

ACCURACY OF EXTERIOR ORIENTATION FOR A RANGE CAMERA

Jan Boehm, Timothy Pattinson

Institute for Photogrammetry, University of Stuttgart, Germany
jan.boehm@ifp.uni-stuttgart.de

Commission V, WG V/2, V/4, SS-2 (range cameras)

KEY WORDS: Range Camera, Time-of-Flight, Orientation, Point Cloud, Alignment

ABSTRACT

While known in principle for over 10 years range cameras have recently received increased attention. As the technology matures, interest is given to the possible applications of such sensors. Clearly the applicability of the sensor is connected to its reliability and accuracy. At the centre of this study lies the question how accurately the exterior orientation of a range camera can be estimated, given a reference point cloud. Before we can assess the accuracy of a quantity derived from processing sensor readings, we have to examine the accuracy of the sensor itself. This leads to an investigation of range camera calibration. In this study we present our results on calibrating a PMD sensor and applying the correction data to the problem of computing the exterior orientation of a range camera. We present results derived from practical experiments and discuss the influence of different correction approaches on orientation through point cloud alignment.

1. INTRODUCTION

Knowing the precise exterior orientation of a sensor is valuable in many different aspects. In this work we aim to know the precision of exterior orientation for the purpose of localization. Localization itself can have wide-spread applications from pedestrian navigation to autonomous vehicles. Therefore it is studied in a wide spectrum of scientific disciplines. From the many approaches suggested to solve this problem, we investigate the use of range cameras. Range cameras have the interesting property, that they simultaneously provide an intensity image and a range map of the scene. The intensity image can be used in existing video-based frameworks. The range map offers additional possibilities, which can either replace or complement intensity based- approaches.

In the neighbouring disciplines the problem is closely related to the problem of determining position and orientation, or orientation in short, from the information given by an imaging sensor, i.e. a camera. The most prominent approaches today are for one structure from motion (SfM) (Hartley, 1993; Koenderink and Vandoorn, 1991; Tomasi and Kanade, 1993), which puts the emphasis on the problem of deriving a representation of the scene and the actual localization is more of a by-product. For continuous online operation over large maps and many observations, simultaneous localization and mapping (SLAM) (Smith and Cheeseman, 1986), or more specific visual SLAM (VSLAM), was introduced, which emphasizes more the localization problem. With the use of 2D laser range finders, the problem was transformed to three degrees of freedom (DOF) for planar movement (Lu and Milios, 1997). Using full 3D sensors this approach was also extended to six DOF (Borrmann et al., 2008).

While known for over 10 years, range cameras have recently gained high interest in the community. Range cameras, also known as Time-of-Flight (TOF) cameras or Focal-Plane-Array LiDAR, provide a dense measurement of range and direction (and thereby a point cloud) at video rate. This class of sensors allows for the determination of the current position of the sensor by aligning the measured point cloud to a reference point cloud.

The alignment is computed by the well-known iterative closest point (ICP) algorithm (Besl and McKay, 1992). This registration method has shown to deliver very precise results in many applications. The algorithm's constraint of requiring an initial approximation of the alignment is acceptable in the proposed application framework. For a dense (in time) sequence of range images, we can use the result of the previous frame as initialization for the consecutive frame.

The approach of locating a sensor using a reference point cloud builds on a previous effort to locate an intensity camera in a point-based environment model (PEM) (Böhm, 2007). The motivation for the PEM approach is based on the expectation, that dense point clouds of large building complexes, industrial facilities and urban areas will become widely available in the next few years. The main drive behind the acquisition of such data is from the computer-aided facility management (CAFM) or building information model (BIM) industry, preservation authorities and safety agencies. The main enabling factor is the recent wide spread availability of reliable sensor systems and service companies. Once this data has been acquired it can serve multiple purposes, our proposed application being only one among many. The PEM can easily be acquired by up-to-date laser scanning systems. In contrast to model-based approaches it avoids any expensive modelling step and directly operates on the raw point data.

To implement a localization system based on range cameras it is of obvious importance to assess the accuracy one can achieve in determining the exterior orientation of a range camera by point cloud alignment. This study differs from other reports about accuracy of range cameras, as it does not only try to examine the accuracy of the point cloud measured by the range camera itself, or the accuracy of a single point within the point cloud, but it assesses the accuracy of the information derived from the point cloud. In this aspect we believe the study is closer to a real world application and thus produces more realistic accuracy values.

Obviously the accuracy of the point cloud will influence the accuracy of the alignment and therefore the exterior orientation.




Figure 1: Range camera on tripod facing the test field.

Therefore the calibration of the range camera has an influence on the accuracy of the exterior calibration. We demonstrate the magnitude of the effect by comparing uncalibrated data, vendor-calibrated data and data calibrated with our own method. Range camera calibration has recently attracted some attention. Several suggestions for calibration strategies and range correction terms were made (Kahlmann, 2008; Lindner and Kolb, 2006). Our newly developed calibration strategy combines a photogrammetric calibration with a range correction term. The range correction term contains sinusoidal elements similar to (Lichti, 2008). It is currently able to reduce the RMS on known control points below 13 mm. However, calibration is not the only influence on the accuracy of localization, and we wish to assess the influence calibration has by practical experiments.

The data used throughout this study is generated with a PMD CamCube 2.0, the latest model in the line of PMD sensors (Schwarte et al., 1997). It operates at a resolution of 204 by 204 pixels with a 40 degree viewing angle at a maximum distance (ambiguity interval) of 7.5 meters. The reference surface data is acquired with a Leica HDS 3000; a well-established lasers scanner with accuracy below 5 mm. Ground truth for the exterior orientation is computed with photogrammetric methods from the intensity images provided by the range camera over known control points.

Parameter	Final Value	Final Std. Error
C (mm)	12.7915	3.93E-03
XP (mm)	-0.1361	3.28E-03
YP (mm)	-0.1504	3.23E-03
K1	2.70E-03	2.85E-05
K2	6.63E-06	1.62E-06
K3	1.69E-07	3.19E-08
P1	-6.02E-05	8.66E-06
P2	9.08E-05	8.59E-06
B1	1.84E-05	4.37E-05
B2	-2.73E-05	4.54E-05

Table 1: Results of photogrammetric calibration using a standard 10 parameter close-range model (Brown, 1971). All parameters were estimated significantly, except affinity and shear

2. PHOTOGRAMMETRIC CALIBRATION

Calibration of a digital camera is standard in photogrammetry (Fraser, 1997) and is easily applied to range cameras. Special care has to be taken for the design of the test field due to the low resolution of range cameras. We established a (nominally) planar test field, with a 7-by-7 grid of large, circular targets. This test field is specially designed for the spatial resolution of the PMD camera used. The white circles on the targets were selected so that they would be represented by several pixels on these camera's images, with small black dots at each centre to allow precise positioning by higher-resolution equipment. The positions of the centres of the 49 targets were measured precisely in three dimensions using a combination of high-resolution optical photogrammetry (using a digital SLR camera) and scanning using a total station. The optical imagery provides high-accuracy relative positioning, and the scanning was used to impose the correct linear scale on the coordinate system.

Figure 1 shows the camera mounted on a tripod facing the test field. Each of the 49 white circles on the plate presents a large, highly-reflective target for the camera. Note that due to the design of the test field, the distance measurements are all made within a uniform, reflective surface. This avoids some of the edge effects encountered with checkerboard patterns, where the distance measurements are taken at the boundary between dark and light sections.

In Table 1 we show the results of bundle adjustment, where all parameters except affinity and shear are significant. Performing bundle adjustment has the advantageous side-effect, that not only internal camera parameters are computed but also external parameters, namely camera stations. We use the external orientations as input for range calibration.

3. RANGE CALIBRATION

By taking the known target coordinates and the position of the camera for each image derived from bundle adjustment, the geometrical distances to each of the targets can be determined. By comparing these geometrical distances with the measured distances from the camera's range image, the errors in range measurement can be determined and evaluated. Figure 2 shows in the top image the desired linear relationship of the photogrammetrically obtained distances and the distances measured by the range camera directly. Further details become visible when the difference of the two is plotted over distance. Almost all the errors are greater than zero, which means that the distance measurements from the camera are generally slightly higher than the photogrammetrically calculated distances. The maximum error is on the order of 12 cm, with most of them below 8 cm. Errors are generally larger for larger distances. There is a marked wave-like structure in the error data, which does not appear to be random. The errors are clustered fairly well in the vertical direction, which implies a large systematic component. Therefore, calibration should be successful in reducing the size of the errors.

The data was analysed by fitting models of varying complexity. An ideal model would match the data accurately leaving a small residual error, while requiring a small number of parameters. Each model was applied to data from each of the four available modulation frequencies used by this camera (18 MHz, 19 MHz, 20 MHz and 21 MHz). After carefully analyzing several possible correction models (linear, polynomials of varying order, look-up table), we establish a model which best fits the data while requiring few parameters. We obtain the following

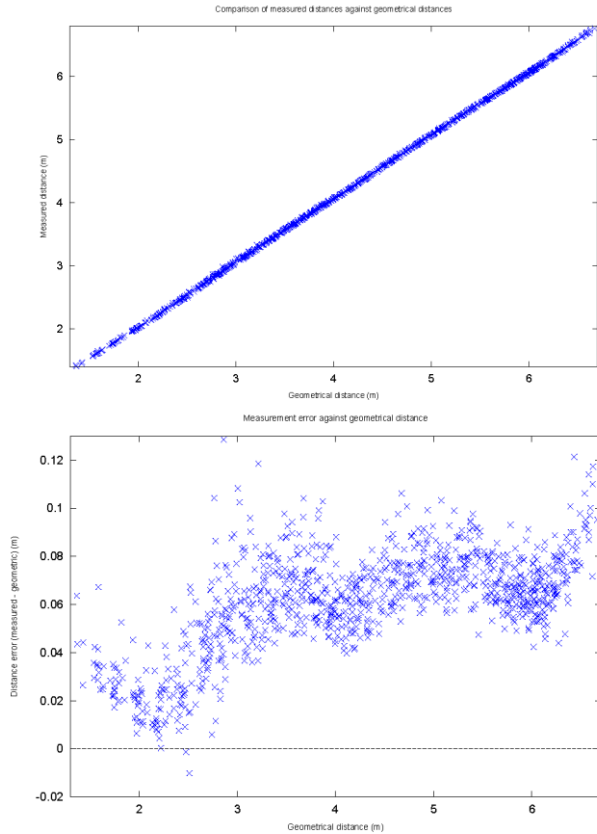


Figure 2: Comparison of photogrammetrically obtained distances and measured range (top). Difference of the two plotted as a function of distance (bottom). Units are in meters.

seven-parameter correction term for the calibration, using the measured distance d and the radial distance in the image plane r :

$$\Delta d = a_0 + a_1 d + a_2 \cos(4kd) + a_3 \sin(4kd) + a_4 \cos(8kd) + a_5 \sin(8kd) + a_6 r$$

where k is the wavenumber related to the camera's modulation frequency f_{mod} :

$$k = \frac{2\pi}{\lambda_{\text{mod}}} = \frac{2\pi f_{\text{mod}}}{c}$$

The meaning of the 7 calibration parameters can be summarised as follows:

- a_0 : Fixed offset term, corresponding to a shift of the measurement origin.
- a_1 : Linear term in distance, corresponding to a scale factor, perhaps due to a shift in modulation frequency.
- a_2, a_3 : Amplitudes of sinusoidal functions at 4 times the modulation frequency. This can occur due to systematic shifts in sensor readings.
- a_4, a_5 : Amplitudes of sinusoidal functions at 8 times the modulation frequency. This can occur due to harmonic components at frequencies $3 f_{\text{mod}}$ or $5 f_{\text{mod}}$ present in the impure modulation signal.
- a_6 : Linear factor of the radial distance. This can be caused by systematic effects at high viewing angles. (Lichti, 2008) attributes linear in-plane parameters to a clock skew error.




Figure 3: Result for fitting the proposed range correction model to the data for 20 MHz. Units are in meters.

Figure 3 shows the result after fitting the proposed model to the range data at 20 MHz modulation frequency. We can clearly see how the systematic pattern is reduced and overall amplitude of the errors is minimized. RMS of residuals is 10.5 mm on the calibration data and 12.4 mm on an independent data set. In Table 2 we compare our proposed 7-parameter sinusoidal model to other candidate models. We see the model outperforms all other models while it still requires a small set of parameters. We hold the number of parameters to be an important factor to the stability of range calibration.

In Table 3 we show the effect of the range correction at different modulation frequencies. This time the correction parameters are applied to an independent dataset. It can be seen that the calibration is indeed successful and the RMS residuals are greatly reduced by applying the given corrections. The residuals here are slightly higher than those fitted directly to the calibration data, as we may expect. But still, the maximum RMS residual of 12.4 mm on completely independent data is a very satisfactory result and a reassuring verification that these calibration parameters are generally valid, also for independent measurements.

4. EXTERIOR ORIENTATION

The experiments compare the performance of three different calibration and correction modes:

- uncorrected (UNC)
- vendor lens correction (VLC)
- our suggested photogrammetric + range calibration (PRC)

Method	Number of parameters	RMS residual (mm)
original data		64.7
fixed offset	1	20.4
offset and scale	2	16.0
quadratic	3	14.6
polynomial	5	13.8
polynomial	7	12.7
polynomial	10	12.2
lookup table	50	12.1
sinusoidal	5	14.0
sinusoidal	7	10.5

Table 2: Comparing our proposed 7-parameter sinusoidal model to other candidate models.

Modulation frequency	RMS residual (mm) without correction	RMS residual (mm) with correction
18 MHz	47	12.0
19 MHz	62	11.8
20 MHz	60	12.4
21 MHz	62	11.4

Table 3: Effect of applying range correction term to an independent dataset over different modulation frequencies.

Uncorrected data is straight forward to be understood as the raw data from the sensor. This does not imply the vendor does not incorporate calibration as part of the production pipeline. The vendor supplied lens correction can be enabled or disabled during data acquisition. The model used for lens correction is not released to the public. We therefore treat it as a black-box procedure. From purely visual impression we observe, that the model seems to successfully correct the strong radial lens distortions present in the sensors images.

For the comparison we use range camera measurements of a realistic test scene. The test scene is realistic in the sense that it does not consist of a planar test field using targets of optimal reflectivity, as we have used for the calibration. The scene rather contains different materials of varying micro structure and reflectivity. An intensity image of the test scene is shown in Figure 4. The range camera captures data at three different positions along a trajectory traversing a corner. Images are captured in static mode, i.e. the camera is mounted on a tripod. This is necessary to enable the switch between the vendor supplied correction and uncorrected mode and still keep the exact same position. Correction terms from our suggested calibration model are applied offline to the uncorrected or uncalibrated data set.

The reference point cloud from the terrestrial laser scanner and the point cloud from the range camera are imported to a commercial registration software (Polyworks by InnovMetric). Each of the range images is aligned independently to the reference point cloud. The standard deviation of the alignment error for each station and for each calibration mode is listed in Table 4. We can confirm from the table, that our correction term yields the best alignment error for each station. Albeit the difference to the uncorrected case is small (improvement of 10% on average). The vendor supplied calibration does not improve the alignment error over the uncorrected mode in each case. The histograms of the alignment errors for the second station for all three correction modes are shown in Figure 6. The graph to the right has the narrowest shape and thus the least points which deviate from the reference surface. This supports our conclusion, that our correction terms yield a point cloud with a shape which fits the reference data the best of all three modes.

Reference markers were added to the test scene and measured with the laser scanner. These points act as control points for photogrammetric resection of the range camera. The exterior orientation obtained by photogrammetric measurements serves as a reference for the assessment of the exterior orientation obtained by point cloud alignment. Table 5 shows the average deviations over all three stations. The coordinate system is defined by the laser scanner and the z-axis is in the vertical direction. Again we can confirm that our correction term yields the best result for each station. However, the improvement over the uncorrected case is again small (20%).

In Figure 7 we show the distributions of the computed stations. We can clearly see that the three correction methods form clusters, which are systematically offset from the wall. Looking at the cluster of the first station in detail, we observe that the PRC method is closest to the true station.



Figure 4: Intensity image of the test scene as captured by the range camera on the third station.




Figure 5: Example of range camera data aligned with laser scan data using ICP approach.

	UNC	VLC	PRC
S1	0.018313	0.021304	0.017972
S2	0.021572	0.026497	0.019505
S3	0.019025	0.018945	0.015322

Table 4: Standard deviation of alignment error after ICP registration for three stations (rows) and three different calibration modes (columns). Units are in meters.




Figure 6: Histogram of the alignment errors at the second station. The calibration mode from left to right is UNC, VLC and PRC. The centre line indicates the mean value, the outer lines indicate the standard deviation. Each graph is scaled on the horizontal axis from/to ± 0.2 meters and from 0 to max on the vertical axis.

	UNC	VLC	PRC
X	-0.27	-0.194	-0.189
Y	0.184	0.265	0.163
Z	0.064	0.109	0.089
Distance	0.333	0.346	0.265

Table 5: Difference of Exterior Orientation obtained by point cloud alignment compared to photogrammetric resection over all stations for three different calibration modes. Units are in meters.

5. DISCUSSION

The experiments described above clearly show, that the determination of the exterior orientation from the point cloud of a range camera is possible, given a reference point cloud which has sufficient geometric properties. However, while calibration on a carefully designed test field shows promising range accuracies, these cannot be transferred to an arbitrary scene. In our experiments we report deviations in position of the exterior orientation in the range of 30 cm. This is an order of magnitude above the range accuracies achieved in calibration.

While careful calibration does show to improve the results of exterior orientation computation, other effects overshadow the influence of calibration. In agreement with other investigations (May et al., 2009) we see an influence from geometric distortions of the point cloud, which is not due to a lack of calibration of the sensor, but is dependent on the scene itself. We attribute these distortions to the multi-path effect in phase computation due to light-scattering. While calibration itself works satisfactory, the true bottle-neck in real-world applications is scene-dependant distortions.

In a critical review of the experiments, we note that the reference data for the exterior orientation would better not be established with the range camera itself. The range camera has a very low resolution and a noisy intensity channel, which might make image measurements inaccurate and thus affects backward resection. Furthermore a higher degree of automation in the evaluation process would enable us to process more stations and possibly different configurations to further substantiate the findings.

The impact of studies on accuracy of exterior orientation goes beyond the applications in positioning. Also applications aiming at modelling scenes from range camera depend on reliable exterior orientation. Therefore new findings should have a broader impact on the applicability of range cameras.

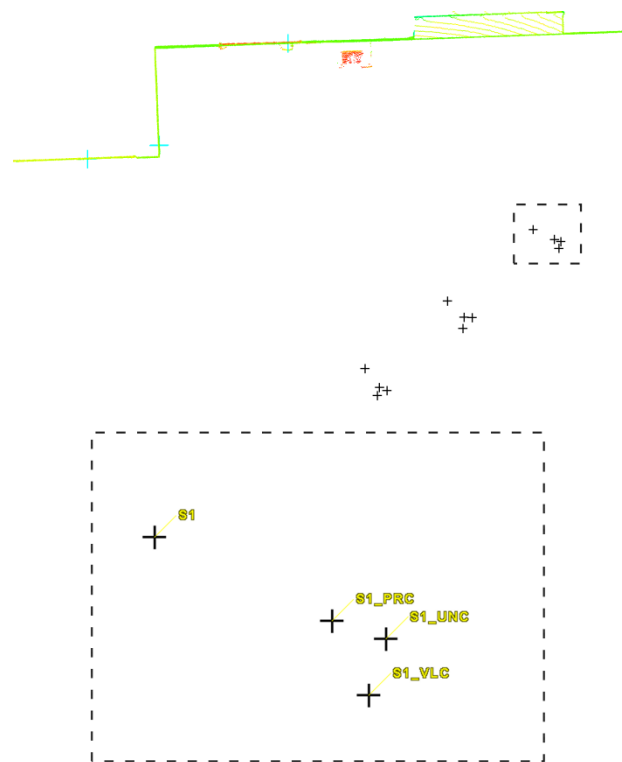


Figure 7: The three stations chosen in relation to the room corner. Each station is shown as a cluster of four vertices representing ground truth and the three correction methods (top). The detail of one cluster shows how all three methods are systematically set back from ground truth, while PRC is closest (bottom).

6. ACKNOWLEDGEMENTS

Grateful thanks go to Konrad Wenzel, who prepared the test field used in this study especially for this camera. We also wish to thank M. Profittlich and T. Ringbeck from PMDTechnologies GmbH for their support in interfacing the PMD device.

7. REFERENCES

- Besl, P.J. and McKay, N.D., 1992. A Method for Registration of 3-D Shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2): 239-256.
- Böhm, J., 2007. Orientation of Image Sequences in a Point-based Environment Model, Sixth International Conference on 3-D Digital Imaging and Modeling (IEEE 3DIM 2007), pp. 233-240.
- Borrmann, D., Elseberg, J., Lingemann, K., Nuchter, A. and Hertzberg, J., 2008. Globally consistent 3D mapping with scan matching. *Robotics and Autonomous Systems*, 56(2): 130-142.
- Brown, D.C., 1971. Close-Range Camera Calibration. *Photogrammetric Engineering*, 37(8): 855-866.
- Fraser, C.S., 1997. Digital Camera Self-Calibration. *ISPRS Journal of Photogrammetry and Remote Sensing*, 52: 149-159.
- Hartley, R.I., 1993. Euclidean reconstruction from uncalibrated views, *Applications of Invariance in Computer Vision*. Springer-Verlag, pp. 237--256.
- Kahlmann, T., 2008. Range imaging metrology: Investigation, calibration and development ETH Zurich.
- Koenderink, J.J. and Vandoor, A.J., 1991. AFFINE STRUCTURE FROM MOTION. *Journal of the Optical Society of America a-Optics Image Science and Vision*, 8(2): 377-385.
- Lichti, D.D., 2008. Self-Calibration Of A 3D Range Camera. IAPRS, Beijing.
- Lindner, M. and Kolb, A., 2006. Lateral and Depth Calibration of PMD-Distance Sensors, *Advances in Visual Computing*, pp. 524-533.
- Lu, F. and Milios, E., 1997. Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, 4(4): 333-349.
- May, S. et al., 2009. Three-Dimensional Mapping with Time-of-Flight Cameras. *Journal of Field Robotics*, 26(11-12): 934-965.
- Schwarte, R. et al., 1997. New electro-optical mixing and correlating sensor: facilities and applications of the photonic mixer device (PMD). In: L. Otmar (Editor). *SPIE*, pp. 245-253.
- Smith, R.C. and Cheeseman, P., 1986. ON THE REPRESENTATION AND ESTIMATION OF SPATIAL UNCERTAINTY. *International Journal of Robotics Research*, 5(4): 56-68.
- Tomasi, C. and Kanade, T., 1993. Shape And Motion From Image Streams - A Factorization Method. *Proceedings of the National Academy of Sciences of the United States of America*, 90(21): 9795-9802.