used for angle estimation; red polygons show the areas used to estimate plane normals.

Angle error To measure the angle error, we recorded two perpendicular planes. Afterwards, they were reconstructed, and the angle between reconstructed planes was estimated as the angle between the normals to the planes. Results are presented in Table 2. They show that depth correction allows to improve the results in both experiments. Our calibration also allows to reduce angular error between planes. The relatively large deviation from 90 degrees can be explained by the fact, that most of the images contain planes, where one plane is almost perpendicular to the image plane (see Fig 6) or too close to the image boundary; this causes bigger normal estimation errors, for example, due to noise. In general, we observed, that if the planes are placed in the middle of the scene, then the error for our calibration is less then 0.2 degrees, however, the closer they are moved to the side, the bigger the error gets.

Table 2. Estimation of the angle between two perpendicular planes. Comparison between manufacturer calibration, calibration without depth correction and calibration with depth correction.

	mean (degrees)	std. deviation (degrees)
manuf. calib.	83.42	3.72
calib. no. corr.	83.61	4.37
calib. with corr.	85.08	2.88

4.2 RGB-ToF alignment

Table 3. Comparison between depth-color image alignment provided by manufacturer, our calibration without depth correction and with depth correction.

	mean (px)	std. deviation (px)
manuf. calib.	2.6979	1.4864
calib. no. corr.	0.8642	0.5638
calib. with corr.	0.8118	0.4391

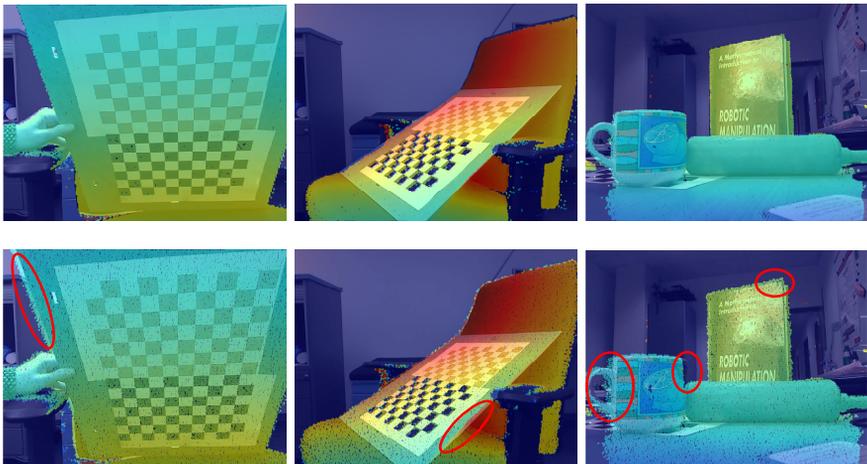


Fig. 7. Examples of the alignment between depth and RGB data for Intel Creative Gesture Camera, obtained using our calibration (top row) and the manufacturer calibration (bottom row); manufacturer calibration often does not give a perfect alignment between depth and rgb image.

In this experiment, we evaluate the improvement in the relative pose estimation between depth and RGB camera.

In Table 3, the corner re-projection error is compared for the manufacturer-provided depth alignment, our relative calibration without depth correction and with depth correction. It can be seen that the alignment provided by our calibration is significantly better than the mapping provided by the manufacturer.

The results can also be confirmed visually. In Figure 7 the alignment achieved by our calibration is compared to the manufacturer-provided mapping. It can be seen, that the borders are preserved much better than in the case of manufacturer mapping.

5 Conclusion

In this work, we presented a calibration approach for a ToF camera, that can be used also to simultaneously calibrate it with a color camera. The proposed method does not require ground truth data to correct ToF camera depth bias and only uses 20 – 30 images of a calibration board. We evaluated our approach for Intel Creative Gesture Camera in two experiments and compared the results with the manufacturer calibration, thus showing, that we are able to obtain more precise calibration results and provide explicit camera model. Due to the non-parametric nature of the depth bias representation, the approach can be used to calibrate other ToF cameras without additional effort. It also can be easily extended by adding additional sensors to the system.

References

1. Cui, Y., Schuon, S., Chan, D., Thrun, S., Theobalt, C.: 3d shape scanning with a time-of-flight camera. In: CVPR. pp. 1173–1180 (2010)
2. Hansard, M.E., Evangelidis, G., Horaud, R.: Cross-calibration of time-of-flight and colour cameras. CoRR abs/1401.8092 (2014)
3. Herrera, D., Kannala, J., Heikkila, J.: Joint depth and color camera calibration with distortion correction. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 34(10), 2058–2064 (2012)
4. Intel: Intel creative gesture camera, <http://click.intel.com/creative-interactive-gesture-camera-developer-kit.html>
5. Kim, Y.M., Chan, D., Theobalt, C., Thrun, S.: Design and calibration of a multi-view tof sensor fusion system. In: *IEEE CVPR Workshop on Time-of-flight Computer Vision*. pp. 1–7. IEEE, Anchorage, USA (2008)
6. Lefloch, D., Nair, R., Lenzen, F., Schäfer, H., Streeter, L., Cree, M.J., Koch, R., Kolb, A.: Technical foundation and calibration methods for time-of-flight cameras. In: Grzegorzec, M., Theobalt, C., Koch, R., Kolb, A. (eds.) *Time-of-Flight Imaging: Algorithms, Sensors and Applications*. LNCS, vol. 8200. Springer (2013)
7. Lindner, M., Kolb, A.: Lateral and depth calibration of pmd-distance sensors. In: *In ISVC* (2. pp. 524–533. Springer (2006)
8. Lindner, M., Schiller, I., Kolb, A., Koch, R.: Time-of-flight sensor calibration for accurate range sensing. *Comput. Vis. Image Underst.* 114(12), 1318–1328 (Dec 2010), <http://dx.doi.org/10.1016/j.cviu.2009.11.002>
9. Nadaraya, E.A.: On estimating regression. *Theory of Probability and its Applications* 9, 141–142 (1964)
10. Schnemann, P.: A generalized solution of the orthogonal procrustes problem. *Psychometrika* 31(1), 1–10 (1966), <http://EconPapers.repec.org/RePEc:spr:psycho:v:31:y:1966:i:1:p:1-10>
11. Unnikrishnan, R., Hebert, M.: Fast extrinsic calibration of a laser rangefinder to a camera. Tech. Rep. CMU-RI-TR-05-09, Robotics Institute, Pittsburgh, PA (July 2005)
12. Zhang, Z.: A flexible new technique for camera calibration. *IEEE Trans. Pattern Anal. Mach. Intell.* 22(11), 1330–1334 (Nov 2000), <http://dx.doi.org/10.1109/34.888718>